

THÈSE DE DOCTORAT

Spécialité: Neurosciences cognitives
École doctorale Informatique, Télécommunications et Électronique (Paris)

Présentée par

Éléonore DUVELLE

Pour obtenir le grade de

DOCTEUR de l'UNIVERSITÉ PIERRE ET MARIE CURIE

RÔLE DE L'HIPPOCAMPE DANS LA REPRÉSENTATION DU BUT

-
Approche Comportementale et Électrophysiologique

Soutenue le 10 décembre 2014 devant le jury composé de:

Rapporteurs:	Jean-Christophe CASSEL Sidney WIENER	Professeur, Université de Strasbourg Directeur de Recherche, CNRS
Examineurs:	Philippe FAURE Kathryn JEFFERY	Directeur de Recherche, CNRS Professeur, University College London
Invité:	Bruno POU CET	Directeur de Recherche, CNRS
Directeurs de thèse:	Angelo ARLEO Étienne SAVE	Directeur de Recherche, CNRS Directeur de Recherche, CNRS

ROLE OF THE HIPPOCAMPUS IN GOAL REPRESENTATION

-
Insights from behavioural and electrophysiological approaches

This doctoral work was conducted under a co-direction involving two different research teams :

In Marseille, under the co-direction of Etienne SAVE,

leader of the **Neural Bases of Spatial Cognition** team,
in the Cognitive Neuroscience Lab (LNC),
UMR 7291, CNRS, Aix-Marseille Université.

In Paris, under the co-direction of Angelo ARLEO,

leader of the **Adaptive NeuroComputation** team
(ANC) until the end of 2013,
in the Neurobiology of Adaptive Processes Lab, UMR
7102, CNRS, UPMC

and leader of the **Aging in Vision and Action** team
(AVA) since early 2014,
in the Institute of Vision,
UMR 7201, CNRS, INSERM, UPMC.

*À mon père, Jacques Duvelle,
à sa mémoire et sa musique.*

Abstract

In humans, the hippocampus is a brain structure known for its role in episodic memory. In rats, this structure is mostly studied for its possible role in spatial cognition. Indeed, the hippocampus is thought to endow an animal with the ability to locate itself and learn the spatial relationships between relevant elements of the environment. The hippocampus is especially involved in flexible navigation towards a non directly visible goal. At the cellular level, the role of the hippocampus in spatial cognition is instantiated by the existence of 'place cells', neurons that fire when the animal is at a particular location in the environment, called the 'place field'. Place cells are proposed to be the neural support for a flexible representation of space which could underlie spatial memory. In the course of navigation, both spatial memory and decision-making processes are combined to direct behaviour, in particular, select a goal, plan a trajectory, and decide to stop upon arrival at the goal. The hippocampus is known to communicate anatomically and functionally with structures involved in decision-making, such as the ventral tegmental area, the prefrontal cortex, and the striatum. However, the nature of the interaction between these structures is still being elucidated. Recently, a possible link between the hippocampal spatial representation and decision-related processes was discovered, in the form of an extra-field discharge of CA1 place cells while a rat is waiting for reward delivery in an uncued goal zone. This 'goal-related' activity suggests that the hippocampus underlies a representation of goals. However, the nature of this representation is the subject of multiple hypotheses and its role remains unclear.

The objective of this thesis was to address the fundamental issue of the goal representation in the hippocampus, in order to better understand what could be the nature of the interplay between the spatial and the decision-making circuits of the brain. In particular, we asked whether the goal-related activity of place cells is modulated by the spatial or the motivational characteristics of the goal. We designed a task that allows a rat to freely choose between two spatial goals. An important aspect of this task is that the value of goals can be modified by changing the magnitude of the reward associated to them. We performed extracellular unit recordings of dorsal hippocampal neurons in rats trained in this dual-goal task. The behavioural results that we obtained showed that rats were not only able to remember the location of each goal, but that they also optimised their behaviour by flexibly adapting their choices to variations in goal values. The analysis of neural activity during this task demonstrated that hippocampal place cells from both CA1 and CA3 expressed a goal-related activity with a specific temporal profile. Surprisingly, this goal-related activity was also expressed by a large proportion of pyramidal 'silent' cells, which did not have a place field in the environment. The combined goal-related activity of both place and silent cells could mediate a different but complementary coding scheme than the one exhibited by the place fields of pyramidal cells. In addition, the results revealed that the main determinant of the goal-related firing is the spatial aspect of the goal, since it was predominantly expressed in one of the two goal zones, thus indicating spatial discrimination. Complementing these results, the population activity at the goal was independent from both the goal value and the behavioural choices of rats.

Overall, the present results allowed us to gain some insight into the possible contributions of the rat dorsal hippocampus to the representation of goals. We suggest that the hippocampus would extract relevant information about the goal, in particular, its spatial characteristics, and create a goal representation independent from motivational aspects. Such a representation might be used in the course of navigation to recognise the match between the stored representation and the current inputs, sending a confirmation signal upon arrival at the goal that could contribute to the decision to terminate navigation.

Résumé

L'hippocampe, chez l'Homme, est principalement connu pour son rôle dans la mémoire épisodique. Chez le Rat, il est surtout étudié pour son possible rôle dans la cognition spatiale. Il permettrait à l'animal de se localiser au sein d'un environnement et d'apprendre les relations spatiales entre les éléments de l'environnement utiles à la navigation. En particulier, l'hippocampe serait impliqué dans les situations nécessitant de naviguer de manière flexible vers un but non directement visible. Au niveau cellulaire, le rôle de l'hippocampe dans la cognition spatiale est illustré par l'existence des 'cellules de lieu', des neurones qui déchargent préférentiellement lorsque l'animal est à un endroit particulier de l'environnement. Cet endroit est appelé 'champ de lieu'. L'activité de population des cellules de lieu sous-tendrait une représentation flexible de l'espace, support neuronal de la mémoire spatiale. L'information mémorisée peut être utilisée par le système de prise de décision à différentes étapes de la navigation dirigée vers un but, en premier lieu, pour la sélection du but lui-même, mais aussi pour la planification de trajectoires permettant d'atteindre ce but ou encore pour prendre la décision de mettre fin au comportement de navigation lorsque le but est atteint. Des liens entre l'hippocampe, en tant que support de la mémoire spatiale, et les structures motivationnelles et décisionnelles, telles que l'aire tegmentale ventrale, le cortex préfrontal et le striatum, ont été mis en évidence. Cependant, l'interaction fonctionnelle entre ces aires dans le cadre de la navigation spatiale reste à élucider. Récemment, l'existence d'une activité extra-champ des cellules de lieu à l'endroit d'un but spatial a été mise en évidence dans la région CA1 de l'hippocampe. Cette 'activité liée au but' suggère que l'hippocampe serait impliqué dans la représentation des buts. Cependant, la nature et le rôle de cette activité sont encore méconnus.

Cette thèse avait pour but d'aller plus loin dans la compréhension des mécanismes sous-tendant la représentation du but par l'hippocampe, pour ainsi clarifier le degré d'implication de l'hippocampe dans la prise de décision. En particulier, nous avons posé la question de l'influence relative des facteurs spatiaux et motivationnels sur l'activité liée au but des cellules de lieu. Pour ce faire, nous avons mis au point une tâche au cours de laquelle un rat peut choisir entre deux buts spatiaux. L'une des caractéristiques de cette tâche est qu'elle permet de faire varier la quantité de récompense associée aux buts, et ainsi, indirectement, la valeur des buts. Nous avons enregistré l'activité unitaire extracellulaire de neurones de l'hippocampe dorsal chez des rats entraînés dans cette tâche. Les résultats comportementaux ont mis évidence le fait que les rats sont capables non seulement de mémoriser la position de deux buts non-indicés, mais aussi d'adapter leur comportement en répondant de manière flexible aux changements de valeur des buts. L'analyse de l'activité électrophysiologique recueillie au cours de la tâche montre que les cellules de lieux, non seulement de CA1, mais aussi de CA3, expriment une activité liée au but avec un profil temporel caractéristique. De manière surprenante, une population de cellules de l'hippocampe, habituellement silencieuse, s'est révélée être active au but avec les mêmes caractéristiques que les cellules de lieu. L'activité liée au but combinée de ces deux populations de cellules pourrait être le support d'une représentation du but dans l'hippocampe qui reposerait sur un niveau de codage de l'information différent, mais complémentaire, de celui exprimé par l'activité du champ de lieu des cellules pyramidales. Enfin, les résultats ont révélé la prépondérance spatiale du codage du but par l'hippocampe. Plus précisément, l'activité au but était majoritairement exprimée de manière spatio-sélective, indiquant une discrimination spatiale des buts. Par ailleurs, cette activité de population était indépendante des variations de valeur des buts, ainsi que des variations comportementales de préférences des rats.

Ces travaux ont permis de clarifier l'une des contributions possibles de l'hippocampe dorsal du rat dans la prise de décision au cours de la navigation spatiale, qui serait d'extraire les aspects spatiaux du but pour créer une représentation indépendante de variations motivationnelles. Cette représentation du but pourrait être utilisée en cours de navigation pour reconnaître la correspondance entre la location visée et actuelle. En cas de congruence entre ces deux informations, un signal serait émis, confirmant l'arrivée à l'endroit visé, qui pourrait contribuer à la décision de cesser la navigation.

Remerciements

Tout d'abord, je tiens à remercier les membres de mon jury de thèse d'avoir accepté de juger ce travail, pour leur patience, pour leur adaptabilité, pour leur disponibilité malgré des emplois du temps certainement très chargés et pour les discussions scientifiques à venir lors du 'D-day', sans aucun doute très enrichissantes. Merci en particulier à mes rapporteurs, Jean-Christophe Cassel et Sidney Wiener, qui ont lu et évalué le manuscrit (et dans les temps !) malgré les très courts délais que je leur ai imposés. Je suis aussi très reconnaissante à mes examinateurs, Kate Jeffery pour avoir traversé la Manche et Philippe Faure pour avoir traversé la Seine. Merci enfin à Bruno Poucet, non seulement pour sa participation du 10 décembre, mais aussi et surtout pour sa contribution précieuse à plusieurs étapes de ce projet.

D'une manière générale, je remercie chaleureusement mes deux chefs (apparemment on dit aussi 'directeurs de thèse'), Etienne et Angelo, pour beaucoup de choses. Merci pour leurs qualités humaines, pour leur encadrement à la fois rigoureux mais qui m'a aussi laissé suffisamment d'espace (en parlant d'espace, voir le reste de ce manuscrit), pour leur soutien dans les moments difficiles (ah, la rédaction de la thèse...), pour leur complémentarité sur certains aspects et leurs points communs sur d'autres. En particulier, on pourrait dire que j'ai eu beaucoup de chance d'avoir des chefs ayant tous deux une approche à la fois éthique et collaborative de la recherche et des rapports humains, qui se diffuse à l'atmosphère des équipes qu'ils dirigent.

Dans l'ordre plus ou moins chronologique, je remercie sincèrement les personnes listées ci-après (et aussi celles que je risque probablement d'oublier). Tout d'abord, ma maman, sans qui je ne serai pas là, à différents niveaux ! Merci de m'avoir supportée, dans les deux sens du terme... et merci à elle d'être ce pilier sur lequel je sais que je peux toujours me raccrocher. Merci à Serge pour son soutien, particulièrement en termes de boissons de qualité. Merci à Matthieu (mon frère) pour les bonbons et les Froot Loops qui m'ont servi de carburant pour la dernière ligne droite, les différentes musiques qu'il m'a fait découvrir, et les discussions où nous ne sommes pas toujours d'accord. Un très grand merci à mon grand-père Robert, qui au-delà de son soutien et de son affection, m'a permis de séjourner à Paris ponctuellement, ou plus longtemps.

Même si c'est un peu hors sujet (mais ici, j'ai le droit !), merci à Angelo de m'avoir accueillie dans son équipe lors de mes stages de master 1 et 2 et de m'avoir donné l'opportunité de découvrir les joies de la modélisation. Un peu plus dans le sujet, je remercie à nouveau Angelo car, sans lui, cette thèse n'aurait jamais commencé ; merci pour son encadrement tout au long de la thèse, pour sa confiance, sa rigueur, ses conseils pertinents, sa relecture du manuscrit, sa faculté de rassurer et d'aller à l'essentiel, sa vision (ça tombe bien !), sa compétence à convaincre les autres (en tout cas, moi) de son point de vue, et enfin, sa capacité à rassembler un environnement humain sympathique et de qualité. Merci pour tout !

Je remercie tout particulièrement Luca, sans lequel énormément de choses seraient différentes (en moins bien, si si, je t'assure !), dans la vie de beaucoup de gens en général, et dans

cette thèse en particulier : merci pour son soutien général et sa disponibilité, pour sa relecture attentive de pratiquement tout le manuscrit, pour son aide à différents moments de la thèse, figures, comptage de cellules, etc., et pour l'organisation à la fin. Merci à Jérémie d'être lui-même, pour sa présence rassurante, pour son soutien moral, pour ses délicieux *cheesecakes*, pour le tennis et le footing, pour son goût immodéré mais quelque peu dissimulé pour le ping pong (mais oui, c'est juste que tu n'oses pas te l'avouer), pour son aide très précieuse (et précise) sur les figures, et pour ses remarques sur un chapitre de la thèse. Merci à Denis pour ses conseils, ses remarques scientifiques toujours directes et pertinentes et pour l'ambiance qu'il met aux soirées. Un grand merci à Laureline pour les nombreuses discussions toujours constructives (scientifiques ou autres), pour son accueil à Grenoble, pour son aide précieuse et efficace en statistiques et en idées d'analyses ; et bonne chance pour la fin, tranquille ! Bonne chance aussi à Béryl, qui m'aura appris certaines choses sur moi-même (notamment, que sous certaines conditions je n'aime pas les bâtons de réglisse). Merci à Marcia-Bouchon pour sa présence dans notre bureau et pour avoir supporté la mienne, pour sa patience et son soutien. Bonne chance pour ta thèse ! Merci à Tim d'avoir relu un bout de manuscrit et d'accepter de partir plus tôt au resto U (oui, il n'y a pas qu'à Marseille que ça arrive), merci à Rémy pour ses « lâche pâs » et ses conseils scientifiques, merci à Charles pour le soutien moral et physique sous forme de nourriture sucrée. Je remercie aussi William pour ses conseils, sa sympathie et sa relecture du résumé anglais de la thèse. Merci de manière générale à tous ceux et celles qui participent au bon fonctionnement de l'Institut de la Vision, un lieu que je n'ai pas encore eu l'opportunité de bien explorer mais où il fait en tout cas bon travailler.

Un très grand merci à Emmanuel Demaury et Sylvie Faivre pour leur gentillesse et leur efficacité à l'ex-NPA. Merci aux précédents docteurs de l'équipe qui sont maintenant loin, Manu, J.-B. et Romain, qui m'auront tous beaucoup appris à leur manière... et merci à Pierre en passant. Merci aussi à Julie et Bénédicte, en particulier pour leurs thèses de qualité qui m'ont été d'une grande utilité. Merci à Jean Mariani, directeur de l'ex-NPA, pour m'avoir donné la possibilité d'y travailler. Enfin, un grand merci à Kamel Megherbi et l'équipe du service de reprographie des thèses de l'UPMC pour leur accueil chaleureux et leur travail de qualité.

De l'autre côté de la France, mais aussi de la thèse, se trouve le Laboratoire de Neurobiologie de la Cognition, apparemment un endroit d'où il est très difficile de partir. Contrairement aux prévisions initiales (« oui, je reste seulement six mois... »), une grande partie de cette thèse s'est déroulée là-bas. Pour m'avoir acceptée dans leur équipe et leur laboratoire, m'avoir donné l'opportunité d'y travailler et d'y apprendre, je remercie chaleureusement Etienne et Bruno. Je ne m'étendrai pas sur l'ambiance qui règne au laboratoire car elle a déjà été généreusement décrite dans la littérature existante des remerciements ; je voudrais juste souligner le caractère à la fois naturel et constructif des échanges qui y ont lieu, intra- ou inter- équipe(s), ce qui permet probablement de faire avancer chacun des membres sur les plans personnel autant que scientifique. Je remercie très sincèrement Etienne, pour son accueil, pour m'avoir appris les multiples aspects liés à la manipulation des rats et à l'enregistrement électrophysiologique *in vivo*, pour son suivi, sa disponibilité exemplaire, son humour (je sais, c'est du réchauffé, mais c'est dur de passer à côté), sa patience, ses enseignements, sa capacité à clarifier (ou alors éliminer) des raisonnements parfois alambiqués, et enfin son

implication dans la phase de rédaction du manuscrit. J'espère ne pas avoir trop augmenté son niveau de stress, ou alors seulement de manière transitoire.

De nombreuses personnes du LNC ont contribué à ce travail, à différents moments. Tout d'abord, merci à Bruno par lequel mon premier contact avec le LNC s'est fait, pour sa vision claire des nombreux projets du laboratoire, sa contribution précieuse aux études de cette thèse, et pour ses programmes en C dont on ne prononcera pas le nom. Merci beaucoup à Luciana pour l'aide logistique de tous les instants, pour son efficacité et sa bienveillance. Un très grand merci à Valérie et Elodie pour leur travail compétent et attentionné qui a permis aux rats d'être dans les meilleures conditions possibles pour produire ce fameux comportement flexible et orienté vers un but. Merci à Didier pour le soutien technique concernant la mise au point des différents dispositifs expérimentaux, pour les nombreux cafés qu'il m'a avancés et notre capacité mutuelle d'attention divisée (ou pas...). Merci aussi à Dany pour son aide notamment sur le dispositif de labyrinthe radial. Merci à Catherine Marra, en particulier pour son importante contribution à la fabrication des plots d'enregistrements. Un très grand merci à Pierre-Yves, pour son aide générale, ses précieux conseils et ses nombreuses démonstrations sur différents aspects liés à l'enregistrement électrophysiologique : la fabrication des électrodes, la chirurgie, l'enregistrement, le spike-sorting (ah ! le spike-sorting...), la difficile phase de perfusion, l'histologie. Je le remercie aussi pour sa bonne humeur (mais aussi sa susceptibilité), son invitation à Allauch et nos discussions. Dans la lignée, je remercie Francesca, pour sa bonne humeur, ses programmes Matlab, sa relecture d'une partie de la thèse et les chocolats de la dernière ligne droite!! Merci aussi à Borís pour une fonction python qui m'a été très utile. Merci à Ingrid pour son dynamisme contagieux et pour avoir supporté de partager sa salle de manip avec moi ;).

Une partie non-négligeable du travail de cette thèse repose sur l'aide de trois ex-stagiaires : merci à Clémence Asselin (qui m'a probablement autant appris que la réciproque), pour son aide ingénieuse lors de la phase (délicate) des expériences de dévaluation (qui l'était encore plus lorsque la stagiaire n'avait pas encore été nourrie); merci à Laurène Save, ma deuxième stagiaire, pour son aide studieuse lors de l'expérience d'alternance (et à qui je souhaite bonne chance dans la voie périlleuse mais ô combien enrichissante de la thèse) ; merci enfin à Julie Gougot, ma troisième et dernière stagiaire, pour son aide tout aussi précieuse lors du premier mois de l'expérience d'électrophysiologie, et à qui je dois une bonne anecdote sur la morsure des rats.

D'autres membres du LNC ont aussi contribué à ma 'formation' : un grand merci à Paul pour son accueil dans l'équipe, pour sa machine, pour son apprentissage des différentes techniques – avec la minutie qui le caractérise –, pour ma première chirurgie, et surtout, pour ses gâteaux aussi beaux que bons. Merci aussi à son frère Jacques pour nos échanges scientifiques très intéressants. Mes sincères remerciements à Vincent, pour son hyperactivité, son humour « subtil », pour ses nombreux conseils toujours très pointus, sa relecture d'un bout de thèse et ses remarques sur la présentation. Tout aussi sincèrement et chaleureusement, merci à ma chère Sophie Renaudineau avec qui nous n'avons finalement passé que trop peu de temps au même endroit en même temps ! Peut-être cela va-t-il changer... Merci pour ses conseils humains et scientifiques, en particulier sur la fabrication des électrodes. Merci à Sophie Dutheil, où la même remarque s'applique, pour nos rares mais précieux échanges et pour sa présentation de thèse qui était un modèle de pédagogie. Un très grand merci à Candie, pour son soutien, son aide tout au long de la thèse, de ses explications sur la fabrication du matériel à sa contribution lors de la phase de la rédaction de la thèse, en particulier pour la

bibliographie. Merci à elle et à Joëlle d'avoir été un point de chute stable et chaleureux à chaque fois que je revenais à Marseille. Merci à Isabel pour sa bonne humeur, à Nico pour sa mauvaise humeur... plus sérieusement, merci à ce dernier pour toutes nos discussions (dont une proportion m'a quand même fait changer d'avis sur certains points, si, c'est vrai), pour Magic (et d'ailleurs... on pourrait reparler de ce booster ?), pour m'avoir (un peu) donné le goût de cuisiner, et pour son soutien précieux. Merci en passant à Audrey pour son soutien et pour être elle-même, tout simplement. Merci aussi à Michel, pour les sympathiques échanges que nous avons eus (c'est la vie !). Merci à Matthew pour sa personnalité attachante, pour ses remarques sur une partie du manuscrit, pour son accueil (et celui de Sarah) lors de moments passés à Marseille. Merci à Mathieu pour nos discussions (scientifiques ou autres). C'est bientôt à toi, et je ne m'inquiète pas ! Merci à Murray pour les moments partagés à Marseille, pour sa vision scientifique très claire, pour ses remarques sur une partie du manuscrit, pour être lui-même, pour le footing et ses délicieux gâteaux. Merci à Asma pour sa vision positive lors de son passage au LNC... et bienvenue à Paris !

Merci à tous les autres membres du LNC de manière générale pour leur accueil et pour l'ambiance à la fois chaleureuse et scientifique qu'ils mettent au laboratoire : Jean-Luc pour ses conseils et le fait qu'après lui avoir parlé, on ressorte toujours plus heureux qu'avant, Laurence Casini, Franck Chaillan, Thierry Hasbroucq, Franck Vidal, Bruno Truchet, Magdalena, Laure, Marianne (aussi pour sa compagnie lors des longues soirées au labo), Sébastien (maintenant bien loin) pour nos footings et discussions, Simon et son aide efficace dès qu'on a besoin de lui (ok, on aurait pu le laisser sur le site..), Jérémie Danna pour son sourire et ses conseils (et d'ailleurs, j'ai toujours ma revanche à prendre au ping), Jérémie Camon, Marieke, Viet, Christine Manrique, Christophe (bonne chance pour ta thèse, et profite bien du bureau), Samira (de même, et merci pour les M&Ms), Emmanuel et Sandie (bonne chance aussi à tous les deux).

A chaque « ex-doctorant » qui se dit que sa thèse ne sera probablement pas lue : je dis merci à Pierre-Yves Jacob, Bénédicte Babayan, Julie Lefort, Sophie Renaudineau, Vincent Hok, Paul de Saint Blanquat, Alice Alvernhe, Tiffany Van Cauter, Louis-Emmanuel Martinet, Aude Retailleau, Michael Loureiro (je n'ai pas *tout* lu non plus).

À l'attention de mes ami(e)s extérieur(e)s au travail, je voudrais les remercier pour leur présence plus ou moins lointaine (dans l'espace) mais en réalité toujours proche et m'excuser pour les nombreuses fois où un « je suis en manip », « je peux pas, faut que j'travaille », ou « j'ai une thèse à écrire » ont interféré avec une potentielle occasion de se retrouver : Lise, Merima et Bahrija, Alizée, Farouk, Anne, Mélanie (L) et Delphine (qui ne liront probablement pas cette thèse, mais on ne sait jamais), Thérèse, Davy, Manu (G), Thomas, Pépette, Assia, Ghislaine. Ceci s'applique aussi à la famille Duvelle : Camille (et nos discussions trop rares mais très appréciées), Lise (ma cousine), Maryvonne, Charles et Jacqueline, en particulier. Merci aussi aux gars du ping, L'USCM à Marseille, mais en particulier l'ESC XV de Paris (allez-y, vous ne serez pas déçus !). Un grand merci en particulier à Denis pour son soutien, Florian pour la nourriture vidéo-ludique, Jérôme pour les discussions, Juan pour les matches et Olivier qui m'a autorisée à en rater quelques-uns (« Eléonore est dispensée d'être capitaine, elle a une thèse à écrire »).

Pour ses conseils sur la dévaluation, je remercie Etienne Coutureau. Un grand merci à Mathieu Wolf pour m'avoir donné l'opportunité de présenter une partie des résultats de cette thèse lors d'un congrès du GDR NeuroMem (en 2014). Merci à la société Viewpoint pour leur mise au point de l'automatisation de la tâche principale de cette thèse. Je tiens à remercier les multiples personnes participant au maintien et au développement du langage de programmation Python, utilisé pour l'immense majorité des analyses de cette thèse. Je remercie chaleureusement la Société Française des Neurosciences pour le soutien financier étudiant qu'elle m'a accordé lors du congrès de 2013 ainsi que l'European Brain and Behaviour Society pour la bourse accordée pour participer au congrès de Munich (en 2013). Merci beaucoup à Mehdi Khamassi et Sidney Wiener d'avoir participé à mon évaluation à mi-parcours et pour leurs conseils pertinents qui ont très probablement joué pour le résultat final.

Last but not least, toute ma reconnaissance va aux deux organismes qui m'ont financée lors de ces quatre années de thèse : la Direction Générale de l'Armement, représentée par Didier Bazalgette, pour les trois premières années et la Fondation pour la Recherche Médicale pour l'aide de fin de thèse en dernière année (code dossier FDT20130928415).

Finalement, on pourra en conclure que si l'expérience de la thèse m'aura permis (ou au moins donné l'impression) de mieux comprendre le fonctionnement du cerveau et des rats¹, elle m'aura aussi, accessoirement, appris à aimer les humains.

¹ Ces fameux rats ayant été traités aussi "humainement" que possible.

Contents

Abstract	i
Résumé.....	ii
Remerciements.....	v
Contents	xi
Foreword	xvii
 PART ONE: INTRODUCTION	 xxi
 Chapter 1 – Memory.....	 1
1.1 Taxonomy of memory systems	1
1.1.1 Short-term versus long-term memory.....	1
1.1.2 Declarative versus nondeclarative memory.....	3
1.1.3 Episodic versus semantic memory.....	3
1.1.4 The particular case of episodic-like memory.....	5
1.2 Neural mechanisms of memory	9
1.2.1 Neural representations	10
1.2.2 Neural code	11
1.2.3 Possible mechanisms of long-term memory.....	13
1.2.4 Possible mechanisms of working memory.....	15
1.2.5 Dynamic mechanisms of memory	17
1.3 Summary: multiple memory systems.....	21
 Chapter 2 – Decision-making	 23
2.1 Goal.....	23
2.2 Goal-directed behaviour and decision-making systems.....	24
2.3 Motivation.....	26
2.4 Value.....	27
2.5 The exploration / exploitation ratio	29
2.6 Vicarious trial-and-error as a marker of deliberation?.....	29
2.7 Conclusion: decision-making in spatial behaviour	30
 Chapter 3 – Spatial cognition.....	 31
3.1 Perception of space	31
3.1.1 Types of sensory information.....	31
3.1.2 Multisensory integration.....	32
3.1.3 The importance of geometry.....	34

3.1.4	Spatial reference frames	35
3.2	Navigation strategies	35
3.2.1	The importance of exploration	36
3.2.2	Path integration	39
3.2.3	Guidance strategy	42
3.2.4	Response strategy	42
3.2.5	Route strategy	43
3.2.6	Place strategy and the cognitive map theory	43
3.2.7	Interactions between navigation strategies across learning	51
3.3	Spatial cognition and memory systems	51
3.3.1	Declarative versus procedural memory in spatial cognition	51
3.3.2	Working versus reference memory in spatial cognition	52
3.3.3	Episodic and semantic memory	53
3.4	Conclusion: spatial cognition and the hippocampus	53
Chapter 4 – The hippocampus: at the crossroads of space and decision-making?		55
4.1	Anatomy of the hippocampus: a brief overview	55
4.1.1	The hippocampal formation	55
4.1.2	Disto – proximal organisation and the trisynaptic circuit	56
4.1.3	Dorsoventral organisation	58
4.1.4	Connectivity with other parahippocampal structures	58
4.1.5	Connectivity with extra-hippocampal region structures	59
4.2	The hippocampus within a spatial processing network	60
4.2.1	The hippocampus and place cells	60
4.2.2	Cortical and subcortical networks with head direction cells	62
4.2.3	Cortical networks with boundary cells	64
4.2.4	The entorhinal cortex and grid cells	65
4.2.5	Other structures with spatially tuned activity	67
4.3	The hippocampus within a decision-making network	69
4.3.1	The hippocampus and the striatum	69
4.3.2	The hippocampus and the amygdala	71
4.3.3	The hippocampus and the VTA	72
4.3.4	The hippocampus and the prefrontal cortex	74
4.4	Conclusion	77

Chapter 5 – Hippocampal place cells: beyond spatial memory	81
5.1 Place cells as an allocentric representation of space	81
5.1.1 Anatomical localisation of place cells	81
5.1.2 Place fields and the coding of space	82
5.1.3 Place fields stability.....	83
5.1.4 Place cells and multisensory integration.....	84
5.1.5 Place cell firing variability.....	86
5.2 Population activity of place cells	87
5.2.1 Place cells are not alone.....	87
5.2.2 The importance of exploration	89
5.2.3 Remapping	89
5.2.4 Pattern completion and pattern separation	91
5.2.5 Replay and sequences.....	91
5.3 Place cells and goal-directed behaviour.....	94
5.3.1 Place cells and attention	94
5.3.2 Place cells and navigational decisions.....	95
5.3.3 Place cells can represent past and future positions	96
5.3.4 Manipulation of the ‘memory engram’	98
5.4 Place cells and decision-making parameters	99
5.4.1 Behavioural and task – related correlates of place cells.....	99
5.4.2 Influence of the goal location on place cell activity.....	100
5.4.3 Reward – related correlates of place cell activity ?	105
5.5 Conclusion.....	106
PART TWO: EXPERIMENTAL CONTRIBUTIONS.....	107
Objectives and working hypotheses	109
Chapter 6 – Introducing outcome devaluation in spatial cognition	111
6.1 Objectives	111
6.2 Radial-arm maze tasks	111
6.2.1 Experiment number one: working memory.....	112
6.2.2 Experiment number two: reference memory	113
6.2.3 Experiment number three: cost increase.....	114
6.2.4 Experiment number four: limited visits	114
6.2.5 Reward selection	115
6.3 Continuous alternation tasks	115
6.3.1 Experiment number five: free alternation and 2:1 rule.....	116
6.3.2 Experiment number six: continuous alternation	116

6.3.3	Experiment number seven: continuous alternation	117
6.4	Conclusion.....	118
Chapter 7 – Goal-directed behaviour in the two-goal navigation task.....		119
7.1	The two-goal navigation task: material and methods.....	119
7.1.1	General paradigm of the two-goal navigation task	119
7.1.2	Material.....	120
7.1.3	Procedures of the experiment.....	124
7.1.4	Further procedures	127
7.1.5	Behavioural analysis methods	129
7.2	Behavioural results	132
7.2.1	General performance.....	132
7.2.2	Time course of behavioural preference	138
7.3	Discussion: behaviour in the two-goal navigation task	142
7.3.1	Summary of results	142
7.3.2	Is the two-goal navigation task hippocampal-dependent?.....	143
7.3.3	Did rats demonstrate goal-directed behaviour?	144
7.3.4	Comparison with existing paradigms.....	145
7.3.5	Finer analysis of natural decision-making.....	145
7.4	Conclusion.....	146
Chapter 8 – Hippocampal goal-related activity in the two-goal navigation task.....		147
8.1	Methods: preprocessing and statistical analysis of electrophysiological data	147
8.1.1	Pre-processing methods.....	147
8.1.2	Goal-related activity.....	152
8.1.3	Peri-event time histograms.....	152
8.1.4	Analysis of the reward consumption-related activity.....	152
8.1.5	Goal-related activity significance.....	153
8.1.6	Analysis of the value-related hippocampal activity.....	154
8.2	Results: hippocampal goal-related activity in reference sessions.....	156
8.2.1	Histological results.....	156
8.2.2	Number and types of cells.....	156
8.2.3	Influence of goal location on place fields.....	158
8.2.4	Goal-related coding by the hippocampal population	159
8.2.5	Temporal profile of goal-related activity and reward consumption-related activity ..	161
8.2.6	Conclusion	164
8.3	Results: is hippocampal goal-related activity modulated by goal value?.....	165
8.3.1	Temporal profile of goal-related activity under value-changing conditions.....	165

8.3.2	Goal-firing preference as a function of goal value	168
8.3.3	Activity of hippocampal interneurons at goal locations	170
8.3.4	Within-session evolution of firing preference.....	171
8.3.5	Conclusion	173
8.4	Discussion	173
8.4.1	Possible sources of hippocampal goal-related signal	174
8.4.2	Information conveyed by hippocampal goal-related activity	178
8.4.3	Use of the hippocampal goal-related signal.....	180
8.4.4	Limits and further investigations	182
Chapter 9 – Conclusions		185
Résumé (long).....		187
Appendices		199
I.	Parameters of devaluation experiments.....	200
II.	Behavioural training of the two-goal navigation task.....	202
III.	Cumulated time maps for all cued sessions	203
IV.	Recorded unit types as a function of electrode depth.....	204
V.	Side and value preference index in subpopulations of hippocampal pyramidal cells.	205
References.....		207

Foreword

The brain is a complex organ able to perform a variety of functions, predominantly dedicated to processing information from the outside world but also from the organism itself in order to produce adapted responses. Among these functions, **cognition** refers to all brain processes related to knowledge, such as learning, memory, attention, decision-making and language. One of the domains in which cognition is expressed under different forms is **spatial cognition**. It is the ability of an organism to acquire or update spatial knowledge, organise it, and use it to express a behaviour adapted to the context. Animal species and mammals in particular have a wide range of spatial learning and navigation abilities that allow them to behave flexibly even in complex and dynamical environments.

Four decades ago, the development of electrophysiological recording techniques in freely moving animals allowed insights into the neural bases of spatial cognition to be gained. At that time, the **hippocampus** was known for its role in memory – among others – as revealed by studies of amnesic patients. Extracellular recordings in the hippocampus of freely moving rats shed light on a striking property of hippocampal cells: their discharge appeared to represent the location of the rat. The fact that such an abstract notion as space seemed to be encoded in their activity was at first sceptically received by the scientific community. However, after scepticism came enthusiasm in the light of the surprising properties of these ‘**place cells**’. This discovery directed a large body of research towards the understanding of neural mechanisms underlying spatial cognition. Such effort evidenced a broad network of neural structures involved in spatial cognition within which the hippocampus plays an important role, particularly when complex spatial processing and behavioural flexibility are involved. Unravelling the mechanisms that endow an animal with this behavioural flexibility is still today a challenge.

The work presented in this thesis aimed at contributing to the understanding of **the role of the hippocampus in spatial cognition and adaptive behaviour**. Since its involvement in spatial memory is rather well established, we asked to what extent it could also be involved in spatial decision-making. Recent works indicate that the hippocampus may not be restricted to purely representing space. In particular, new insights were recently gained about how places of interest in our environment, or goals, could have a specific importance in the hippocampal spatial representation.

In the first part of this manuscript, we introduce the fundamental notions involved in spatial cognition and flexible behaviour. We focus on the role of the hippocampus in those functions that were covered by the experimental investigation carried out during this doctoral thesis. The first chapter concerns memory and it reviews the current hypotheses about the neural mechanisms underlying this brain function. The second chapter introduces some of the fundamental notions involved in decision-making, by relating them to spatial cognition. The third chapter is devoted to spatial cognition itself, in an attempt to evaluate the specific conditions that require a functioning hippocampus. The fourth chapter describes the anatomy of the hippocampus and its relationships with surrounding structures. In particular, we ask to what extent the hippocampus could interact

with structures generally considered as belonging to a ‘decision-making’ circuit. We stress the fact that the way by which the brain combines spatial and decision-making parameters to direct behaviour is still not fully understood. Because single cell activity can bring new insights into the mechanisms underlying the combination of these parameters, the fifth and last introductory chapter reviews current knowledge on hippocampal place cells. It relates the activity of these cells to the above-mentioned functions, namely, spatial memory and decision-making. We present recent studies that shed light on a possible role of place cells in representing spatial goals. However, the nature of this goal-related signal and its possible significance for goal-directed behaviour remain open questions.

In the second part of the thesis, we present the experimental work and the new evidence provided by this research. The sixth chapter overviews a first series of studies in which we used a specific method to modify goal value, namely, outcome devaluation. We used several paradigms to assess the influence of a decrease in reward value on the spatial behaviour of rats. The seventh and eighth chapters both focus our main experiment based on a novel paradigm called the two-goal navigation task, in which we performed electrophysiological recordings in the hippocampus of rats. The behavioural methods and results of this experiment are exposed in chapter 7 while the electrophysiological results (and the associated methods) can be found in chapter 8. The results and their interpretation are discussed at the end of each chapter. Finally, the ninth chapter concludes by highlighting the new insights gained about the role of place cells, and of the hippocampus in general, in goal-directed navigation.

A roadmap of the thesis is presented in Fig. 1.




Fig. 1: Roadmap of the thesis.

PART ONE: INTRODUCTION

Chapter 1 – Memory

Memory can be defined as any physical change that carries information about the past (Redish and Mizumori, 2014). In the brain, these physical changes are thought to be implemented inside and between neurons, as will be seen in the present chapter. Memory is tightly related to **learning**, which is the process by which memories are created or modified. Several criteria can be used to categorise memory: the neural systems involved, the time during which it holds information, whether or not it can be deliberately expressed, and the behavioural flexibility it mediates.

1.1 Taxonomy of memory systems

The work of Scoville and Milner (1957) was influential in defining the current taxonomy of memory systems in relation to their neural bases. Scoville and Milner focused on a patient, originally known by his initials (H.M.), who had intractable epilepsy. His seizures could not be alleviated by any existing medicine and eventually led him to lose his job. As an attempt to cure his condition, he underwent surgery to remove the estimated foci of his epilepsy, which were localised in his temporal lobes. Most of this part of his brain was bilaterally removed, including the major portion of the hippocampus, but also the amygdala and parts of the parahippocampal region. H.M. recovered quite well from his surgery and the frequency of his seizures was largely reduced. However, he was left with both retrograde and anterograde amnesia. Retrograde amnesia affects the memory of events experienced prior to the cause of amnesia, while anterograde amnesia impairs the long-term remembering of events experienced from the onset of the amnesia. Interestingly, retrograde amnesia is generally temporally graded, meaning that recently stored events are completely forgotten but not older ones. This was the case for H.M., whose retrograde amnesia initially extended to events encoded 2 years before his surgery, a duration which evolved with time. Aside from these severe long-term memory deficits, H.M.'s short-term memory, language abilities, and previously learned motor skills were operational. He was also able to learn new motor skills, although with lower performance than normal subjects (Corkin, 1968). The series of studies about H.M. had two major impacts: first, they refined the classification of memory systems, by suggesting that short-term and long-term memory were independent, as were declarative and non-declarative memory, and second, they shed light on a possible role of the hippocampus in declarative memory². We will now address the currently used taxonomy of memory systems.

1.1.1 Short-term versus long-term memory

Short-term and long-term memories are differentiated by the period of time during which information stays in memory. Information held in **short-term memory** can be easily disrupted whereas long-term memories will require more energy to be modified. As an example, if one must temporarily retain a phone number before dialling it, the memory for the phone number will be available for a few seconds but will have disappeared at the end of the conversation. On the contrary, one can recall one's own phone number at any moment even though many events

² A review of the recent findings from H.M., who died in 2008 at the age of 82, can be found in Squire, 2009.

happened between learning and recalling. The capacity of short-term memory is limited and was postulated to be constant (Miller, 1956; Cowan et al., 2004). However, chunking, which is the grouping of items, can be used to increase the number of elements held in short-term memory.

On the contrary, **long-term memory** has a seemingly unlimited capacity and enables the retention of information for very long durations, possibly for the whole lifespan of an individual. As an example, Thorpe and collaborators recently showed that elderly persons were able to recognise the characteristic music from the opening credits of TV shows that they had last seen more than 40 years before the test³. A common distinction is made between recent or remote memory (Frankland and Bontempi, 2005; Moscovitch et al., 2006). Recent memory concerns information that was freshly learned but that has not yet been consolidated, quite similarly to short-term memory. Remote memory concerns memory traces that outlast short-term processes: they can be recalled days or even decades after the memorised event occurred (Moscovitch et al., 2006).

1.1.1.1 *Short-term memory*

A wide range of models of memory have been proposed by cognitive psychologists and neuropsychologists, supported by results from studies of normal or amnesic subjects like H.M. (to cite some of the most influential: Atkinson and Shiffrin, 1968; Baddeley and Hitch, 1974; Cowan, 1988; Tulving, 1995; Baddeley, 2000; Eustache and Desgranges, 2008). An important distinction that we will retain is the one made between short-term and working memory: short-term memory concerns the retention of information over a brief period of time while **working memory** is the temporary manipulation and use of this information to guide behaviour (Aben et al., 2012). Moreover, two classes of models can be defined according to the putative structures supporting working and long-term memory. Systems models consider that short and long-term memory are supported by separated structures (Atkinson and Shiffrin, 1968; Baddeley and Hitch, 1974; Baddeley, 2000). On the contrary, state-based models propose that working memory and long-term memory correspond to different activation states of identical structures (Cowan, 1988, 2008; Oberauer, 2002; see Larocque et al., 2014 for review). Nevertheless, both types of models converge on the fact that a fundamental role in working memory is played by attention. **Attention** can be defined as the focus of cognitive resources towards a subset of information available at a given time to the organism, usually with the effect of enhancing the processing of that information.

1.1.1.2 *Long-term memory*

Among the various existing models of long-term memory, we chose to focus on the one proposed by Squire and collaborators (Fig. 2), as it is still widely used today (Squire and Zola, 1996; Squire, 2004). Each functionally different type of memory is associated with the brain structures thought to be central to it, forming **memory systems**. Different memory systems are thought to function in parallel: an example given by Squire is how the childhood experience of being chased by a dog can give rise to a declarative memory for this event as well as an emotional memory that will be expressed as a phobia for dogs. Moreover, several systems can interact (Squire, 2004). From this taxonomy, we can

³ Communication presented at the NeuroMem annual meeting, May 2014.

already retain that the hippocampus, which is part of the medial temporal lobe, is thought to be involved in declarative memory. Other structures which we will briefly address in the current work are the striatum, involved in procedural memory, and the amygdala, which deals with emotional memory. This taxonomy focuses on long-term memory.




Fig. 2: Example of taxonomy of long-term memory systems.

Categories of memory systems. From Squire, 2004.

1.1.2 Declarative versus nondeclarative memory

The distinction between declarative versus nondeclarative memory is also often addressed in terms of explicit versus implicit memory. This distinction concerns the recall of a memory. **Implicit memories** are characterized by the fact that they are expressed by action (knowing how) rather than by conscious recollection (knowing that). Emotional responses are an example of implicit memories. They are thought to mainly be processed by the amygdala. Among implicit memory systems, we will mostly be confronted, in the present work, to **procedural memory**, which encompasses skills and habits. Procedural memories take time to be learned but, once acquired, they have the advantage of not requiring attention, which can then focus on other tasks (e.g., having a conversation about where to spend your next holidays while driving on a familiar road).

On the other hand, recall of explicit or **declarative memory** usually requires attention and, in humans, can be expressed verbally. Declarative memory is representational, insofar as it provides a way to model the external world (Squire, 2004). It comprises the memory for facts, or semantic memory, and the one for events, also called episodic memory.

1.1.3 Episodic versus semantic memory

The distinction between episodic and semantic memory was first made by Tulving (1972). **Episodic memory** concerns specific personal experiences that took place in a particular context and at a given moment. The ‘personal’ aspect is defined by Tulving as ‘autonoetic awareness’, which is the feeling of first-person subjectivity. An important feature of episodic memory is that it allows mental time travel in subjective time (Tulving, 2005, 2002). Mental time travel can consist in projecting oneself in the past to re-experience a scene, or projecting oneself towards the future to envision multiple

possibilities. On the contrary, **semantic memory** concerns the knowledge of facts, dissociated from the context (whether spatial, temporal or social) in which learning took place. As an example, knowing that Paris is the capital city of France is a semantic memory, whereas the memory of the last time one visited Paris is of episodic nature. The difference between both memory types is often made in human experiments by asking the subjects “do you remember ...?”, which addresses episodic memory and contains a sense of time travel, or “do you know ...?”, which refers to semantic memory and is dissociated from time (Knowlton and Squire, 1995). Similarly, other authors state that episodic memory is expressed through recollection whereas semantic memory is expressed through familiarity (Diana et al., 2007; Eichenbaum et al., 2007).

Several views coexist about the relationship between episodic and semantic memory. One of them states that semantic memories would result from a progressive process extracting consistent features of episodic memories (Squire and Knowlton, 1995; Eustache and Desgranges, 2008). Conversely, another hypothesis is that episodic memory relies on semantic memory and cannot exist without it. In this view, episodic memory would result from the association of semantic memories and the addition of auto-noetic awareness (Tulving, 2005; Klein, 2013). Dere and collaborators (2010) proposed that this auto-noetic feeling could be generated by the recall of the emotional state associated to a particular memory.

The separation between semantic and episodic memory is supported by psychological studies of patients with brain lesions to the medial temporal lobe such as H.M. (Tulving, 1985; Tulving et al., 1991; Vargha-Khadem et al., 1997; Klein and Nichols, 2012; Klein, 2013). These patients usually suffer from an amnesia that selectively impairs episodic memories. These impairments sometimes go along with an inability to plan future events (Tulving, 1985). When focusing on the anatomical extent of lesions from these amnesic patients and others, a common brain region reliably involved in episodic memory deficits is the hippocampus (Zola-Morgan et al., 1986; Tulving, 2002; Squire, 2004; Klein, 2013).

In order to allow the neural bases of such intriguing phenomena to be further studied, a major issue remains as to whether nonhuman animals have declarative memory abilities. In humans, this memory system is usually probed using language. However, language is thought to be an exclusively human skill, although other animals are not denied forms of communication and the characteristics specific to human language are still a matter of debate (Hauser et al., 2002; Pinker and Jackendoff, 2005; Corballis, 2007). One way to bypass the language issue in declarative memory is to focus on the defining characteristics of this type of memory, such as its flexibility, the fact that it allows the remembered material to be manipulated and compared, that it consists of relationships among multiple items and events, and that it is representational, in contrast to implicit memory which is termed dispositional (Squire, 2004).

For example, the ability to perform transitive inference (more generally, inferential reasoning) is indicative of the use of declarative memory. **Inferential reasoning** is the logical process by which elements of individual memories are retrieved and combined to answer novel questions (Zeithamova et al., 2012). Transitive inference is a way to test for inferential reasoning, by first presenting a set of

items with relationships between each other, and second testing for the ability of the subject to infer new relationships between two items that were never presented simultaneously before. For example, if $A > B$ and $B > C$ in the training phase ('>' meaning 'should be selected over'), then when presented with A and C, the subject must select the item A. Rodents can solve transitive inference tasks (usually with different odours as items) and lesion studies in rats and mice have demonstrated the involvement of a number of structures in the ability for transitive inference, among which the hippocampus (Bunsey and Eichenbaum, 1996; Dusek and Eichenbaum, 1997; DeVito et al., 2010a) and the prefrontal cortex (DeVito et al., 2010b). However, it is unclear whether the hippocampus is involved in the acquisition of the overlapping memories necessary for transitive inference or in the flexible use of these memories necessary to infer new relationships, as demonstrated by the opposite conclusions of Van der Jeugd et al., 2009 and DeVito et al., 2010a. It also remains unclear whether inferential reasoning relies specifically on episodic or semantic memory, but it definitely holds characteristics of declarative memory.

The ability for episodic memory in nonhuman animals, for which the auto-noetic awareness aspect is still difficult to define in humans, remains an open question. Nonetheless, attempts were made to define an 'episodic-like' memory whose existence could be tested without relying on language abilities.

1.1.4 The particular case of episodic-like memory

Episodic-like memory can be defined by the characteristics of what is held in episodic memory, independently from whether it can be verbalised or it involves auto-noetic consciousness. These characteristics are the nature of the memorised episode, the spatial context where it happened, and the time when it happened – as opposed to semantic memory which stores information independently from its spatial or temporal context. The characteristics of episodic-like memory are generally summarised as the 'what, where and when' components of episodic memory (Nyberg et al., 1996; Clayton and Dickinson, 1998). Using these criteria, a series of studies tackled the question of the existence of episodic-like memory in nonhumans, initiated by the seminal study of Clayton & Dickinson in scrub jays (1998).

Scrub jays are birds that have a natural inclination to hide food in secured locations. Moreover, they love worms. Clayton and Dickinson relied on these two characteristics to question whether scrub jays could remember *what* kind of food they had cached *where* and *when*. In this study, the birds first learned that worms degrade overtime and become inedible after a long delay (124 h) but not a short delay (4 h). Then, they were required to cache either peanuts (which do not degrade) or worms in distinctive compartments of a sand-filled storage tray. Either 4 or 124 h after the caching phase, they were allowed to retrieve the food. If the worms had been cached 4h ago, birds unambiguously directed their visits towards the worm caches, showing that they had a marked preference for this type of food. However, after the long delay of 5 days (124 h), they first went to the peanuts compartment, thereby demonstrating that they had combined the 'when' with the 'what' information in order to guide their choice (Fig. 3).




Fig. 3: Episodic-like memory in scrub jays.

Left: scrub jay (*Aphelocoma coerulescens*) about to cache a wax worm. From Griffiths et al., 1999.

Right: Number of inspections for nuts or worms caches depending on the post-caching delay.

At the 4h delay, the worm cache was more inspected than the nuts cache. At 124h delay, it was the other way round. The overlaying percentages indicate the number of animals whose first visit was directed towards the worms cache. Note that the tests, contrary to training, were performed in extinction, i.e., with both types of food removed from the caches. Adapted from Clayton and Dickinson, 1998.

Thus, scrub jays demonstrate memory for the ‘what, where and when’ components of episodic-like memory (Clayton and Dickinson, 1998; Griffiths et al., 1999). A possible alternative interpretation of these results is that scrub jays could have solved the task using the fact that memory trace fades as time passes. Different trace strengths of the caching memories might have been sufficient to guide their choices (Eichenbaum and Fortin, 2003). Nonetheless, this study, by explicitly exposing three testable parameters of episodic memory, paved the way for further attempts to study episodic-like memory in nonhuman animals.

In particular, episodic-like memory was tested in rats, using an adapted radial arm maze paradigm (Babb and Crystal, 2006a). Similarly to the scrub jays study, the authors investigated whether rats would be able to remember what kind of food was located where, and when this specific type of food had previously been encountered. Because this experiment was important for our work, we will go into its details. First, the radial-arm maze is an apparatus composed of 8 arms with a cup at their extremity and connected by a central stem (Olton and Samuelson, 1978). The cups can be either baited with food pellets or not. The radial maze paradigm is used to test different forms of spatial memory, as will be exposed in further details in Sec. 3.3 (p. 51). Relevant to the episodic-like memory issue is the fact that a version of the paradigm allows to test the long-term memory of the ‘where’ component, i.e., whether rats remember which of the arms are baited. Indeed, if the rats remember that specific arms are always baited while others are not, they will first visit the baited arms. The episodic-like memory study relied on this apparatus using a particular protocol, which consisted in two parts. In the first part, the authors addressed whether rats were able to combine the ‘where’ and ‘when’ components of episodic-like memory. It involved a two-staged paradigm with a study and a test phase separated by a delay (Fig. 4). During the study phase, only half of the arms of the radial maze, randomly chosen, were accessible; two of them were baited with normal food while the remaining two contained specific, flavoured pellets. Upon entering the maze, the rat was allowed to visit each of the 4 opened arms before being removed. After either a short or long retention interval,

the test phase occurred. Rats were put back in the maze, all eight arms being accessible, but baited differently depending on the delay. After either delay, the locations baited with normal food replenished. In addition, only after the long delay would the two flavoured locations replenish. To perform optimally in this task, rats had to learn to revisit the flavoured pellets locations only after the long delay. After training, they managed to do so, demonstrating that they had somehow memorised the temporal component of the task. However, this part of the experiment did not assess whether they remember the type of food previously encountered.




Fig. 4: Episodic-like memory in rats – protocol.

After the study phase, during which only selected arms are baited, two test phases could occur. If the delay was short (1h), only chow pellets were accessible at previously non baited locations. After the long delay (6h), both chow (C) and flavoured (G and R) pellets locations replenished.

Adapted from Babb and Crystal, 2006a.

In the second part of the study, the authors tested the ‘what’ aspect of episodic-like memory, i.e., memory for the attributes of the reward. This is similar to the scrub jays experiment when the birds learned that worms degrade over time. Here, the authors used a devaluation procedure to selectively diminish the expected value of one of the flavoured pellets. The **devaluation procedure** is commonly used in the instrumental learning domain (examples will be given in Sec. 2.2, p. 24). Devaluation can be performed by two means: satiety devaluation or conditioned taste aversion. The first procedure consists in free-feeding a rat with a large amount of a given type of reward. In the second type of devaluation, the rat is allowed to consume the reward and is then injected with a lithium chloride solution that produces malaise. In both cases, the value of the reward is considered to be lowered as rats will stop or drastically reduce their consumption when subsequently presented with that reward. Babb and Crystal tested the two types of devaluation on one of the flavoured pellets during the retention interval to assess whether or not rats would remember the ‘what’ component of the study phase. Rats indeed selectively decreased their visits to the arm that provided the devalued reward. This was the result held for both types of devaluation procedures.

Thus, similarly to scrub jays, rats seem to demonstrate episodic-like memory. Episodic-like memory was also assessed in other species, such as birds, rodents or apes (for reviews, see Crystal, 2010; Pause et al., 2013). However, several aspects of the ‘what-where-when’ paradigms were criticised, beginning with the unclear nature of the ‘when’ component (which might be either a ‘how long ago’ component or an estimation of absolute time). This component was proposed to be replaced with a broader, ‘which’ component, as a memory for the specific occasion when the event was experienced

instead of the specific time (Eacott et al., 2005; Eacott and Easton, 2010). In this framework, the duration of the event is defined by the duration of the associated occasion, which can be a few seconds (e.g., ‘when I had this car accident’) or several hours (e.g., ‘when I visited London’). Another criticism is that these paradigms usually require intense training to learn the task rules (whether degradation of worms or replenishable food sources), and that during encoding the animals knew they had to remember a series of information as they would be tested later. On the contrary, true episodic memory relies on incidental and one-trial encoding. To address this issue, Zhou and collaborators (2012) tested the ability of rats to incidentally encode information about the presence or absence of reward. The authors first taught rats to ‘answer’ a question by letting them choose between two arms on a T-maze (a T-shaped maze with an entry arm and two lateral arms). In the entry arm, rats were either given food or not. If they had encountered food, rats had to then go left on the maze to get further reward; otherwise, reward would be delivered at the end of the right arm. Rats were also trained in a radial maze task. During probe tests, rats were first given the radial maze task where food could be either present and absent, and then the T-maze test where they had the opportunity to ‘report’ the previously encountered presence or absence of food. In these ‘unexpected’ tests, they answered with accuracy levels similar to their training performance. Thus, they were able to incidentally encode information about the occurrence of food and report it later. Interestingly, a lesion to the hippocampus (specifically the CA3 field) impaired the ability of rats to answer the unexpected question, but not an expected one (Zhou et al., 2012).

Other means to model and assess episodic memory in nonhuman animals have been proposed, including testing memory for the temporal order of items (where items can be either odours or places: Fortin et al., 2002; Eichenbaum and Fortin, 2003; Fouquet et al., 2010), memory for the source of the memory (Crystal et al., 2013), the ability for future planning, or prospective memory (Roberts, 2012; Wilson and Crystal, 2012). As an example, the study of Fortin and collaborators (2002) set aside the spatial aspect of episodic memory and instead tackled its temporal sequence aspect, relying on a definition of episodic memory as “sequences of events that unfold over time and space”, arguing that an episodic memory is composed not only of the item one is trying to recall but of the experience of preceding and following events (Eichenbaum and Fortin, 2003). In this experiment, rats were presented with a sequence of five different odours. During a test phase, they were given the choice between two of the previously encountered odours and had to select the one that had first been presented in the sequence. Control rats were able to learn this task while hippocampal lesioned rats were impaired. Importantly, both control and lesioned rats expressed normal performance in a simple odour recognition test, meaning that the hippocampus is not necessary for simple memory of events (Fortin et al., 2002). These kinds of tasks are quite similar to inferential reasoning tasks but require an additional memory for the time when events occurred.

An important thing to take into account is that each species evolved according to the specific environmental constraints and selection pressure it was confronted with, leading to multiple survival strategies. These different strategies often engage different neural systems. The challenge, if one is interested in knowledge transfer from nonhuman to human species, is to find neural mechanisms and abilities that can be expressed in both. For example, social species might have evolved a developed episodic memory system that takes into account the social context where an event

happened, while food-storing species might be more sensitive to the temporal context where food was cached (e.g., Clayton and Dickinson, 1998). As an example, rats tested in a food-storing task did not show any memory of the temporal context concerning when a specific type of food was cached (Bird et al., 2003). Conversely, when the ability of rats to scavenge for food was tested, as in Babb and Crystal's study (2006a), they demonstrate time-related memory. Another natural behaviour expressed by rats is the spontaneous exploration of novel objects (Ennaceur and Delacour, 1988). Using this spontaneous behaviour, the 'what-where-which' memory was assessed in rats (Eacott et al., 2005), and the 'what-where-when' memory demonstrated in mice (Dere et al., 2005). Finally, a study evaluated the 'what-where-when' memory of male mice by using female mice as the 'what' component of memory (Fellini and Morellini, 2013). Interestingly, these authors also showed that the hippocampus was necessary for the long-term retention of episodic-like memory. Overall, this underlies the importance of employing an ethological perspective, using tasks that address natural behaviour in the species one uses as a model (Eacott and Easton, 2010; Templer and Hampton, 2013).

The attempts to move toward a clear working definition of episodic memory, besides being useful from an ethological and neuroscientific perspective, can also be applied to refine nonverbal episodic memory tests in humans. In control populations, asking if a subject remembers or knows a given fact might not be such a rigorous test. Indeed, subjects can be made to believe that they do remember a given word while it was not actually present (Eichenbaum and Fortin, 2003). Thorough tests of episodic memory are specially needed when studying non-speaking populations (young children, language disabled patients) or for early detection of Alzheimer's disease, as episodic memory deficits are one of the major characteristics of the disease (Morris JC et al., 2001; Nestor et al., 2004; Storandt, 2008). Attempts to adapt nonhuman tests to human populations have been made and tend to indicate that the 'what-where-which' tests indeed require remembering rather than knowing (e.g., Easton et al., 2012). To this day, a reliable test of episodic memory which includes the auto-noetic awareness aspect has yet to be established (Pause et al., 2013).

The fact that nonhuman animals have memory abilities, both implicit and explicit, makes them appropriate models for the study of the neural bases of memory. Some of the known or hypothesised mechanisms of memory will now be outlined.

1.2 Neural mechanisms of memory

Memory, and the way the brain implements it, remain one of the major challenges in neuroscience (Kandel et al., 2014). However, after several decades of research on the neural bases of memory, some of the mechanisms responsible for implicit memory are now fairly well understood, and possible mechanisms that could underlie explicit memory have been proposed. We will focus on the concepts and mechanisms that are thought to support explicit memory.

1.2.1 Neural representations

Central to modern cognitive psychology and neuroscience is the concept of representation. It was first defined as a theoretical concept, but the subsequent development of extracellular recording techniques enabled neural correlates of mental representations to be found.

Roitblat (1982) defines a **representation** as “*a remnant of previous experience that allows that experience to affect later behavior*”. Another definition was proposed by Gallistel (1990, 2001): a representation is a functional isomorphism between an element of the environment and a neural process, that is to say, the representation and the element it represents have the same mathematical form, and the same relationships are shared between the symbols that compose the representation and between the items these symbols represent. Yet another definition, from deCharms and Zador (2000), states that a representation is a message that uses the rules defined by the neural code to carry information. Two characteristics can be used to clarify the concept of representation: its content and its function. The content of a representation is the information it carries. The function of a representation (when considering representation as a neural signal) is the effect it may have on further processing and eventually on behaviour (deCharms and Zador, 2000).

The concept of neural representations was commonly used in pre-behaviourism psychology but was literally banned during the behaviourist period, under the reasoning that they were not directly observable and that behaviour could be explained without them. However, cognitivists subsequently showed that certain animal species could express behaviours that are difficult to explain without internal representations (an idea already present in the work of Tolman, 1948). Nowadays this concept is quite well accepted and frequently used, sometimes in different forms such as the concept of memory **engram**, being the physical change(s) encoding a particular long-term memory (Chklovskii et al., 2004⁴). If such mental representations exist, they must be somehow implemented in a specific aspect of neural activity.

It is quite difficult to demonstrate the existence of neural representations, since a true representation must be observable in the absence of the stimulus that is encoded, in order to ensure that it is not a mere automatic and transient response to that stimulus. Nonetheless, a starting point for the study of representations is to assess whether neurons that can specifically fire in response to the presentation of a stimulus exist. Such neurons have indeed been found in the brain. Their activity is said to be ‘tuned’ to a given stimulus, i.e., their probability of firing increases when this specific stimulus is presented. The seminal works of Hubel & Wiesel (1959) exemplify this. The authors evidenced cells in the primary visual cortex (V1) of anaesthetised cats that responded specifically when the animal was presented with a visual stimulus (a bar with a specific orientation) at a particular position in its visual field. Similar ‘tuned’ neurons have been observed in many brain regions. It can be said that each of these ‘tuned’ neurons carry information about a precise parameter, either an input stimulus or an output action. The information about the bar position, for example, can be used by downstream brain areas to perform more complex operations (e.g., eventually recognise a specific face). Indeed, neural representations of more integrated concepts

⁴ See also the early theory on memory of Richard Semon, reviewed in Schacter et al., 1978.

have been found. A striking example comes from a series of studies using single neuron recordings in epileptic human patients. Quiroga and collaborators (2005) found cells that responded to the visual presentations of faces of famous persons such as the actress Jennifer Aniston, to drawings of this person, and even to visual presentations of his or her name. Interestingly, a majority of these 'concept' cells were found in the hippocampus. The authors proposed that they could be involved in declarative memory (Quiroga, 2012). The hypothesis that such cells are part of a mental representation is supported by the fact that they not only selectively discharge to the presentation of their preferred stimulus, but also do so during conscious recall in the absence of the stimulus (Gelbard-Sagiv et al., 2008).

1.2.2 Neural code

From a computational point of view, neuroscientists often aim at 'cracking the neural code', i.e., decoding the stimulus that led to a specific neural activity. The **neural code** can be defined as a system of rules and mechanisms by which a signal carries information (deCharms and Zador, 2000). The fact that specific neurons respond to specific stimuli probably facilitates computations and information processing.

It is an ongoing debate as to whether neural networks perform distributed or local coding of information (Bowers, 2009). **Local** (or localist) **coding** involves a relation between a single unit and a meaningful equivalence class of entities in the world (e.g., 'cat') (Hummel, 2000). Local coding does not imply that only one neuron would code for this class, but it does mean that one can infer the signification of the representation from the readout of a single neuron (Bowers, 2009). Local coding is often exemplified (in a caricatured manner) by the 'grandmother cell', which would represent the memory of one's grandmother. On the contrary, in the framework of **distributed coding**, the readout of a single neuron is generally not sufficient to assess the item it encodes. One must take into account the activity of a network of neurons to precisely decode the meaning of the representation. Distributed coding actually encompasses three types of coding: dense, coarse and sparse (Bowers, 2009). In **dense distributed** coding, each neuron is involved in the representation of multiple concepts and little information can be extracted from the activity of a single neuron. In **coarse distributed**, contrary to dense distributed coding, a single neuron is not generally assumed to participate to many representations. Rather, it has a broad tuning curve (such as V1 neurons), meaning that not only the 'preferred' item will make this neuron fire but also items that are close to it. In that case, pooling of several units helps reading out the encoded item or stimulus. A classic example of coarse distributed coding can be found in the primary motor cortex (Georgopoulos et al., 1986). Georgopoulos and collaborators trained monkeys to reach a visual target presented on a screen using a joystick. Single neurons from the primary motor cortex recorded in this study fired around the time of movement for several directions of movement, making decoding of the direction difficult from single neurons. However, the actual direction of movement could be inferred by computing a population vector, i.e., by taking into account the whole population of recorded neurons. Finally, in **sparse distributed** coding, a stimulus is coded by the activation of a small number of units, and each unit contributes to the representation of a few stimuli. The main difference between dense and sparse distributed coding is the number of neurons involved in the

representation. They would also have different properties: sparse coding could underlie the rapid acquisition of new knowledge without erasing previously stored information, but would be poor at generalisation. On the contrary, dense distributed representations would have generalisation abilities but would be prone to interference due to overlap between representations. The ‘concept’ cells recorded by Quiroga and collaborators (Quiroga et al., 2005; Quiroga, 2012) could be an example of sparse distributed coding (Quiroga et al., 2008; see also Wixted et al., 2014 for another example of sparse distributed coding in the hippocampus). Fig. 5 schematically exemplifies how a population of neurons would respond to different stimuli in the case of dense distributed coding versus local coding.

Test Items	Dense distributed coding								Local coding							
	-----Units-----								-----Units-----							
	1	2	3	4	5	6	7	8	1	2	3	4	5	6	7	8
<i>Clock</i>	0.1	0.2	0.2	0.1	0.2	0.2	0.2	0.1	0	0	0	0	2	0	0	0
<i>Chief</i>	0.2	0.1	0.2	0.2	0.1	0.1	0.2	0.2	0	2	0	0	0	0	0	0
<i>Map</i>	0.2	0.1	0.1	0.1	0.1	0.2	0.1	0.2	0	0	0	0	0	0	2	0
<i>Umbrella</i>	0.2	0.2	0.2	0.2	0.1	0.1	0.2	0.2	2	0	0	0	0	0	0	0
<i>Leather</i>	0.1	0.1	0.2	0.2	0.2	0.1	0.1	0.2	0	0	0	0	0	2	0	0
<i>Navy</i>	0.2	0.2	0.2	0.1	0.1	0.1	0.2	0.1	0	0	2	0	0	0	0	0
<i>Milk</i>	0.2	0.1	0.2	0.2	0.1	0.2	0.2	0.1	0	0	0	2	0	0	0	0
<i>Coffee</i>	0.1	0.2	0.1	0.1	0.2	0.2	0.2	0.2	0	0	0	0	0	0	0	2

Fig. 5: Illustration of dense distributed versus local coding.

Schematic representation of the normalised spike count for eight neuronal units in response to the presentation of a stimulus (clock, chief, etc.) in case of a dense distributed code (**left**) or local code (**right**). One item (e.g., clock) can be encoded by either a slight activation of the whole population of cells (the whole population must then be read out to decode the item) or by the strong firing of only one unit.

Adapted from Wixted et al., 2014.

Overall, the discussion about whether the brain performs any of the various forms of distributed coding or even local coding is not settled yet (Bowers, 2009). It is probable that several sorts of coding are implemented by the brain at multiple levels of processing. For example, the hippocampus would rather use sparse distributed coding while the cortex would use dense distributed representations (McClelland et al., 1995). Moreover, primary sensory areas could implement coarse distributed coding while sparse distributed coding would be preferred in higher-order areas (Quiroga and Kreiman, 2010). We will see in Chapter 5 that results from single cell recording in spatial cognition tasks can bring more arguments into this debate.

Another source of debate among neuroscientists is whether neurons perform temporal or rate coding (Thorpe et al., 2001). **Rate coding**, in which the information is contained in the firing rate of neurons, is less prone to noise but requires time to extract the information. In this scheme, the time required to precisely process a piece of information can be too important compared to the rapidity with which the organism completes certain tasks such as image categorisation (Thorpe et al., 1996). However, it could be used in other types of processes than such rapid decision-making. **Temporal coding**, where the timing of spikes encodes information, can require setting a reference to compare the time of spikes, but can also rely on synchrony between spikes or patterns of spikes. It enables

decisions to be taken very quickly but is sensitive to noise. Again, the study of neural bases of spatial cognition (see Chapter 5) can provide answers about this debate, mainly that both rate and temporal code coexist.

1.2.3 Possible mechanisms of long-term memory

A well-accepted idea in neuroscience, maybe the most fundamental one, is that neurons transmit information via action potentials. In a simplified sensory-motor loop, information from the outside world enters the brain by the way of sensory receptors, is processed and transmitted from neuron to neuron, and eventually leads to an action through motor neurons that drive actuators such as muscles. The output of the system can also be the absence of action, or further internal processing. Memory, in this framework, is embedded in neural networks: when a given experience modifies the way further information will be transmitted and processed, it means the brain just learned something. There are several ways to modify information processing by neurons: synapse efficiency can be modified (increased or decreased), connections between neurons can appear or disappear, the integration of information by a given neuron (i.e., how its inputs will trigger an output spike) can be modified.

The fact that changes in the connectivity between neurons could underlie learning is part of the connectionist approach in neuroscience and dates back to Ramón y Cajal (1894, p. 466) – for the French-speaking reader:

“on peut admettre comme une chose très vraisemblable que l’exercice mental suscite dans les régions cérébrales plus sollicitées un plus grand développ[e]ment de l’appareil protoplasmique et du système des collatérales nerveuses. De la sorte, des associations déjà créées entre certains groupes de cellules se renforceraient notablement au moyen de la multiplication des ramilles terminales des appendices protoplasmiques et des collatérales nerveuses; mais, en outre, des connexions intercellulaires tout à fait nouvelles pourraient s’établir grâce à la néoformation de collatérales et d’expansions protoplasmiques.”

In summary, he proposes that mental training could, in specific brain regions, induce the reinforcement of connections between neurons and the creation of new connections between those neurons. This idea was later specified by Hebb (1949, p. 62) in those terms:

“When an axon of cell A is near enough to excite a cell B and repeatedly or persistently takes part in firing it, some growth process or metabolic change takes place in one or both cells such that A’s efficiency, as one of the cells firing B, is increased”.

Thus, according to Hebb, the connection between two neurons should be reinforced if the presynaptic neuron always or often fires before the postsynaptic neuron.

The theoretical hypothesis that the synaptic connection between two neurons could be modified, namely, synaptic plasticity, was later proved true. A core process of synaptic plasticity that might underlie long-term memory is long-term synaptic plasticity, which encompasses long-term potentiation and long-term depression. **Long-term potentiation (LTP)** was first discovered by Lømo (1966) and subsequently described by Bliss and Lømo (1973). It is an increase in the efficiency of the synapse between a postsynaptic and a presynaptic neuron, meaning that once LTP is induced, a presynaptic spike will more easily trigger a postsynaptic spike. LTP was first induced by high frequency stimulation and its effects were shown to last for several days (Bliss and Lømo, 1973). This discovery was made in the hippocampus. It was later demonstrated that many of the synaptic connections involving hippocampal cells support LTP, which can also be induced in other structures such as the amygdala, the subiculum, the cerebellum and the prefrontal cortex, among others (Lynch, 2004). However, if memory is dependent on synaptic strength and the strength of synapses can only increase, the system will eventually saturate and no further learning would be possible. In addition to Hebbian plasticity, there is need for a mechanism that can selectively decrease the strength of synapses which do not participate in firing. Such a mechanism, termed **long-term depression (LTD)**, was also first demonstrated in the hippocampus (Stanton and Sejnowski, 1989; Dudek and Bear, 1992). Subsequent work found that LTD could also be expressed in many different brain regions (Malenka and Bear, 2004). Several types of both LTP and LTD actually exist and rely on multiple molecular and cellular mechanisms.

High-frequency stimulation, the experimental paradigm originally used to induce LTP or LTD, is unlikely to happen in a healthy brain. A more biologically plausible way of inducing long term plasticity is spike-timing dependent plasticity (STDP, see Fig. 6). As its name indicates, the temporal aspect is highly important in the STDP mechanism. First, if a presynaptic neuron repeatedly discharges just before a postsynaptic neuron, the synapse between the two neurons will be potentiated. Conversely, if presynaptic spikes are emitted in a short time window following the discharge of the postsynaptic neuron, the synapse will be depressed. Indeed, in the latter case, the presynaptic neuron did not participate in the firing of the postsynaptic one.




Fig. 6: Spike timing dependent plasticity.

Illustration of the STDP mechanism. If the presynaptic spike regularly occurs after the postsynaptic spike (left part), the synapse will be depressed, inducing LTD. Conversely, if the presynaptic spike regularly occurs just before the postsynaptic spike (on the right), the synapse will be potentiated (LTP). In both cases, the smaller the time interval between the two spikes, the stronger the effect.

From Zaehle et al., 2010.

With STDP, neurons often firing together at the proper timing will be more strongly associated. As a consequence, a spike from the presynaptic ('predicting') neuron will more strongly contribute to the generation of an action potential in the postsynaptic neuron. In this context, a network of cells using

STDP could 'learn' specific patterns of activation that could be neural representations. This mechanism has proven to be very efficient for learning in computational models of temporal coding (Gerstner et al., 1996; Song et al., 2000; Masquelier et al., 2009). As an example, Masquelier and Thorpe (2007) trained an artificial neural network using STDP rules by repeated presentations of natural visual scenes. After learning, many neurons of the network had become tuned to visual features such as specific orientation bars, similar to V1 neurons. Depending on the class of input stimuli, the features could differ, for example being generic features of faces if the input stimuli were a set of human faces (Masquelier and Thorpe, 2007). In biological organisms, STDP-induced long term potentiation was demonstrated by Markram and collaborators (1997). It was subsequently thoroughly investigated (e.g., in the hippocampus: Debanne et al., 1998; the amygdala: Bauer et al., 2001; the visual system: Meliza and Dan, 2006; for review, Markram et al., 2012). In particular, Bi and Poo (1998) demonstrated that the time window within which STDP can be induced is between 5 to 40 ms.

The fact that LTP can be induced in biological organisms does not prove that it is the support of long-term memory. Many studies specifically addressed the link between LTP and memory, initiated by Morris and colleagues (1986) who pharmacologically blocked LTP and showed that it impaired learning in a spatial memory task. Subsequent studies added support to the hypothesis that synaptic plasticity is one of the mechanisms underlying long-term memory (e.g., Wilson and Tonegawa, 1997; Whitlock et al., 2006). Although LTP was shown in specific conditions to last for over a year (Abraham et al., 2002), in general it decays after a few hours. It is actually probable that LTP is only one among several plasticity mechanisms underlying long-term memory (Stuchlik, 2014; Takeuchi et al., 2014).

Other approaches to the study of long-term memory mechanisms involve the manipulation of immediate-early genes (IEG), genes whose expression is rapidly and transiently triggered under different conditions, among which, LTP. The main IEGs thought to be involved in synaptic plasticity are *c-fos*, *zif268* and *Arc* (Abraham et al., 1991; for review, see Ramirez-Amaya, 2007). Genetic manipulations now enable the targeting of the expression of these genes in regions of interest (e.g., the hippocampus), as a means to go deeper in the study of their involvement in memory, and even to apparently modify stored memories (Ramirez et al., 2013).

1.2.4 Possible mechanisms of working memory

The most common hypothesis about the neural mechanisms of working memory is that it relies on persistent activity. In this view, a subpopulation of neurons spiking recurrently could temporarily hold an item in memory. Neurons with this kind of activity have been observed in the prefrontal cortex (Fuster and Alexander, 1971; Funahashi and Kubota, 1994; Goldman-Rakic, 1995). As an example, Funahashi and collaborators (1989) recorded individual neurons in the dorsolateral prefrontal cortex of monkeys during a working memory task. Monkeys had to fixate on a central point on a screen during the brief presentation of a visual cue indicating the position of a target on the screen. After a short delay without the cue, subjects had to saccade towards the remembered target position. The activity of an example neuron recorded during this task is shown on Fig. 7. This neuron strongly fired during the whole delay period, almost exclusively for one out of the eight

target positions. Such neurons are different from those that can be recorded in motor areas, as they are not active during the movement itself but during the delay period, in the absence both of any cue and movement. This kind of activity could be a form of neural representation.




Fig. 7: Delay and directional specific prefrontal neuron.

The activity of one neuron recorded in the dorsolateral prefrontal cortex of a rhesus monkey is presented as a function of time for each of eight possible angular locations of the target (surrounding plots). For a given target location, the spikes are shown for each trial in the upper part and the cumulative histogram is shown in the bottom part. The first two vertical lines indicate the cue presentation (that indicated where the target location would be) and the last one is the response (eye movement). The middle plot shows the angular position of targets. The duration of the delay period was 3s. Adapted from Funahashi et al., 1989.

Recurrent spiking would be facilitated at the molecular level by short-term synaptic plasticity mechanisms (Zucker and Regehr, 2002; Mongillo et al., 2008). More recently, new hypotheses on the role of this persistent activity have emerged, notably that it could take part in decision-making processes (Curtis and Lee, 2010). Indeed, working memory and decision-making probably share common features, as working memory is the manipulation of stored information to direct behaviour.

Note that we focused here on what could be termed explicit short- or long-term memory. For a review that includes molecular aspects of neural bases of implicit memory, including work on *Aplysia* and *Drosophila* or the eye-blink reflex of the rabbit, the reader is referred to Kandel et al., 2014. Long-term plasticity mechanisms such as LTP or LTD are likely to be involved in implicit memory, for example in the tuning of fine motor movements probably performed by the cerebellum (Ito, 2001).

1.2.5 Dynamic mechanisms of memory

1.2.5.1 Memory consolidation

Memories, once encoded, are not carved in stone. They can be updated or forgotten, and disturbed by interference or trauma. They appear to be more easily modified just after learning and gradually stabilise with time (McKenzie and Eichenbaum, 2011). This **memory consolidation** process is defined as the progressive postacquisition stabilisation of long-term memory (Dudai, 2004). There are actually two types of consolidation: synaptic consolidation (which occurs within the first minutes to hours after learning) and systems consolidation. The idea of systems consolidation emerged from the study of amnesic patients such as H.M. and the fact that their retrograde amnesia following medial temporal lobe lesions was temporally graded. Systems consolidation involves the progressive transfer of memories from a 'labile' state to a more permanent one, engaging the gradual reorganisation of brain networks supporting memory. One of the models of memory consolidation is the 'standard model' (Alvarez and Squire, 1994; Squire and Alvarez, 1995; Frankland and Bontempi, 2005), presented in Fig. 8. The standard model states that, first, the hippocampus temporarily gathers, organizes and stores information from either sensory input or existing information in cortical areas. Second, multiple reactivations of the stored patterns (possibly during sleep or rehearsal) reinforce the connections between cortical modules involved in the memory, while hippocampal-cortical connections are weakened. Eventually, the memory becomes independent of the hippocampus.




Fig. 8: The standard model of consolidation.

The first stage of encoding is performed by the hippocampus, which gathers information from several primary and associative cortical areas and fuses them in a memory trace. Upon successive reactivations, the hippocampal-cortical network is strengthened. With time, memories are integrated with pre-existing cortical memories and become independent of the hippocampus. In this model, connectivity modifications are rapid and transient within the hippocampus but slow and long-lasting in the cortical modules. From Frankland and Bontempi, 2005.

A possible way by which changes of different speeds and durations might be implemented in neural networks could be weight plasticity and wiring plasticity (Chklovskii et al., 2004, see Fig. 9). Weight plasticity could underlie rapid, hippocampal-dependent learning, whereas wiring plasticity might change during the consolidation process which takes place on a slower timescale. Plausible biological mechanisms that could underlie both types of plasticities are synapse formation and elimination,

dendritic growth, axon remodelling and long-term potentiation and depression (Chklovskii et al., 2004).




Fig. 9: Weight and wiring plasticity.

Each circle represents a neural unit. The illustration of weight plasticity is represented in the **upper panel**: the postsynaptic unit receives inputs from a subset of presynaptic units. Learning causes changes in the weights of the existing connections (as represented by the thickness of connections).

The **lower panel** illustrates wiring changes resulting from learning: in this case, there is no change in the connection weights, but new connections are created and existing connections are deleted after learning. From Chklovskii et al., 2004.

The standard model of consolidation explains the gradual retrograde amnesia effect of medial temporal lobe lesions, but it does not account for the differences between episodic and semantic memory. Another theory of memory consolidation makes this distinction: the multiple trace theory (McClelland et al., 1995; Nadel and Moscovitch, 1997; Harand et al., 2012). According to this view, the hippocampus plays a permanent role in memory storage and retrieval of episodic but not semantic memory. Each time an episodic memory is retrieved, it is re-encoded (possibly with minor variations), leading to multiple traces involving both the hippocampus and the neocortex. Older episodic memories that have been recalled more often are then more widely encoded. These memories are less prone to amnesia but still require at least a partially functioning hippocampus to be retrieved. On the contrary, much like the standard model, semantic memories gradually become consolidated in the neocortex and become independent from the hippocampus. This model is supported by a number of arguments, among which the fact that the number of reports of temporally graded amnesia are actually equivalent to reports of non temporally graded ones. Importantly, the consolidated memory is not the exact replicate of the initial one: the two memories have different characteristics, namely, that hippocampal-dependent memories are context-specific whereas extra-hippocampally represented memories are non-contextual (Winocur et al., 2010a).

It is generally assumed that systems consolidation concerns declarative memory. The consolidation of implicit memory is considered to be restricted to the same circuits as the ones used during learning (Dudai, 2004).

Consolidation was first thought to occur only once, after learning, but it appears that under some circumstances, reconsolidation can occur. **Reconsolidation** is the process by which a long-term memory transiently returns to a labile state and then gradually stabilises (Squire, 2009). It seems to involve other molecular mechanisms than consolidation (Kandel et al., 2014). According to

Eichenbaum and collaborators, however, consolidation and reconsolidation can be merged, as a new memory is never really learned from a blank slate but probably always relies on the reorganization of already learned material (McKenzie and Eichenbaum, 2011). Indeed, when faced with a new situation, the system must somehow choose between encoding the information from scratch and updating an existing memory.

1.2.5.2 Pattern separation, pattern completion and attractor networks

When faced with a given situation, episode, place or item that can be memorised, the brain can either encode the features of this item as totally new, or slightly modify the existing memory of a previously encountered, similar, item. If a new memory were created for each minor modification of a given item, even though long-term memory is said to have a quite large capacity, this would risk saturation of the system. It is not optimal either to consider that two items are similar if they are indeed sufficiently different to necessitate different behavioural responses. Thus, if we consider that a given memory is encoded by the connections between neurons and that the recall of a specific memory is mediated by the firing of a specific network of neurons, two mechanisms seem to be of particular interest: pattern completion and pattern separation. First, the encoding of two different memories must not be stored in the same cellular network to avoid interferences between the two memories during recall. This is called **pattern separation**, defined as the ability to make the stored representations of two similar input patterns more dissimilar in order to decrease the probability of errors in recall (Guzowski et al., 2004). The degree of difference between the two input patterns necessary to consider that the two memories are different probably depends on a wealth of parameters. Studying memory within a spatial cognition framework enables researchers to tackle some questions about these parameters, as we will see in Chapter 5 (in particular, Sec. 5.2.4, p. 91). The second mechanism that is thought to be involved in memory (most probably in its recall phase) is **pattern completion**. It refers to the ability of a network to respond to a degraded input pattern with the entire previously stored output pattern (Guzowski et al., 2004), that is to say, being able to make the entire network of neurons fire following the discharge of a subpopulation of it.

A simple way to implement memory properties in an artificial neural network is the Hopfield network (Hopfield, 1982). It is a recurrent network, in which each neuron is connected to all others, apart from itself. A basic property of Hopfield's associative network is to make the state of the recurrent neural dynamics evolve in a way that minimises a given energy function. Memories are modelled as specific patterns of activity of the network which minimise this energy function. In Hopfield networks, such patterns are learned through Hebbian plasticity. If the stored memories are sufficiently independent from one another, the network can retrieve the full memory when presented with a partial, even noisy, clue (see Fig. 10). Hopfield networks are able to perform pattern completion.




Fig. 10: Hopfield network and pattern completion.

A: Stored pattern. The memory corresponds to a given state of the system where each ‘property’ has a given set of values.

B: Cue given as an input to the system.

C: Dynamic retrieval of the stored ‘memory’.

Adapted from Hopfield, 2007.

The Hopfield network is an example of **attractor network**. Such networks are composed of nodes (e.g., neurons) that can have given values (e.g., firing activity). Among the entire set of possible combination of node values, some are the attractors, stable states where the network does not evolve any more. Surrounding these attractors (still in the state space) are basins of attraction, i.e. states from where the network will eventually converge to a given attractor (Fig. 11).




Fig. 11: Attractors and basins of attraction.

The current state of the system is presented in green. Two attractors are implemented in the network: they correspond to state A and B. Since the current state is in the basin of attraction of state A, the system will evolve to minimise the energy, eventually reaching state A.

Adapted from Poucet and Save, 2005.

Pattern completion, pattern separation and attractor networks are probably not restricted to one type of memory but rather seem central to how the brain processes information (Hunsaker and Kesner, 2013).

1.3 Summary: multiple memory systems

In summary, we saw that the idea of multiple memory systems is currently well established and that different mechanisms at the molecular, cellular and systems level are involved. Memory systems can broadly be divided into three categories:

- i. A procedural, implicit, slow learning and inflexible memory system, in which the dorsal striatum plays a major role.
- ii. A declarative, explicit, rapid learning and flexible memory system, which probably involve medial temporal lobe structures, including the hippocampus.
- iii. Emotional memories, which are probably supported by the amygdala and associated structures.

White and McDonald proposed a classification of memory which highlights these three systems (McDonald and White, 1993; White and McDonald, 2002; McDonald and White, 2013, see Fig. 12). The declarative, procedural and emotional systems can work independently but the accent is put on their parallel functioning. Depending on the form of the information arriving as an input to the system (e.g., stimulus – response association, or stimulus – stimulus association), stronger and more coordinated activity will be generated in the system most compatible with this set of relationships. A highly coherent activity will durably modify a given system (possibly using synaptic plasticity mechanisms such as LTP or LTD), altering the way future information will be processed: this can be considered as memory formation. Importantly, the three systems can interact: a given system can facilitate or impair information processing performed by the others; and the output of these systems can either converge towards cooperative facilitation of the same response or competitive facilitation of different responses (White and McDonald, 2002; White et al., 2013). This view, although schematic, can be applied to spatial cognition, and these different types of memory systems can be probed by spatial tasks, as will be seen in Chapter 3.




Fig. 12: Multiple parallel memory systems.

Types and putative neural substrates of memory systems according to White and McDonald, 2002. Only the central structures for each system are presented but each of them relies on afferent and efferent structures. Adapted from White, 2007.

Memory alone cannot drive behaviour, but the memory of a situation can help in deciding which action to take when confronted with a similar situation. How human and nonhuman animals select actions in the face of choice is addressed by the decision-making domain of research.

Chapter 2 – Decision-making

The spatial cognition and decision-making domains are often investigated separately. Because our work addressed both, this chapter aims at proposing clear definitions concerning the concepts and parameters useful for decision-making that can play a role in spatial cognition.

Decision-making, as defined by Wang (2008), is the process of choosing an opinion or an action among a set of two or more alternatives. It involves the notion of choice, which is not explicitly present in all spatial tasks. Nonetheless, even if only one alternative exists (e.g., going straight in a line), there is always the other alternative of doing nothing, or going backwards. While researchers modulate rats' motivation to make them express the expected behaviour, decision-making is still an unavoidable aspect of spatial cognition.

2.1 Goal

According to Elliot and Fryer (2008), a **goal** is a cognitive representation of a future object that the organism is committed to approach or avoid. Ramnerö and Törneke (2014) quite similarly define a goal as a mental event that represents properties of the future and that can be considered to have causal influence on behaviour. Interestingly, they propose that, because the future has no objective features, goals would rather be 'pre-presentations' than representations. However, one could reply that the dissociation between experienced events and imagined events is rather blurry in the world of representations.

Tolman (1925) also contributed to the definition of the 'goal object' as the object or situation towards which (or away from which) the organism moves. He proposed that two factors come into play when evaluating the strength of a particular goal object: its inherent value (whether positive or negative), and the current physiological state of the organism. He also distinguished between the **ultimate goal object** as a physiological state that a subject persists to attain or eliminate and the **subordinate goal object**, which is the means of getting to or from the ultimate goal object.

Elliot and Fryer (2008) specify that the goal is not restricted to the object of the goal. It also encompasses the approach or avoidance commitment with regard to the object. Indeed, different motivations can drive different goals concerning the same object. Moreover, in this definition, the notion of commitment is important. It relies on the fact that the subject must explicitly commit to the goal. The representation of a future object becomes a goal when the subject commits to reaching this object.

The spatial cognition approach is a convenient way to instantiate goal objects in a space that can be quantitatively described since goals, rather than being abstract, can be embodied by specific places. For example, in the radial arm maze task, the ultimate goal object could be to eat all food pellets available with the least possible physical expense. Subordinate goal objects would be instantiated by

the food cups placed at the end of each arm, combined with the approach commitment when this food cup is selected as the current goal. The term goal in spatial cognition is mainly used to refer to a location in the subject's environment that is associated with reward (Burgess et al., 1997). A few studies addressed the question of spatial goal representation in the hippocampus of rats, and the existence and nature of a goal representation in the hippocampus seems to depend upon the task used (e.g., Hollup et al., 2001a; Hölscher et al., 2003; Jeffery et al., 2003; Hok et al., 2007a). These studies are central to our work and will be reviewed in Chapter 5 (Sec. 5.4.2, p. 100).

Importantly, not all types of behaviour are believed to rely on a representation of the goal. Rats (and humans) can express different types of behaviour when confronted to the same situation. This can underlie the use of different decision-making systems, reflected in the dissociation between habitual and goal-directed behaviour.

2.2 Goal-directed behaviour and decision-making systems

Precise distinctions among different types of behaviour were made in the instrumental learning framework, a domain of research usually relying on tasks in which rats have to press levers to get a reward. These highly controlled tasks allowed to dissociate different decision-making systems that rely on different types of learning (Balleine and Dickinson, 1998; Redish, 2013). A first distinction is made between Pavlovian and instrumental learning. Through Pavlovian learning, also termed **classical conditioning**, a stimulus will elicit an unconditioned response, i.e., a behavioural response that has been learned over an evolutionary timescale. No action is required to get a reward. Still, this system allows the subject to anticipate the reward or a reward-predicting stimulus. The simplest and most famous example of such learning comes from Pavlov, who trained a dog in associating a sound with the arrival of food. At the end of training, the sound was sufficient to make the animal salivate, even if food was not subsequently presented. The fact that the dog salivates when presented with food is a reflex; the fact that this response can be produced following a specific sound is the result of Pavlovian learning. That is, the Pavlovian decision-making system learned to associate the sound with future reward. However, in this type of learning, the subject has no control over the source of reward. On the contrary, **instrumental learning** allows the subject a degree of control over motivationally significant events.

Instrumental learning yields instrumental behaviour that can in turn be classified in two types of behaviours: on the one hand, habitual behaviour, also termed **stimulus-response behaviour**, which relies on a **procedural decision-making** system; and, on the other hand, goal-directed, or **action-outcome**, behaviour, which relies on a **deliberative decision-making** system (Balleine and Dickinson, 1998; Redish, 2013). Interestingly, this dissociation parallels the distinction between procedural and declarative memory systems that we treated in the previous chapter.

Goal-directed behaviour is proactive and influenced by the mental representation of the goal. An action is said to be goal-directed when performance is mediated by the knowledge of the contingency between the action and the goal or outcome (Dickinson and Balleine, 1994). Goal-directed behaviour is different from other forms of behaviours because it relies on the end state that

an operation should achieve, whereas it does not depend on the performance of a specific operation or procedure (Verschure et al., 2014). This is the class of behaviours we will mostly be interested in. In the instrumental learning framework, goal-directed behaviours rely on two subtypes of learning. The first is **contingency learning**, i.e. the acquisition of information about the relationship between the instrumental action and the reward. The second is **incentive learning**, in which the subject associates an incentive value to the reward (Dickinson and Balleine, 1994; Balleine and Dickinson, 1998). Importantly, goal-directed behaviour relies on a representation of the goal (Verschure et al., 2014).

Goal revaluation (Dickinson, 1985) is a paradigm that is largely used to assess the goal-directed or habitual nature of behaviour. Most of the time, goal revaluation consists in outcome devaluation such as satiety devaluation, the same procedure than the one used in the previous chapter to assess the 'what' component of episodic-like memory in Babb and Crystal's study (2006a). **Satiety devaluation** consists in free-feeding a subject with a specific reward to selectively lower the value of this reward. Rats are often trained to press a lever to get a specific reward, and then they undergo reward devaluation. Those rats that are sensitive to devaluation tend to selectively decrease their performance when presented with the lever corresponding to the devalued reward. Rats showing this adaptive response are said to demonstrate goal-directed behaviour. Rats that do not adapt their policy, i.e. that continue to press the devaluated lever with similar performance than a non-devaluated one, demonstrate habit behaviour (Balleine and Dickinson, 1991; Dickinson and Balleine, 1994).

Using this paradigm, Corbit and Balleine (2000) studied the involvement of the hippocampus in goal-directed behaviour. They first tested whether rats that were lesioned in the dorsal hippocampus would still demonstrate incentive learning. Lesioned rats were able to learn a lever-pressing task with similar performance than controls. Following satiety devaluation, their behaviour was still not different from controls: they selectively reduced their responding to the lever associated with the devalued reward (Fig. 13). Another task was used to assess contingency learning. After retraining the rats in the normal paradigm, the response-outcome contingency for one of the levers was degraded. This was done by delivering the same reward than the one obtained by pressing on the lever, but independently from lever pressing performance (i.e., with a given probability per second). In a choice test performed in extinction, control rats showed reduced pressing on the lever with degraded contingency, while hippocampal rats responded similarly to both levers (Fig. 14). These results suggest that the dorsal hippocampus does not seem to be involved in incentive learning, whereas it may be instrumental to contingency learning, i.e. assessing the causal relationship between an action and its outcome.




Fig. 13: Incentive learning in hippocampal rats.

After specific devaluation of a type of food, both controls (**right**) and hippocampal-lesioned rats (**left**) selectively diminish the number of lever presses for the devalued reward in an extinction choice test.

'Error bars' represent the SED for each group.

Adapted from Corbit and Balleine, 2000.




Fig. 14: No contingency learning in hippocampal rats.

Contrary to controls (**right**), hippocampal rats (**left**) did not adapt their behaviour following contingency degradation.

'Error bars' represent the SED for each group.

Adapted from Corbit and Balleine, 2000.

Importantly, in a follow-up study, Corbit and collaborators managed to reproduce both effects with similar, electrolytic lesions of the hippocampus. However, more selective NMDA lesions did not have any effect in the sensitivity to contingency degradation, i.e. hippocampal rats were not different from controls (Corbit et al., 2002). Indeed, electrolytic lesions can impair fibres of passage and this might be the cause of the observed deficit and of the discrepancy between the results from different types of lesions. Further tests based on chemical lesions (either of the subiculum or of the entorhinal cortex, which are output and input structures of the hippocampus, or of efferent fibres of the hippocampus) reproduced the previously observed hippocampal deficit. The authors concluded that the effects reported above (Corbit and Balleine, 2000) actually damaged these efferent fibres and that the hippocampus, contrary to the entorhinal cortex, might actually not be involved in contingency learning. The putative absence of hippocampal role in incentive learning seems clear, as the absence of effect of hippocampal lesions on sensitivity to devaluation was later reproduced (Reichelt et al., 2011).

Finally, it might be worth mentioning that the distinction between stimulus-response and action-outcome learning is also addressed by human decision-making scientists or computational scientists in terms of model-free versus model-based systems. A **model-free** system relies on learned values of situation-action associations; it is efficient but inflexible, because stored action values reflect past experience rather than current goals. **Model-based** systems, on the other hand, rely on internally generated expectations of action outcomes, are more computationally demanding but more flexible (van der Meer et al., 2010; Doll et al., 2014).

2.3 Motivation

In the spatial cognition domain, we usually focus on how a specific behaviour is performed and which mechanisms enable it. The reason why animals perform this behaviour is not the main point of studies. However, without motivation, rats would not complete spatial tasks. What does make a rat explore his environment and not simply stay still? Why do rats push on levers in instrumental training

tasks? Why would a rat want to find its way in a maze? The concept of motivation is central to these questions.

First, one must make the distinction between general motivation (a behavioural state) and motivation for something (reason for engaging in a specific behaviour). The first notion of motivation is that of an all-purpose energy that can be directed towards some destination (Higgins, 2011). However, the ‘all-purpose’ idea can be misleading and Higgins favours the second definition of **motivation** as ‘preferences directing choices’. In other words, motivation is the force that drives actions (Hull, 1943; Redish, 2013). It enables *i*) to set the current goals of the organism and *ii*) to determine how hard one will work to attain these goals. It can be said to emanate from the need to reduce drives. **Drives** can be of two kinds: primary (innate, e.g., thirst, hunger, sex, curiosity, safety, or maintaining body temperature; Hull, 1943; Wise, 1987) or secondary (learned, e.g., money). Drives or needs must not be confounded with the goal itself. Drives work as an energizer of behaviour while goals direct this energy in a way that flexibly serves drives (Elliot and Church, 1997). A given drive (eat) can prompt the use of different goals (go to the restaurant or order a pizza) whereas different needs can prompt the use of the same goal. Drives can also directly affect behaviour without requiring the intermediary of a goal, which is the case for habitual behaviour (McClelland et al., 1989).

In spatial tasks, researchers usually rely on motivation for food and deprive animals so as to prioritise eating over other goals. Water deprivation is also used. However, for a subset of tasks such as the water maze, where the rat must escape out of a pool, no specific deprivation is needed: motivation to get out of the water is sufficient. Other tasks such as free exploration of objects do not need to bias natural motivations because they rely on spontaneous behaviour.

2.4 Value

Value is a central concept in decision-making, as it has a major impact on action selection. However, it can be applied to different parameters of decision-making, namely, outcome, goal, decision or action (Peters and Büchel, 2010). The **outcome value** refers to the value of a reward upon consumption. It is unrelated to the cost associated with getting the reward. **Goal value** does not include the costs either. Rather, it refers to the value of a stimulus in a more abstract currency than outcome value. As an example, in humans, it is evaluated by asking how much money one would be willing to spend in order to get a given reward. Goal and outcome value are highly correlated but are supposed to be represented differently. **Decision value** stands for the net value of a decision. It includes costs, which can be delays, effort, and perhaps distance in spatial tasks. Finally, **action value** refers to the pairing of an action with a particular type of value, whether outcome, goal or decision. Prior to choice, values are assigned to the available actions depending on the complexity of the decision context and these value actions are then compared in order to execute the action with the highest value.

In a seminal series of studies, Schultz and collaborators (1997) shed light on the reward-related coding of dopaminergic neurons from the ventral tegmental area and the substantia nigra (Fig. 15).

An instrumental conditioning paradigm was used in which monkeys were trained to touch a lever following the appearance of a light in order to get a reward in the form of fruit juice. At first, dopaminergic neurons discharged following reward delivery. Once the monkey mastered the task and reached the lever as soon as the light was illuminated, the increased firing of dopaminergic neurons took place at the onset of light, and no more at reward delivery. Yet more interestingly, if the reward was not delivered at the predicted time, the neurons decreased their spontaneous firing. Thus, these dopaminergic neurons did not only code for the reward itself but rather for an error of prediction in the reward, i.e., a difference between the expected reward and the real outcome, increasing their firing if the obtained reward was higher than expected and decreasing firing when the outcome was less valuable than expected. Such prediction error signal is a major actor of decision-making models as it enables to update the expected value of a goal and then influences action selection (Sutton and Barto, 1998).




Fig. 15: Reward prediction error coding of a dopaminergic neuron.

For each inset, the upper plot represents the cumulated firing of the neuron with respect to time, where firing is aligned with either reward occurrence (R) or light onset (CS). Each line under this histogram stands for the firing of the neuron during one trial (raster plot, where each dot is a spike).

Top: after learning, when no light predicts the reward, the neuron discharges at the (unpredicted) occurrence of reward.

Middle: when the light predicts the reward, the neuron discharges just after light onset and no more at reward occurrence.

Bottom: after learning, if no reward is delivered, the firing is depressed at the expected time of occurrence of reward.

From Schultz et al., 1997.

Value-related signals of different types were also found in the ventral tegmental area of rats (Roesch et al., 2007) and the dopaminergic and noradrenergic neurons of monkeys (Tobler et al., 2005; Bouret et al., 2012). Neural representations of different categories of value can be found in the activity of neurons from several other brain regions. In particular, the dorsal striatum (a structure of the basal ganglia) and the orbitofrontal cortex (a portion of the frontal cortex) both contain neurons that seem to represent certain types of value, in rats (in the striatum: Bissonette et al., 2013, Lavoie and Mizumori, 1994; Kim et al., 2009; Roesch et al., 2009; Howe et al., 2013; in the orbitofrontal cortex: Sul et al., 2010) and primates (in the striatum: Yamada et al., 2013; in the orbitofrontal cortex: Padoa-Schioppa, 2007). As far as the hippocampus is concerned, its involvement in goal value coding is, to date, controversial, as some studies have not found any value-related signal (e.g., Tabuchi et al., 2003) but a recent study found such correlates in a subset of hippocampal neurons (Lee, Ghim, et al., 2012). The issue of representation of goal value in the hippocampus is one of the

main questions of our work and will be addressed in Sec. 5.4.3 (p. 105) and in the experimental part (Sec. 8.3, p.165).

In contrast to value, which is usually assessed on a quantitative scale, **valence** is qualitative. It is the motivational direction of an object and it can be either positive, triggering attraction, or negative and causing repulsion (Higgins, 2011). The qualitative aspect of valence does not mean it is all or none: a stimulus can have high or low, positive or negative valence.

2.5 The exploration / exploitation ratio

In models of decision-making, a frequent way to describe the behaviour of subjects confronted to choice is the exploration / exploitation ratio. It refers to the balance between choices that allow gathering new information about the environment and choices that exploit the already acquired knowledge to maximise the outcome. Exploitation allows to maximise the expected reward on a short timescale, but exploration might be more rewarding on the long term (Sutton and Barto, 1998). The exploration / exploitation ratio of subjects can be assessed in tasks such as the n-armed bandit task, where n actions are possible, each leading to a given reward (Sutton and Barto, 1998; Daw et al., 2006). Since this ratio can be quantitatively measured in humans as well as nonhumans, it is a proper tool to compare results across species and possibly gain insight into the neural bases of inter-individual differences in decision-making. This ratio can be modulated by multiple parameters such as stress or satiety (Luksys and Sandi, 2011). For example, a higher level of satiety would induce a more exploitative behaviour (see Inglis et al., 2001, for a modelling approach).

2.6 Vicarious trial-and-error as a marker of deliberation?

Few behavioural correlates of decision-making are accessible in the rat. Vicarious trial-and-error, or VTE, is thought to be one of them. Muenzinger (1938) and Tolman (Tolman, 1938, 1939) first referred to VTE as a conflict-like behaviour expressed at choice points, consisting in head movements oriented towards potential options. Tolman (1939) observed that rats produced more VTE when confronted to more difficult problems, in his case, colour discrimination between a white door and either a black, medium grey, or light grey door. The hardest the discrimination, the more rats showed a VTE behaviour.

A few studies tried to confirm this link between VTE and problem complexity, as well as the link between VTE and hippocampus. VTE behaviours decrease in hippocampal rats, which learn a discrimination task more slowly than controls (Hu and Amsel, 1995). Moreover, on hippocampus-dependent tasks, the amount of VTE is correlated with improved performance and with hippocampal activity, as measured by cytochrome oxydase, an index of neuronal metabolic activity (Hu et al., 2006). We will indeed see that neural activity corroborates the hypothesis of VTE as a marker of deliberation over potential options, materialised as paths in spatial tasks (Sec.5.3.3.2, p.96).

2.7 Conclusion: decision-making in spatial behaviour

Applied to spatial behaviour, at least two different levels of decision-making can be determined: first, when an animal enters an environment, it is confronted to the exploration / exploitation dilemma; it can either explore to gather knowledge, or exploit, provided it already knows where to find reward in the environment. Second, if the animal 'chose' (perhaps in an implicit manner) to exploit, it must decide which set of rules to follow so as to attain the goal (termed 'policy' in the reinforcement learning framework). In the spatial cognition domain, these different policies can be said to correspond to the possible navigation strategies that a subject can use. Different navigation strategies rely on different neural systems. Moreover, as we will see in Chapter 3, a parallel can be drawn between navigation strategies and decision-making systems, i.e., habitual versus goal-directed behaviour.

In addition to the exploration / exploitation dilemma, other parameters can influence decision-making in spatial tasks. The value of the expected outcome associated to a given goal is important in the computation of goal value and in the final selection of a goal. Space can also be a parameter involved in the computation of the action value in the form of a cost. If two spatial goals are available, but one is farther away, the subject will be more likely to choose the closest goal, all else being equal. Spatial cost is not only about distance: it can also be implemented under the form of obstacles that require climbing (e.g., Hillman and Bilkey, 2010).

In the next chapter, we will focus on spatial cognition, a domain that combines memory – in the form of spatial memory, or memory of the task – and decision-making – in the form of navigation strategies. Moreover, spatial cognition allows for evaluating the whole perception – action loop, by assessing which type of information is used by animals to navigate, and how different neural systems can use this information to select the action.

Chapter 3 – Spatial cognition

3.1 Perception of space

Space is a notion to which we are confronted daily. However, as an abstract concept, it is quite hard to define. From the Encyclopaedia Britannica, **space** is the boundless three-dimensional extent in which objects and events occur and have relative position and direction. One does not ‘directly’ perceive space; it is much easier to consider the objects within space, or one’s movement through space, than absolute space by itself. Without going into the debate of the existence of absolute space and whether the perception of space is innate or learned (issues detailed in the introduction of O’Keefe and Nadel, 1978), we will focus on those parameters of space suited to self-localise and navigate, and on what sensory channels convey this spatial information.

3.1.1 Types of sensory information

To perceive space, animals need to interact with their environment. Information gathered through this interaction flows through different sensory channels and is generally classified as either idiothetic or allothetic. **Idiothetic inputs** correspond to self-motion related signals, i.e. generated by one’s movement. It encompasses vestibular, proprioceptive, motor command efferent copies, and sensory (e.g., optic) flow information. **Allothetic cues** inform about the external environment and can be conveyed through visual, olfactory, auditory, and somatosensory channels (Arleo and Rondi-Reig, 2007). A given sensory modality can provide both types of information: for example, vision can transmit allothetic information about static environmental landmarks as well as idiothetic cues through the optic flow generated during self-motion.

Idiothetic signals are always available (e.g., even in darkness) and are sometimes sufficient for an animal to estimate distance and orientation parameters. In the example given in Fig. 16, rats were moving around a circular table-top in either light or complete darkness conditions (Wallace & Whishaw, 2004). Although their speed was lower in the dark, rats managed to head to their departure point with a precise direction in both conditions. In addition, rats demonstrated knowledge of the distance to the goal, as their speed significantly decreased at the midpoint of the homeward trip, regardless of the length of the trip. In this case, both direction and distance controlled the trajectory, independently from the availability of allothetic information (Wallace & Whishaw, 2004). The ability of animals to keep constant track of their position with respect to a departure point is termed **path integration** and will be detailed in the section dedicated to navigation strategies (Sec.3.2.2, p.39).




Fig. 16: Knowledge of distance and direction in homing behaviour.

Velocity (**left**) and path (**right**) are indicated for outward and return trajectories with respect to a departure point (square) either in light (**top**) or dark (**bottom**) conditions. Note the precise occurrence of peak velocity at the middle of the return path, even in darkness, indicating knowledge of direction and distance to the goal.

From Sutherland and Hamilton, 2004, originally adapted from Wallace and Whishaw, 2004.

Although rats are able to use idiothetic signals to navigate, they usually rely heavily on allothetic cues when available. In general, the visual modality is the most used but olfactory and tactile signals can also help localisation. Indeed, these other sources of information can improve performance when visual cues are less salient (e.g., olfactory-based navigation in the dark, Lavenex and Schenk, 1998; cooperation of olfactory and vision, Rossier and Schenk, 2003; auditory cues, Rossier et al., 2000; see Jacob, 2014 for a recent review). However, landmarks can be unstable and allothetic cues might sometimes not be sufficient to disambiguate two similar environments (such a situation is probably more likely to happen in laboratory conditions). In natural situations, animals combine allothetic and idiothetic signals to navigate, depending on the reliability of the available cues, a process named multisensory integration (Arleo and Rondi-Reig, 2007).

3.1.2 Multisensory integration

To assess the relative contribution of each type of information to self-localisation, a common paradigm consists in causing a conflict between different sensory sources. For instance, in the experiment by Etienne and collaborators (1990), hamsters first learned to go from their nest to a feeder located in the middle of a circular arena, of diameter 220 cm, by following a baited spoon directed by the experimenter. Once there, the hamsters filled in their cheek pouches with food and came back to the nest (Fig. 17 A). During training, a light spot was presented at the opposite side of the nest. During the test, this visual cue was rotated by either 90° or 180°. If hamsters relied on idiothetic cues only, they would still directly return to their nest. If they relied on the visual cue only, they would aim at the spotlight. What happened is that they did neither one nor the other, but they seemed to combine visual and allothetic cues (as their final position was intermediary between the one indicated by self-motion cues and the one deduced from the visual cue). Interestingly, the deviation from the actual nest position was not proportional to the rotation of the visual cue: the highest deviation was obtained for the 90° rotation (Fig. 17 B & C). In this condition, the visual cue was given a larger weight than idiothetic cues for the estimation of position. Under the 180° rotation condition, conversely, the main source of information used to perform homing came from self-

motion signals, as the final position of the hamsters was closer to their nest. Thus, navigation relies on a weighted multisensory integration process, in which cues contribute roughly in proportion of the confidence that can be attributed to them.




Fig. 17: Hamsters perform multisensory integration.

The homing vector represents the mean return trajectory performed by a population of hamsters. Adapted from Arleo and Rondi-Reig, 2007; originally adapted from Etienne and Jeffery, 2004.

Other species, such as ants for example, can also associate idiothetic (odometry) and allothetic (polarisation pattern of natural light) information to return to their departure point after an outward trip (Wehner, 2003). There are several other examples suggesting that animals (both humans and nonhumans) perform multisensory integration to navigate, by combining either allothetic and idiothetic cues or sensory sources within a single type (Berthoz and Viaud-Delmon, 1999; Rossier et al., 2000; Etienne and Jeffery, 2004; Arleo and Rondi-Reig, 2007). There is a general preference for more precise sensory sources such as vision over olfactory or self-motion cues (Maaswinkel and Whishaw, 1999).

Concerning the visual modality – the most studied one – this integration is performed according to the relative reliability that can be attributed to each set of cues: rats will only use a given set of visual cues for navigation if they are stable, at least relative to other cues (Biegler and Morris, 1993; Biegler, 1996). Visual cues are often separated in proximal versus distal cues. According to Knierim and Hamilton (2011), **distal cues** are the ones present on the walls of the laboratory or otherwise removed from the behavioural apparatus, whereas **proximal cues** are part of the apparatus itself. Although this definition focuses on laboratory environments, it has the advantage of clarity. Distal cues are the ones most preponderantly used for navigation, but the importance of proximal cues must not be neglected. More precisely, distal cues seem important to provide a general sense of direction to the navigating subject while proximal cues are used to specify precise locations (Knierim and Hamilton, 2011). We note, however, that the frontier between proximal and distal cues can be blurry when one tries to generalise to natural environments. Some authors consider proximal (or

local) cues as being items that are concurrent with the goal position (Morris, 1981). In this sense, cues on the walls of the experimental apparatus, such as a cue card, would be considered distal.

Rodents can perform navigation using a combination of allothetic and idiothetic information, in a flexible manner, since the relative importance given to sensory sources depends on their availability and reliability. Thus, one of the essential characteristics of a neural correlate of spatial representation, if it exists, should be to result from multisensory integration and to rely on stable cues.

3.1.3 The importance of geometry

An environment is said to be anisotropic if it has different properties according to the direction, as a rectangle, and unlike circular arenas. Experimental evidence points towards the fact that the geometry of anisotropic environments can be used for navigation. The geometrical arrangement of items can also be used as an orienting cue. Cheng (1986) postulated the existence of a brain module dedicated to geometry that would be impenetrable to information about local features. He trained rats to locate food in the corner of a rectangular environment (see Fig. 18). The walls of the environment had distinctive colours, odours and texture. However, when disoriented before being put back in the environment, rats often made rotational errors, i.e., they searched for food in the opposite corner compared to where the food was actually located. These errors were almost as numerous as correct responses. Thus, geometrical cues (shape and distance properties) seemed to dominate over feature-related cues (such as odour, colour, or texture). The author proposed that this behaviour could be explained by the existence of an encapsulated geometric module in the brain. The **encapsulated** nature of this putative brain process means that it would represent the surface geometry of the environment in a totally independent manner with respect to other information.




Fig. 18: Geometrically ambiguous apparatus

The panels at each corner of the rectangular apparatus have different visual, tactile and olfactory characteristics. The filled circle stands for the location of the hidden food. On geometry grounds alone, the open circle (wrong location) cannot be distinguished from the true goal location. In this task, disoriented rats showed high rates of search both at the correct and the rotationally equivalent location. From Cheng, 1986.

The geometric module proposal arose the interest for the encoding of geometry by animals (humans and nonhumans). Arguments were brought to support the idea that geometry information was indeed preferably used for navigation over the specific features of individual objects (Gallistel, 1990, Hermer and Spelke, 1996 in young children, Benhamou and Poucet, 1998 in rats). Others argued against the specific existence of a geometric module, stating that snapshot views could explain the results (Cheng himself, 2008; Sheynikhovich et al., 2009). The original hypothesis had to be revised, stating that, on the one hand, a geometric system could be used by an animal to specify its position

(and orientation) and, on the other hand, landmarks and patterns were used to locate items such as a goal (Lee and Spelke, 2010). Moreover, the configuration of an array of items can also be used if it helps polarising the environment. A recent study tested the ability of mice to navigate using either specific items features or the geometric arrangement of these items (Fellini and Morellini, 2011). The authors concluded that mice could use the geometrical arrangement of cues if it was non ambiguous. However, if mice could only rely on the features of particular cues, they were impaired. This had previously been demonstrated in the rat (Benhamou and Poucet, 1998).

Overall, geometric information seems to be used as a backup system after disorientation, or if it is the most salient cue that enables to polarise the environment (Burgess, 2006; Knight et al., 2011). However, the ‘geometric module’ would not be strictly encapsulated as animals can combine both geometry and featural cues when not disoriented (Maurer and Derivaz, 2000). Anisotropic geometry of the room, the experimental apparatus, or the arrangement of cues can be used for localisation. Actually, what Cheng called geometry might be related to landmarks configuration in natural environments (Benhamou and Poucet, 1998; Skov-Rackette and Shettleworth, 2005; Sutton, 2009). Distal landmarks change relatively little as an animal moves and thus provide an orienting framework. Proximal cues, on the contrary, are more prone to instability with respect to the animals’ motion and might be encoded in a different way (Poucet, 1993; Skov-Rackette and Shettleworth, 2005). Nowadays, the importance of boundaries of an environment has found echoes in the neural activity of neurons, as we will see in Chapter 4 (Sec. 4.2.3, p. 64; Burgess, 2008).

Lesion studies showed that the hippocampus might be involved in shape-based navigation (McGregor et al., 2004), in particular, evaluating that a goal is at a certain direction and distance from a boundary (Horne et al., 2010).

3.1.4 Spatial reference frames

It is commonly accepted that space, or one’s position, or external world items, can be represented in two different manners: in an **egocentric framework**, i.e., with respect to the subject, and/or in an **allocentric framework**, i.e., with respect to an element of the external world, for example, the magnetic north, or a distal landmark (Arleo et Rondi-Reig, 2007; Burgess, 2008; Hartley et al., 2014). Manipulations of different parameters (such as departure position in a maze) can help determine which reference frame is used. Notably, it was shown that rats are able to use different spatial reference frames simultaneously (Fenton et al., 1998; Sutherland and Hamilton, 2004).

3.2 Navigation strategies

A **navigation strategy** can be defined by a set of rules to follow when in a given situation in order to reach a spatial goal. Navigation strategies can involve spatial information processing of different degrees of flexibility and complexity. For example, turning left at the green sign is a response strategy whereas going to a specific place defined by its relationships with surrounding cues is a place strategy. There are different, but similar, ways to categorize strategies (e.g., O’Keefe and Nadel, 1978; Gallistel, 1990; Redish, 1999). The one we will rely on is presented on Fig. 19 (Arleo and Rondi-Reig, 2007).

To these navigation strategies, we will add exploration and the path integration strategy.




Fig. 19: Taxonomy of spatial navigation strategies.

From Arleo & Rondi-Reig, 2007.

3.2.1 The importance of exploration

One of the first things that rodents will tend to do when confronted with novelty is to explore the source of novelty. This behaviour consists in moving towards unknown places or objects and collecting different types of information from several sensory modalities, for example, visual (through rearing), olfactory (via sniffing) or tactile (using whisker contact). Exploratory behaviour diminishes with habituation, albeit in different ways depending on species (Poucet et al., 1988). Interestingly, this diminution seems to depend on the integrity of the hippocampus (Roberts et al., 1962; Leaton, 1965; Save et al, 1992). The cerebellum seems also involved in exploratory behaviour (Caston et al., 1998). Exploration is fundamental in spatial cognition as it would be used to gain spatial knowledge and build representations of environments (Poucet et al, 1986; O'Keefe and Nadel, 1978). It is a form of **latent learning**, which refers to the acquisition of knowledge occurring in the absence of explicit reward and without generally observable changes in behavioural performance (Johnson and Crowe, 2009).

3.2.1.1 Organisation of exploration

Whether it is a new environment, a new object in a well-known environment (Save et al, 1992), the new spatial arrangement of objects (Wilz and Bolton, 1971, Save et al, 1992), or even a change in the topology⁵ of an environment (Alvernhe et al., 2012), many types of perceptible novelty will trigger exploration. This will be the case even in a paradigm where a specific task must be performed and can sometimes interfere with the interpretation of results. It seems that novelty detection acts on the current goal of the animal, prioritising the gathering of new knowledge over feeding or other

⁵ The **topology** concerns, basically, the connectivity between different places in an environment.

behaviours. In the decision-making framework, the simplest way to model exploration (in the exploration / exploitation ratio sense, see Sec. 2.5, p. 29) is to randomly choose between available choices (but see Arleo and Gerstner, 2000, for a different approach). However, behavioural studies demonstrate that it is actually quite organised (Avni et al., 2006; Fonio et al., 2009) while still enabling the expression of inter-individual differences (e.g., Guillette et al., 2009). Basically, when exploring a new environment for the first time, a rodent will make excursions from its departure point to unexplored parts of its environment, most often following the borders, and regularly returning to a place termed ‘home base’ (Eilam & Golani, 1989; Golani, 1993; Draï et al., 2000). The **home base** is operationally defined as the place where an animal spends a disproportionate period of its time and from which it makes excursions (Whishaw et al., 2006). There can be several home bases. Specific behaviours such as rearing or grooming are more likely to occur at the home base (Eilam and Golani, 1989). The home base is usually the place where the animal was first released in the environment, above all if close to salient cues (Nemati and Whishaw, 2007), but not if it does not provide sufficient shelter (Whishaw et al., 2006). Whishaw and collaborators observed rat’s observation in wide environments such as a parking lot. They suggested that exploration would mainly serve to optimise safety. Exploration has similar patterns in the absence of visual cues: in the dark, rats placed in a new environment will still organise their displacements around a chosen home base. Their displacements show invariant characteristics, e.g., a dissociation between the outward trajectory and return trajectory (Wallace et al., 2006).

More recently, a thorough characterization of mice exploratory behaviour was performed by Fonio and collaborators (2009). They demonstrated that exploration of a new and large circular arena could be broken up in several behavioural patterns, the order of which was highly reproducible among individuals. These behavioural patterns progressively involve increasing spatial dimensions (see Fig. 20). In summary, mice first make short back and forth trips (1-D) from their home base following the wall. Once they complete a full turn, they begin making small incursions inside the environment (2-D) that progressively become independent from the home base. They end up performing jumping movements (3-D).




Fig. 20: Free exploration sequence in mice.

Developmental sequence of free exploration across a 3 h period in a 2.5 m diameter arena. The black dot indicates the home-cage (here chosen as a home base by the mouse) and red lines indicate mouse trajectories (yellow stands for the return trip in home-related shuttle). This order of exploratory behaviours is highly reproducible from individual to individual. The data is from the neophobic strain of mouse (BALB/c) but the classic C57BL/6 strain mice show similar exploration patterns, albeit with more inter-individual variability.

Adapted from Fonio et al., 2009.

The authors highlighted the fact that in their experiment, exploration was free: the departure point of exploration trips was the home cage of the mouse, where *ad libitum* water and food was provided,

and the time left for exploration was quite long (45h in total). In usual rodent experiments, exploration is forced and constrained in time and space, which might explain why the full pattern of exploratory behaviour is not usually observed. The importance of environmental limits (and probably geometrical information) is evidenced by the necessity for the mice to first entirely explore the borders before performing incursions towards the centre. In addition to providing shelter, borders and geometrical layout, as seen previously, probably serve as anchorage points to build the map of an environment.

3.2.1.2 Object exploration

As previously mentioned in Chapter 1, the spontaneous exploration of objects can be seen as another proof that nonhuman animals most certainly memorise and manipulate representations of space and objects in space (e.g., Thinus-Blanc et al., 1987 in hamsters; Ennaceur and Delacour, 1988 in rats). Indeed, the selective exploration of new objects in a known environment can only be possible if one has stored the arrangement of objects in this environment and is able to compare the current layout with the memorised representation. Many studies rely on spontaneous exploration to assess the memory for the nature or the position of objects, which relates to the ‘what’ or ‘where’ aspects of episodic-like memory (Dere et al., 2005; Eacott et al., 2005). The hippocampus appears to be involved in processing spatial memory in the case of object exploration (Save et al., 1992).

3.2.1.3 Link between exploration and performance

Interestingly, Olton and collaborators showed the importance of exploration (also termed ‘shaping’ in that context) prior to testing. Rats that were not given the opportunity to explore a radial arm maze before testing did not perform better than chance in the task (Olton et al., 1977).

Exploration, even in the absence of food, or pre-exposure, seems necessary for proper performance in navigation tasks (Olton et al., 1977; Ellen et al., 1982; Sutherland et al., 1987; Chai and White, 2004; and Hamilton et al., 2002 for studies with comparable results in the rat and human). As an example, Chai and White tested rats in their ability to discriminate neighbouring locations in a radial arm maze (Chai and White, 2004). In this task, rats were confined to a specific arm of the maze, where they could either find food or not. When later tested with a free choice between adjacent arms that include the food-paired arm, rats demonstrated preference for this arm only if previously exposed to the entire maze (Chai and White, 2004; Gaskin et al., 2005). In the other case, it seems that the knowledge acquired when restrained in an arm was not sufficient to build a representation of the environment and of the spatial configurations of the maze arms. Interestingly, if the dorsal hippocampus of rats was temporarily inactivated during the pre-exploration phase, learning was not impaired, which seems to be in opposition with the current view that the hippocampus is needed to build the spatial representation of the environment. The authors postulated that the learning performed during exploration might be supported by extra-hippocampal, cortical, structures. Another interpretation could be that the remaining hippocampus was sufficient to learn the structure of the environment (see Moser et al., 1995).

Consistently with a role for exploration in building a representation of space, in complex environments, exploration is not homogeneous: rats spend more time exploring the topologically relevant parts of a maze (i.e., the intersections), probably reflecting encoding of information on the connectivity layout of the environment (Poucet and Herrmann, 2001; Alverne et al., 2012).

3.2.2 Path integration

Path integration is the ability of an animal to integrate its own translations and rotations along navigation in order to maintain an estimate of its position relative to a departure point. This mechanism is schematically illustrated in Fig. 21. Contrary to path reversal (see Fig. 22), path integration enables the animal to create a return vector that provides the shortest path to a departure point. This return vector, which is also called homing vector when the departure point is a home base, will be straight even if the outward trip was tortuous. The path integration strategy is used when an animal relies on the path integration mechanism to navigate.




Fig. 21: Path integration.

Path integration from B to A after an indirect trajectory from A to B. S_{1-3} represent the length of segments 1 to 3 of the outbound journey; ϕ_{1-3} represent the corresponding head direction. Path integration consists in summing the displacement vectors in order to produce a direct return trajectory.

Adapted from McNaughton et al., 2006, which was adapted from Mittelstaedt and Mittelstaedt, 1982.




Fig. 22: Path reversal and path integration.

Illustration of the difference between the indirect trajectory produced by path reversal (dotted line) and the homing vector resulting from path integration (dashed line) after a trajectory from A to B (full line).

From Arleo and Rondi-Reig, 2007.

Path integration is thought to play a role in homing behaviour, when an animal must go back to a departure point after some time. It can be used in the absence of allothetic signals. However, because the path integration process is cumulative, errors in the position estimate are bound to accumulate (Etienne and Jeffery, 2004). An example is provided in Fig. 23. It sums up results from experiments using subjects from different species. All of them were required to navigate without allothetic cues from a departure point. All species, including humans, demonstrated an error in the homing vector.




Fig. 23: Errors in idiothetic-based navigation.

Subjects from three different species had to follow a fixed path from S to P indicated by the full line, then, attempt to return to S on their own. All three species were deprived of visual, auditory, olfactory and tactile spatial references from the environment throughout the excursion. The heavy arrows (for spiders and hamsters) indicate the direction of return when the subject was at a standard distance from P. The heavy line with a dot (humans) indicates the return vector. Numbers indicate relative distances.

Adapted from Etienne and Jeffery, 2004.

Path exploration can also be used without prior knowledge on the environment, as is the case during exploration. It would be the mechanism supporting the direct trips back to the home base (Wallace et al., 2006, see Fig. 16, p. 32). Such return trips could be used to correct the accumulated error by recalibrating the self-localisation system (Etienne and Jeffery, 2004). Path integration is believed to work automatically and continuously and to depend on hardwired rules of information processing (Etienne and Jeffery, 2004). It can be expressed by a variety of species, including rats, mice, ants and hamsters (Elduayen and Save, 2014; Etienne and Jeffery, 2004; Müller and Wehner, 1988).

3.2.2.1 Which definition for path integration?

It is often implicitly assumed that path integration only relies on idiothetic cues. However, it is possible to use allothetic information to update (i.e., adjust) or recalibrate (i.e., reinitialise) the estimation of position. Actually, as mentioned by Poucet and collaborators (2014) and others (Fenton et al., 1998), the term path integration is used under two different meanings in the literature:

- i. the ability of an animal to compute, step by step, the sum of a sequence of displacement vectors from a departure point so that a return vector can be computed at any time, enabling a direct return to the departure point; or
- ii. any form of navigation in which the animal's location is updated on the basis of self-motion information alone.

The first definition emphasises the way the position is computed (by using geometric summation of displacement vectors) without excluding that path integration could rely on other cues than purely idiothetic ones. The second definition stresses the fact that the navigating animal uses only idiothetic signals to navigate, but the way position is computed does not matter. As previously stated, we adopted the first definition of path integration.

3.2.2.2 Role of the hippocampus in path integration

There seems to be a debate around the neural bases of the path integration based strategy. We note that this might be related to the above-mentioned differences in definition. In particular, the

involvement of the hippocampus in this process is not so clear. A series of experiments performed by Whishaw and collaborators (e.g., Maaswinkel et al., 1999, Whishaw et al., 2001) showed that the hippocampus was necessary for idiothetic-based path integration. However, these studies relied on lesions of the fimbria/fornix, one of the main hippocampal input/output pathways, but which also includes fibres from other structures. Another study, using pharmacological lesions of the entire hippocampus, did not show any impairment of hippocampal rats in path integration (Alyan and McNaughton, 1999). However, subsequent studies reproduced a deficit of path integration in rats with hippocampal lesions: this was the case either following lesions of the entire hippocampus (Wallace and Whishaw, 2003; Kim et al., 2013) or lesions of the dorsal hippocampus (Save et al., 2001). As an example, Wallace and Whishaw (2003) tested the homing behaviour of rats under light or dark conditions. Rats with entire hippocampal lesions demonstrated very different paths from controls, specifically during the homeward trips: instead of being as direct as those of controls, hippocampal rats' trajectories were circuitous and did not cross the centre portion of the table, whether in light or dark conditions (Fig. 24).




Fig. 24: Trajectories of rats performing homing behaviour in the dark.

Outward trajectories were similar between groups, but homeward trips were different: direct for controls, they followed the boundaries of the arena in hippocampal rats. The black square represents the normalized home base location.

Adapted from Wallace and Whishaw, 2003.

Although the outward trips were similar between controls and lesioned rats, one could argue that these results were not caused by a deficit of path integration but rather by a disruption of the motivation of subjects. Indeed, the ventral hippocampus is known to be involved in motivational processes. It seems that even dorsal lesions of the hippocampus affect the general behaviour of rats. When tested in path integration tasks, they usually express a wandering-like behaviour, as if they

were not aiming at reaching the goal anymore (E. Save, personal communication). This could either underlie a motivational problem or maybe the forgetting of the goal itself.

Actually, the hippocampus might not be the only structure involved in path integration and extrahippocampal structures seem to contribute to this process, such as the parietal cortex and the entorhinal cortex (Save et al., 2001; Parron and Save, 2004; Van Cauter et al., 2013). According to this last study, the medial, but not lateral, portion of the entorhinal cortex would be involved in path integration. Interestingly, the cerebellum seems also to be involved in idiothetic-based path integration (Rochefort et al., 2011; Passot et al., 2012). Overall, it seems that the apparently simple mechanism of path integration involves a variety of structures, each probably involved in a different processing along the way from sensory input to action. In our definition, path integration is closer to a skill and probably relies on procedural memory. However, it requires an estimation of one's movements, which is probably where the cerebellum might be involved, for example via the interpretation of efference copy signals (Pynn and DeSouza, 2013). Finally, the estimation of position, or return vector, resulting from path integration is likely to be used in the building of more declarative representations of space, which might be where the hippocampus plays a role (Etienne and Jeffery, 2004; McNaughton et al., 2006).

3.2.3 Guidance strategy

In certain navigation situations, the goal is either directly visible or cued. In that case, the best strategy, or at least the less cognitive demanding, is to orient towards the goal and approach it. This type of strategy is termed target approaching (when the goal itself is visible) or beacon approaching (if a cue is located at the goal position), or more generally cue, guidance or taxon strategy. It only requires learning of a single stimulus-response association.

Contrary to most functions described here, it is generally accepted that the hippocampus is **not** involved in guidance strategy, or at least that hippocampal lesions do not impair performance in cue-guided tasks (e.g., Morris et al., 1982; Jarrard et al., 1984; Rasmussen et al., 1989; Packard and McGaugh, 1992). The ability for hippocampal-lesioned rats to perform a guidance strategy is often used as a control for non-spatial aspects of behaviour (e.g., sensory or motor abilities).

3.2.4 Response strategy

If the goal is neither visible nor directly cued, but it can be reached by means of associations between elements of the environment and actions (each association being independent from the others), a response (or stimulus-triggered response) strategy can be used. It has also been termed egocentric strategy (Kesner et al., 1989). As an example, this strategy can be employed in the cross maze task (Fig. 25), when food reward is located at the end of one arm (say, the east arm) and a rat is starting at the south arm needs to reach it. Turning right (action) when faced with the intersection (stimulus) will be sufficient to reach the goal. A sequence of response strategies can also be used to solve the starmaze task (Rondi-Reig et al., 2006), a sort of five-arm maze filled with water (Fig. 26). A rat or a mouse is generally released from one of the arms and must find a goal located in another arm. If

specific visual patterns are present as proximal cues on the walls, a response strategy can be used to solve the task by associating a movement (left or right) to each proximal cue.




Fig. 25: The cross maze.

From Passino et al., 2002.




Fig. 26: The starmaze.

Adapted from Passot et al., 2012.

Similarly to the cue strategy, hippocampal lesions do not impair performance when the response strategy can be used to navigate towards a goal. The striatum is likely to be one of the structures involved in this strategy (e.g., McDonald and White, 1993; Jacobson et al., 2012).

3.2.5 Route strategy

When specific actions can be associated to specific states as in the response strategy, but that the knowledge of the state is not sufficient to select the action, one can use a route strategy. The route strategy has also been termed sequential egocentric strategy or sequence-based navigation (Rondi-Reig et al., 2006; Babayan, 2014). It relies on a sequence of stimulus-response actions and it can also be used in a version of the starmaze in which neither proximal nor distal cues are present (Fig. 26). Then, to go from the departure point to the goal, the subject is confronted to three identical intersections. The action to be performed at each of these intersections is different. A sequence of stimulus – action responses must be learned and each choice must be taken according to its position in the sequence. We note that a route strategy is more complex than a succession of cue and response strategies, because the order of the sequence is important.

Many structures are likely to be involved in this strategy, which holds a sequential (and possibly a timing) component. The CA1 field of the hippocampus would be one of the structures involved, along with cortical structures (Rondi-Reig et al., 2006); the reader is referred to Babayan (2014) for a review on sequence-based navigation.

3.2.6 Place strategy and the cognitive map theory

The strategy which probably requires the highest level of spatial information processing is the place (or map-based) strategy. It consists in localising the goal and oneself using the spatial relationships between elements of the environment. Contrary to the response strategy, it enables flexible

behaviour, i.e., adaptability in the face of environmental changes. It was postulated to rely on a ‘cognitive map’, as defined by Tolman (1948).

3.2.6.1 *The cognitive map theory*

Tolman (1948) proposed that these animals could manipulate representations of their environment and that they were not simply stimulus-response “machines”, in contradiction with the general view of behaviour at that time. Namely, he suggested that rats could rely on a **cognitive map** to navigate, i.e., a neural representation of places and of the relationships between these places, independent of the current position of the subject. Tolman advanced several arguments to support this view. The ability of rats to find **shortcuts** and to perform detours is one of them. The occurrence of **vicarious-trial-and-error** behaviour in the face of choice is another (see Sec. 2.6, p. 29). A third argument is the existence of **latent learning**⁶; several forms of latent learning are indeed expressed by rats, among which the mere fact that exploration improves further performance in a task (Kimble and Greene, 1968), but also that they can incidentally learn what type of reward is available even when not currently motivated for this reward (Spense and Lippitt, 1946). The **hypothesis-based** (or strategy-based) behaviour is another argument. It corresponds to a form of learning that shows a sudden shift from a near-random to near-perfect performance, contrary to trial-and-error learning. Such a change in behaviour would underlie a non-incremental neural process, i.e., a change of hypothesis about the world (see Johnson and Crowe, 2008, for a review of these arguments).

Later on, the concept of cognitive map was updated following discoveries on its putative neural bases (namely, the hippocampus; O’Keefe and Nadel, 1978). This updated theory, supported by neural data, led to a large amount of research to be directed towards the role of the hippocampus in spatial cognition. This proposal also received a few criticisms (e.g., Bennett, 1996; Gibson, 2001). Reformulations or precisions were proposed (e.g., Poucet, 1993; Eichenbaum et al., 1999; Jacobs and Schenk, 2003). Overall, this concept seems currently quite well accepted.

The use of a cognitive map in spatial cognition can be exposed by two types of tasks: those demonstrating that rats can rely on a configuration of distal cues to locate a goal and those studying their ability to create shortcut and detours.

3.2.6.2 *Localisation with respect to distal cues*

Experiments with nonhuman animals using a number of spatial tasks as, for instance, the radial arm maze, suggest that animals are able to extract and memorise the relationships between features of the environment. Then, they can use this knowledge to locate themselves as well as a hidden goal within the environment. Animals seem then to learn mental representations of space, or at least representations of the spatial relationships between relevant elements of space. For example, in the eight-arm radial maze task, each arm has a specific position with respect to distal cues (Olton and Samuelson, 1976). In the original task, all eight arms were baited, thus rats just needed to remember

⁶ Latent learning has been defined earlier in this work as learning occurring in the absence of explicit reward and without generally observable changes in behavioural performance (Johnson and Crowe, 2009).

previously visited arms in order not to visit them again. Controls were made to prevent the use of proximal cues such as olfactory traces. Interchanging arms or rotating the maze did not have any significant effect on rats' behaviour, as they kept visiting arms according to their absolute spatial location (i.e. in the room's framework). Re-baiting some of the already visited arms did not influence the order of visits either, demonstrating that rats did not rely on the smell of food. Spreading an aftershave lotion on the whole apparatus did disturb rats' behaviour on the first day of test but did not impair subsequent performance. Because rats are very good at solving this task, it could mean that they are able to store individual memories of each spatial location with respect to distal cues. However, one could argue that rats just need to store a 'snapshot' view corresponding to each arm as seen from the central stem. This would rather underlie a response strategy than a place strategy.

The most conclusive paradigm that demonstrated that rats were able to use a place strategy is the **Morris water maze** task (Morris, 1981). This task takes place in a large circular pool filled with opaque water and surrounded with distal cues (Fig. 27). A hidden platform, which is placed at a precise location just under the surface of the water, can be located with respect to distal cues. The first time a rat is released in the water, it explores the environment, using a tortuous path, and eventually finds the platform and climbs on it. After a few seconds left on the platform, it is removed from the water maze. Across trials, the goal-directed paths followed by rats become more and more direct. Again, one could argue that rats use a response strategy. But when released from different departure points, rats will show direct paths towards the platform, as shown in Fig. 27c (Morris, 1981).




Fig. 27: The Morris water maze.

Left: The water maze apparatus. In a pool filled with opaque water, a hidden platform is positioned at a constant place. Distal cues are available.

Adapted from Passot et al., 2012.

Right: Schematic trajectories. A rat released from a constant point (a) will learn a direct path to the platform in a few trials (b). When the release point changes (c), the trajectory is still direct.

Adapted from Alvernhe, 2010.

Further studies also demonstrated that rats were able to go directly to the platform even when released from a part of the water maze that they could not explore during learning (Matthews and Best, 1997), although the precise conditions under which such performance can be expressed are controversial (see Sutherland and Hamilton, 2004).

The fact that rats are able to locate the platform without local cues, which are defined by Morris as those cues that are spatially concurrent with the goal, excludes a cue strategy. The fact that rats go directly to the platform location when the departure point is changed, instead of simply reproducing the same trajectory than the one previously learned, excludes a response strategy. Thus, these results support the use of an **allothetic representation of space**, meaning that rats are able to express other forms of learning than just sensory-motor associations, and that they can locate a goal using its position with respect to distal landmarks. The water maze task is now widely used to assess spatial memory abilities in rodents and to investigate the role of different brain structures in the place or the cue strategy (as the platform can also be cued). One of the interests of this paradigm is that animals have no possibility to rely on olfactory traces and that they are highly motivated to solve the task, without the need for food deprivation.

The place strategy can also be used to solve other tasks such as the radial maze and the cross maze (Fig. 25). Similarly to the water maze, if the rat is placed on the north arm instead of the south one at the beginning of a trial, it will still go to the proper goal location by relying on a place strategy. However, if the rat uses a response strategy, it will turn right at the intersection and end up in the west, wrong arm. In other terms, if the rat uses a response strategy, it might be because it learned 'what to do', whereas if it uses a place strategy, it might be because it knows 'where to go' (Packard and McGaugh, 1996).

3.2.6.3 Role of the hippocampus in the place strategy

Lesion studies relying on rats solving either the 8-arm radial maze (Rasmussen et al., 1989, McDonald and White, 1993), the cross maze (Packard and McGaugh, 1996) or the water maze (Morris et al., 1982) demonstrated that the hippocampus is specifically involved in the place strategy but not in the response or cue strategies (see Poucet and Benhamou, 1997 for a review). Lesions of the dorsal hippocampus impaired the performance of rats in the water maze as well as in a dry version of this maze, the Oasis maze (Clark et al., 2005; Moser et al., 1993; Moser and Moser, 1998). It was recently shown that vicarious-trial-and-error, which is considered to rely on a functioning hippocampus (see Sec. 2.6, p. 29), was more expressed when a task required a place strategy compared to one where a response strategy was required (Schmidt et al., 2013).

Experiments specifically addressed the role of different subregions of the hippocampus in the place strategy. Moser and collaborators (1995) characterised the extent of remaining hippocampus that was needed to support place learning in a water maze: only 26 % of the total volume of rat hippocampus can support the use of a place strategy, provided that it is in the dorsal hippocampus. Challenging these results, Ferbinteanu and colleagues (2003) found that the ventral hippocampus was sufficient to support spatial memory. A more recent study specified that the involvement of either the ventral or the dorsal hippocampus in memory retrieval actually depends upon which anatomical part was available during learning: if the entire hippocampus is available at the time of encoding, then a lesion of the ventral hippocampus will impair memory retrieval (Loureiro et al., 2011). Coherently with the Moser study (1995), the ventral hippocampus alone is not sufficient to learn the task whereas the dorsal hippocampus can support spatial memory by itself. There are

indeed many arguments supporting the fact that the dorsal hippocampus is rather involved in spatial processing while the ventral hippocampus is involved in emotional processing such as stress or anxiety (Moser and Moser, 1998; Bannerman et al., 2004). However, we will see in the next chapter (Sec. 4.1.3, p. 58) that this dual view of the hippocampus must be refined to account for more than two subdivisions of the hippocampus (Bast et al., 2009; Strange et al., 2014).

Interestingly, with prolonged training or specific training procedures, rats with lesions to the hippocampus or the subiculum (an output structure of the hippocampus) are able to find the platform even in the hidden version of the water maze task (Morris et al., 1990; Ramos, 2010). However, if both the hippocampus and the subiculum are lesioned, even prolonged training does not allow the animals to learn the task (Morris et al., 1990). In the study by Ramos (2010) the goal was indirectly cued during training by using a light bulb placed above and slightly behind to goal position. When tested in the absence of the cue, hippocampal rats could still solve the task. However, they failed to solve a transfer test when the maze was rotated of 180° (see Fig. 29 below), and they were also impaired in a long-term memory test performed 15 days after learning. In agreement with the general view of hippocampal functions, the author concluded that hippocampal animals lacked flexibility and that the hippocampus was necessary for memory consolidation.

The fact that place learning can be possible without a functioning hippocampus indicates that other structures could support this learning. Indeed, at least one other brain region, connected to the hippocampus, was shown to be involved in place learning: the entorhinal cortex (Steffenach et al., 2005; Traissard et al., 2007).

3.2.6.4 *What are the cues used during place navigation?*

In tasks implying a place strategy, when the apparatus is rotated with respect to its centre, rats generally still search for the goal with respect to distal cues (Olton and Samuelson, 1976). Rats are able to find the platform even when a subset of the distal cues is missing. Two distal cues appear sufficient for correct performance (Fenton et al., 1994), which suggests the ability to perform pattern completion (e.g., Fellini et al., 2009; Mei et al., 2011). Although rats do rely on distal cues to navigate, an important question is whether they locate the goal by evaluating its absolute position in space (with respect to the room reference frame) or if they use directional information combined with local boundaries (e.g., distance from the pool wall in the water maze, distance from the end of an arm in the radial or cross mazes). Hamilton and collaborators (2007) addressed these issues in the water maze. After learning, they translated the pool (Fig. 28). In the new position, the rats could either search for the platform in its absolute position (the same position with respect to distal cues) or in a relative position (relative to the pool). The results demonstrated that rats chose the relative position, probably relying on directional rather than absolute place information (Hamilton et al., 2007). In a follow-up study, the authors showed that rats actually seemed to shift from absolute to directional responses with training (Hamilton et al., 2009).




Fig. 28: Search for the platform in a relative position.

Left: Initial localisation of the pool (black circle) within the experimental room. The hidden platform is represented by the black square.

Right: New position of the pool (grey circle): the absolute position of the platform is the same (black square) but rats search for the platform in a relative position with respect to the pool (grey square). A representative trajectory is plotted where black dots indicate the beginning of the trajectory.

Adapted from Hamilton et al., 2007.

The study by Ramos (2010) also allowed for testing the absolute versus relative goal localisation. Rats were trained to find the goal in a cross maze using a place strategy. To force the use of this strategy, no proximal cues were available and the departure position was randomly chosen while the goal position remained the same with respect to distal cues. In a transfer test, the cross maze was pivoted by 180°, with the goal remaining in the same absolute position (Fig. 29). In this paradigm, normal rats visited the arm that corresponded to the absolute position of the goal.




Fig. 29: Flexibility test with the cross maze: layout of the experimental room.

The cross maze used for testing could be in the training position or in the transfer position. In both, the goal arm position was the same with respect to distal cues. Hippocampal rats in this study could not solve the transfer task after being trained in the training position, contrary to controls.

Adapted from Ramos, 2010.

Thus, the issue of absolute versus directional relative coding in place strategy is not settled yet (see also Horne et al., 2007). This diversity of results probably underlies inter-individual differences as well as possible shifts in strategy with learning.

The ability of rats to evaluate a new, direct route to the goal when introduced in the water maze from a new position could be viewed as a kind of shortcut ability. However, such behaviour could also rely on pure trajectory optimisation based on the knowledge of one's position and the goal position. This does not necessarily mean that rats have abstract and quantitative knowledge of spatial information such as distance or topology. A few studies tested these abilities in rats.

3.2.6.5 Shortcut and detour ability

Demonstrating shortcut ability, according to Chapuis and Scardigli (1993), requires evidence from three criteria:

- i. the goal should not be perceptible by the subject, at least from the starting point;
- ii. the shortcut should consist of a new connection between two places;
- iii. the shortcut path should be in a straight line from the starting point to the goal (optimising distance and direction).

Tolman was one of the first researchers to study shortcut and detour planning in the rat. One of the paradigms that he used to assess shortcut behaviour is the sunburst maze (Fig. 30). In this experiment, the rats were first trained to reach a reward through an indirect sequence of alleys (Fig. 30, left). After learning, the original path was closed and rats could choose between different radial paths, one of which would lead directly to the reward site. A proportion of rats, higher than chance, chose the proper shortcut path (reported in Tolman et al., 1946).




Fig. 30: The sunburst maze.

Left: the configuration used for training. Rats had to go from the starting point S to the reward location R.

Right: the test configuration. The middle path was closed and rats could choose between twelve different arms. The arm directly aiming at the goal location was chosen by the highest proportion of rats.

From Johnson and Crowe, 2009, adapted from Tolman et al., 1946.

However, issues were raised about this experiment. For example, there was a light bulb above the reward site, which rats could possibly use as an indirect cue (O'Keefe and Nadel, 1978). Moreover, although higher than chance, the number of rats choosing the correct path was not the majority (19 / 53), probably due to the difficulty of the task. Also, several replication attempts were made with more (Muir and Taube, 2004) or less (Gentry et al., 1948) success. More recently, the shortcut ability of rats was successfully assessed in other tasks (Alvernhe et al., 2008; Winocur et al., 2010b).

Tolman and Honzik (1930) designed another maze (Fig. 31) to assess the ability of rats to make detours. Contrary to the sunburst task, this maze could not be solved by the use of directional information only. The rationale of the detour maze was as follows: if rats are given the choice between two paths of different lengths that do not directly aim at the goal, one can assess whether or not they possess a representation of the topology of the maze, including the distances between

different elements of the environment. In the maze designed by Tolman and Honzik, three paths could lead to the goal. One was straight, and the two others were detours of different length and ending location. After rats had been trained on this maze (and had explored all paths), a barrier was set in two different positions along the straight path of the maze (P1 and P2 in Fig. 31). Rats tended to preferentially choose the shortest available path, while also accounting for the topology of the environment (Tolman and Honzik, 1930).




Fig. 31: The detour maze.

Version of a detour maze from our laboratory, adapted from the original Tolman and Honzik (1930) maze. Pellet dispensers are located in place A or B.
From Alvernhe, 2010.

Some of the experiments relying on the Tolman and Honzik maze, including the original one (Tolman and Honzik, 1930), were the subject of controversy, as summarised by Ciancia (1991). However, recent experimental studies reproduced the rats' detour behaviour (Alvernhe and collaborators, 2011). In addition, computational models proposed putative neural architectures, accounting for the interplay between the hippocampus and the prefrontal cortex, which could subserve the detour planning capability (Martinet et al., 2011).

To our knowledge, few lesion experiments addressed the involvement of the hippocampus in the detour or shortcut ability. In a task where rats could take several alternative paths to reach a goal, Buhot and collaborators (1991) showed that dorsal hippocampus lesions impaired the ability of rats to find the most direct pathway. Interestingly, lesioned rats tended to use more the paths along the periphery of the apparatus. More recently, Winocur and collaborators (2010) showed that, upon introduction of a barrier in a well-known environment, control rats could find the shortest path to their goal, whereas rats with hippocampal lesions (to the entire hippocampus) took less direct trajectories. Interestingly, albeit impaired in this shortcut ability, lesioned rats demonstrated learning of the goal location relative to distal cues. The authors postulated that the hippocampus is needed to create a coherent representation that links distinct places together. The results from this last experiment actually reflect those of Ramos (2010), in which hippocampal rats could learn to locate the goal with respect to distal cues, but were impaired in a flexibility test.

To conclude, an increasing body of evidence points towards the ability of rats to use a flexible shortcut or detour strategy, although such experiments usually demonstrate higher inter-individual variability than those that require less flexible strategies.

3.2.7 Interactions between navigation strategies across learning

For a given spatial learning task, inter-individual differences in the use of navigation strategies can generally be found (Barnes et al., 1980; Devan and White, 1999). With prolonged training, rats have a tendency to shift from a place strategy to a response strategy (e.g., in the cross maze: Packard and McGaugh, 1996; Chang and Gold, 2003; Botreau and Gisquet-Verrier, 2010; in the water maze: Hamilton et al., 2009, but see Botreau and Gisquet-Verrier, 2010). A similar pattern is observed in the instrumental learning framework. As stated previously, the goal-directed nature of behaviour can be assessed by its sensitivity to outcome devaluation (Adams and Dickinson, 1981). Using this paradigm, instrumental-learning scientists demonstrated that rats can be sensitive to outcome devaluation in early stages of learning but not when overtrained (Dickinson, 1985; Killcross and Coutureau, 2003).

The behavioural expression of a response strategy does not mean that information underlying the place strategy has been forgotten. In the cross maze, Packard and McGaugh (1996) showed that rats with inactivations of the hippocampus performed prior to the test used a response strategy. By contrast, rats with their caudate nucleus (a part of the striatum) inactivated continued to display a place strategy even after prolonged training. Moreover, Pearce and collaborators (1998) showed that control rats were able to flexibly use either a cue-guided or a place strategy depending on the requirements of a variant of the water maze task. In this task, hippocampal rats could only express a cue-guided strategy.

Overall, these results highlight the fact that different (sometimes concurrent) navigation strategies are likely to operate in parallel and that normal subjects can flexibly use them according to the demands of the task (Arleo and Rondi-Reig, 2007; Jacobson et al., 2012). One can note the similarity between this idea and that of multiple memory systems exposed in Chapter 1.

3.3 Spatial cognition and memory systems

3.3.1 Declarative versus procedural memory in spatial cognition

In the spatial learning context, **declarative memory** consists in the encoding of spatio-temporal relationships between spatially relevant items, whether cues or events (O'Keefe and Nadel, 1978; Eichenbaum, 2001). It enables to flexibly retrieve the encoded information. **Procedural memory**, in the spatial domain, encompasses the acquisition of sensorimotor procedures used to navigate and the fine-tuning of trajectories. It also concerns the learning of associations between environmental stimuli and behavioural responses (Passot et al., 2012).

One can note the parallel between declarative spatial memory and the deliberative (or goal-directed) decision-making system, and between procedural spatial memory and stimulus-response decision-making. They might even be different facets of the same neural system. The similarities and differences between memory systems and decision-making systems are addressed in the recent works by Redish (2013) and Redish and Mizumori (2014).

3.3.2 Working versus reference memory in spatial cognition

Another widely used distinction of memory systems parallels the short-term versus long-term taxonomy of memory. It distinguishes between spatial working memory and reference memory (e.g., Wirsching et al., 1984; Templer and Hampton, 2013). In a given task, working memory concerns information continuously held in memory (on a short-term time scale) while reference memory concerns the long-term retention of regularities in the environment and in the task.

Rats are generally very good at solving **working memory** tasks. The seminal example is, again, the radial arm maze task. In the original study, all eight arms were baited and visits of already explored arms were counted as errors (Olton and Samuelson, 1976). In this task, rats chose on average more than 7.5 different arms within the first eight choices following only five days of learning. They could also solve a 17-arm radial arm maze task, with performance for the last (17th) trial still being above chance. Cole and Chappell-Stephenson used mazes with up to 48 arms and concluded that the limits of spatial working memory, in rats, is between 16 and 24 locations (Cole and Chappell-Stephenson, 2003). The radial-arm maze is still much used today in many variants (e.g., Babb and Crystal's series of studies about episodic-like memory, mentioned p. 6).

Another way to assess working memory in the spatial domain is the alternation task (Fig. 32), in which the obtention of reward is contingent to the fact that the rat visited a different arm than the one previously visited. Thus, to perform correctly, a rat has to remember either the movement that was performed or the arm that was visited in the previous trial. At least in the case of remembering the previously visited arm, this ability probably relies on working memory. Rats are usually very good at this task (Wood et al., 2000; Lipska et al., 2002). Indeed, they have a natural tendency for spontaneous alternation (Dember and Fowler, 1958; Lalonde, 2002).




Fig. 32: The continuous alternation task.

The rat is released in the central stem. Two reward sites are located at the end of the left or right arms. The rat can either go left (blue arrow) or right (red arrow) on each turn. It is rewarded when performing a turn different from the one he did on the previous trial.

Adapted from Wood et al., 2000.

Similarly to the memory taxonomy, in the spatial domain, it is unclear whether working memory is only of declarative nature. It is actually sometimes considered a form of short-term episodic memory (e.g., Wood et al., 2000). We could also say that the path integration mechanism relies on working memory since it consists in the continuous update of information and its manipulation to drive behaviour.

A classical test for **reference memory** uses the radial maze paradigm (e.g., Wirsching et al., 1984). In this case, only a subset of the arms is baited. There is a training phase, during which subjects learn that some of the arms are never baited. With learning, the number of visits directed towards the

non-baited arm decreases, showing that the information about baited and non-baited arms is being memorised. Successive trials on the radial maze are generally performed after a 24 h delay, which allows a form of long-term memory to be tested (although long-term memory in humans can involve information stored for more than a few hours).

3.3.3 Episodic and semantic memory

Buzsáki and Moser proposed a link between navigation strategies and memory systems (Buzsáki, 2005; Buzsáki and Moser, 2013). It relies on the reference frame underlying different navigation strategies. They highlighted the similarities between egocentric-based navigation and episodic memory, in the sense that both relate events with respect to the individual. Conversely, both allocentric-based navigation and semantic memory manipulate facts independently from a personal perspective. The authors supported the view of semantic knowledge that would be acquired through multiple episodes and drew a parallel with map-based information that requires some time to emerge, as can be seen through the need for exploration of new environment prior to proper task performance.

3.4 Conclusion: spatial cognition and the hippocampus

Spatial tasks can address various forms of memory and allow different decision-making systems to take control of behaviour, as is demonstrated by the variety of navigation strategy expressed by animals. Throughout this chapter, several aspects of spatial cognition were overviewed and the role of the hippocampus in each of these aspects was questioned through results from lesion studies. The hippocampus does not appear to be essential for a subset of navigation strategies, which indicates that this brain region is not tightly involved in direct sensory or motor processing. Rather, the hippocampus seems instrumental to performance in very specific conditions, when **flexibility** following environmental change is involved, and when **complex spatial processing** is required, in particular, when a hidden goal must be located with respect to environmental cues. It also seems to have a more general role in exploration and spatial novelty detection (Save et al., 1992; Johnson et al., 2012). In addition, the hippocampus was also hypothesised to be important for a whole range of processes, among which, generalisation (Kumaran and Mc Clelland, 2012), working memory (Olton et al., 1979), representational memory (Thomas and Gash, 1986), temporal contiguity or association of discontinuous events (Rawlins, 1985, Wallenstein et al., 1998), memory for sequences of events (Fortin et al., 2002), contextual encoding (Winocur and Gilbert, 1984), recognition memory (Olsen et al., 2012) and declarative memory (Eichenbaum and Cohen, 2014). Overall, the role of the hippocampus in spatial information processing (Hartley et al., 2014) is generally not challenged, but this view might be just one of the angles under which the role of the hippocampus can be examined.

In the following chapter, we will overview the anatomical and functional interactions of the hippocampus with its surrounding structures, in an attempt to better delineate the type of information which is processed within the hippocampus.

Chapter 4 – The hippocampus: at the crossroads of space and decision-making?

To better understand the mechanisms underlying the role of the hippocampus in spatial cognition, this chapter first reviews the anatomical relationships between the hippocampal formation and other cortical and subcortical structures. Then, we provide an overview of the functional properties of the ‘spatial circuit’ of the brain, where different types of neurons with spatially-selective discharges were discovered. Finally, we describe the interactions between the hippocampus and the network of structures known to be involved in decision-making processes.

4.1 Anatomy of the hippocampus: a brief overview

The **hippocampal region** encompasses brain structures from the **hippocampal formation** (Fig. 33) and the **parahippocampal region** (Fig. 34). The hippocampal formation consists of an archicortex while parahippocampal structures are made of neocortex. The parahippocampal structures regroup the entorhinal cortex, the peri- and postrhinal cortices, and the pre- and parasubiculum. We note that the dorsal portion of the presubiculum is sometimes termed postsubiculum (Witter & Amaral, 2004). In the hippocampal formation, one can find the dentate gyrus, the hippocampus proper and the subiculum (Amaral and Witter, 1989; van Strien et al., 2009). Note, however, that other authors include the entorhinal cortex and the pre- and parasubiculum in the definition of hippocampal formation (eg: Lever et al., 2009; Hartley et al., 2014).

4.1.1 The hippocampal formation

The hippocampus is made of two hippocampi, one per hemisphere, connected by fibres of the corpus callosum and the anterior commissure. It is a major structure of the rat brain as its unfolded surface is approximately 1.2 cm² whereas the surface of the rat neocortex would reach 1.5 cm² (Swanson et al, 1987). The hippocampus proper, also called Ammon’s horn, can be divided in three subfields: CA1, CA2 and CA3 (Lorente de Nó, 1934).

4.1 – Anatomy of the hippocampus: a brief overview




Fig. 33: The hippocampus within the brain.

Position of the hippocampal formation (banana-shaped area) within the rat brain. The cortical surface covering the hippocampus has been removed. The yellow part represents a slice of the hippocampal formation, enlarged on Fig. 35A.

Adapted from Amaral and Witter, 1989.




Fig. 34: The hippocampal region.

Lateral (**left**) and caudal (**right**) views of the hippocampal region, composed of the hippocampal formation and the parahippocampal region. The hippocampal formation consists of the dentate gyrus (DG), the hippocampus proper (CA1-3), and the subiculum (Sub). The parahippocampal region contains the presubiculum (medium blue, visible on the right panel) and parasubiculum (dark blue), the entorhinal cortex (lateral: dark green, and medial: light green), the perirhinal cortex (pink for Brodmann area 35 and purple for area 36) and finally the postrhinal cortex (blue-green). The section corresponding to the 'b' dotted line is shown on Fig. 35C.

Adapted from van Strien et al., 2009.

4.1.2 Disto – proximal organisation and the trisynaptic circuit

One of the most frequently described anatomical circuits in the hippocampus is the **trisynaptic circuit** (Fig. 35). The axons of cells from the superficial layer II of the entorhinal cortex, forming the perforant path, connect to cells in the dentate gyrus. Dentate cells connect to pyramidal cells in the CA3 subfield. CA3 cells send collaterals to CA1, whose cells project in turn to the subiculum. Closing the loop, both subicular and CA1 cells innervate cells in the deep layers (IV and V) of the entorhinal

cortex. In addition to this circuit, CA3 receives projections from the direct perforant path, emitted by layer II neurons. Moreover, neurons from the layer III of the entorhinal cortex project to CA1 via the temporo-ammonic pathway. The trisynaptic circuit is often schematised as being unidirectional but back projections, albeit less substantial, can be found at several steps of the pathway (Naber, 2001, van Strien et al, 2009).




Fig. 35: The trisynaptic circuit and connections with the parahippocampal region.

A: Enlarged portion of the hippocampus showing the trisynaptic circuit (corresponding to the slice in Fig. 33).

S: Subiculum; DG: Dentate Gyrus; pp: perforant path; mf: mossy fibres; sc: Schaffer collaterals.

Adapted from Amaral and Witter, 1989.

B: The trisynaptic loop and connections with neocortical regions. CA2 is not shown, however, it receives input from layer II and III of the MEC, and reciprocal input from CA3. It sends projections to CA1 and back to layer II of the MEC (Witter et al., 2014).

From Clark & Squire, 2013.

C: Horizontal section showing the hippocampal formation and the parahippocampal region, corresponding to the 'b' dotted line in Fig. 34. Sub: subiculum; PrS & PaS: pre- and para subiculum; MEA & LEA: medial and lateral entorhinal cortex; A35 & A36: Brodmann areas of the perirhinal cortex.

From van Strien et al., 2009.

The CA3 subfield has a very specific connectivity insofar as it contains numerous recurrent collaterals fibres, i.e., fibres that connect cells in the same subfield. A given CA3 cell is directly connected to at least 2% of other CA3 cells (Miles and Wong, 1986). This feature of CA3 led Marr (1971) to propose that this subfield could perform pattern completion. Recurrent connections are actually also present, but to a lesser extent, in the dentate gyrus, CA1 and the subiculum (van Strien et al., 2009). CA3 cells also receive connections from the contralateral hippocampus (through the corpus callosum). This is also the case for CA1, although involving fewer connections, and for the dentate gyrus.

The relatively small CA2 region is located between CA1 and CA3 on the proximo-distal axis. It has directed interest only recently and would be involved in olfactory memory, at least in mice (Kohara et al., 2013; Hitti and Siegelbaum, 2014; Stevenson and Caldwell, 2014).

4.1.3 Dorsoventral organisation

Different hippocampal domains can also be delineated along the dorsoventral axis, also termed septotemporal, longitudinal, or simply long axis. The dorsoventral and septotemporal axes of the hippocampus are often interchanged although they do not exactly refer to the same axis, as can be seen on Fig. 34. We will use the dorsoventral terminology as it seems to be the most common one. The dorsoventral organisation is quite complex as three different types of dissociation of the hippocampus, either sharp or continuous, can be made according to intrinsic or extrinsic connectivity and genetically detectable domains (Strange et al., 2014). Distinctions relevant to our work are mainly that the dorsal hippocampus is strongly connected to the dorsolateral entorhinal cortex while the ventral hippocampus is more strongly connected to the ventromedial entorhinal cortex (Fig. 36). This difference in connectivity parallels a decrease in spatial information resolution in both areas (exemplified in the next chapter, Sec. 4.2.4, p. 65). These two hippocampal subparts also differ in their function: the dorsal hippocampus appears to be involved mainly in spatial processing whereas the ventral hippocampus seems to be linked to stress and emotional functions (Kjelstrup et al., 2002; Moser et al., 1995).




Fig. 36: Differential entorhino-hippocampal connectivity along the dorsoventral axis.

Regions of similar colour are more strongly connected. EC: Entorhinal Cortex.

Adapted from Strange et al., 2014, originally adapted from Bast et al., 2009.

Importantly, intrahippocampal longitudinal connections allow neurons from different dorsoventral levels to communicate. However, there is no evidence for direct links between the dorsal part and the most ventral part. Thus, communication between those areas probably only relies on the intermediary hippocampus (Bast et al., 2009).

4.1.4 Connectivity with other parahippocampal structures

The entorhinal cortex projects to the hippocampus either ipsilaterally or contralaterally. It is connected to structures of the hippocampal region and other cortical and subcortical structures (for reviews, see Kerr et al., 2007, Canto et al., 2008). It is thought to process different kinds of inputs in a topographically organised fashion (Strange et al., 2014): the medial entorhinal cortex would process

visuo-spatial information while the lateral entorhinal cortex rather receives non-spatial (olfactory, items and emotional) information.

The hippocampus also shares connections with other structures of the parahippocampal region (Kloosterman et al., Naber et al., 2001):

- i. All subfields of the hippocampal formation, with the exception of the subiculum, receive projections from the **pre-** and **parasubiculum**. CA1, CA3 and the subiculum project in return towards both the pre- and parasubiculum.
- ii. Reciprocal connections have been described between CA1 and both the **perirhinal** and **postrhinal** cortices. Similarly, reciprocal connections also exist between the subiculum and perirhinal and postrhinal cortices.

For an extensive review of the connectivity of the hippocampal region, the reader is referred to Van Strien et al., 2009.

4.1.5 Connectivity with extra-hippocampal region structures

The main input to the hippocampal formation is the entorhinal cortex. However, cortical and subcortical structures also share connections with specific subregions of the hippocampal formation (see Hok, 2007, for review). A general description of this connectivity is that the number of afferent and efferent structures increases along the proximo-distal axis.

The **dentate gyrus** receives projections from subcortical structures such as the septum, the mammillary bodies of the hypothalamus, a few fibres from the ventral tegmental area, and brainstem structures such as the locus coeruleus (which releases norepinephrine) and the raphe nuclei (which release serotonin). CA1 and **CA3** receive subcortical inputs from the septal nuclei, in higher density towards CA3, and the ventral tegmental area (with stronger projections to CA1; Scatton et al., 1980; Gasbarri et al., 1994). The ventral and intermediary regions of **CA1** receive inputs from the amygdala. CA1 receives projections from the nucleus reuniens, preferentially at the intermediary level. The **subiculum** receives projections from the amygdala, the septum and thalamic structures such as the nucleus reuniens, and the ventral tegmental area (Scatton et al., 1980; Verney et al., 1985). Interestingly, the substantia nigra, which also sends dopaminergic signals, projects to the dorsal hippocampus (Scatton et al., 1980). Finally, the hippocampus also receives projections from the claustrum (Witter et al., 1988; Zhang et al., 2013).

Concerning the output structures, interestingly, the dentate gyrus does not project to extra-hippocampal structures: its only output is CA3. Both CA1 and **CA3** project to the septal area, which is the major extrahippocampal output of CA3. By opposition, **CA1** projects to multiple other cortical and subcortical structures, in a topologically organised way. The dorsal CA1 projects towards the retrosplenial cortex, the nucleus accumbens, the diagonal band of Broca, while more ventral CA1 sends fibres towards the dorsal peduncular cortex, the striatum, the prelimbic and the infralimbic medial prefrontal cortices (Jay and Witter, 1991). Interestingly, the prefrontal cortex does not project directly back to the hippocampus (Laroche et al., 2000). The **subiculum** also projects towards the pre- and infralimbic prefrontal cortices (more strongly than CA1) and, in addition, to the medial and

ventral parts of the orbitofrontal cortex and the cingulate cortex. Its other, subcortical, connections are directed towards the nucleus reuniens, the septal complex and the nucleus accumbens. Importantly, the subiculum also sends fibres towards the mammillary bodies via the fimbria/fornix. This output initiates a circuit first described by Papez (1937), schematised in Fig. 37, which encompasses extrahippocampal region structures where spatio-selective cells can be found, as will be seen in Sec. 4.2.2 (p. 62).




Fig. 37: Papez circuit.

The subiculum projects via the fornix to the mammillary bodies. Information flows through the mammillothalamic tract to the anterior nucleus of the thalamus, then to the cingulate gyrus via anterior thalamic radiations. Finally, the cingulum projects to the entorhinal cortex (Shah et al., 2012). Initially thought to be involved in emotions (Papez, 1937), this circuit might actually be related to learning and working memory (Jankowski et al., 2013; Vann, 2013).

To summarise this anatomical overview, the hippocampus is a cortical structure organised along different axes (proximo-distal, but also dorsoventral) that mainly receives input from the entorhinal cortex but also communicates with other cortical and subcortical structures. From its connectivity, it can be said to be involved in two main circuits, a spatial one and a decision-making one. The idea of a spatial processing circuit is supported by a wealth of electrophysiological data that will be briefly exposed in the following section.

4.2 The hippocampus within a spatial processing network

4.2.1 The hippocampus and place cells

Lesions studies indicate that the hippocampus has a major role in spatial cognition, specifically in tasks that require locating oneself and one's goal using the spatial relationships between elements of the environment. Coherently, cells from the hippocampus were discovered whose activity is highly correlated with the position of an animal in space. They were termed **place cells**. Their existence was first demonstrated by O'Keefe and Dostrovsky in a seminal publication (O'Keefe and Dostrovsky, 1971). The authors were recording the extracellular activity of dorsal hippocampal neurons in freely moving rats and realised that some of the cells would fire only when the animal was in a specific place of the environment:

"These 8 units responded solely or maximally when the rat was situated in a particular part of the testing platform facing in a particular direction."

O'Keefe and Dostrovsky, 1971

For example, the neuron presented in Fig. 38 selectively fired whenever the rat was in positions A, B or C, all of them corresponding to the bottom-right corner of the recording environment.




Fig. 38: Place-related firing of a hippocampal cell.

Left: experimental arena and possible positions of the rat. Each letter corresponds to a position where the rat was coaxed (or pushed) to go and stay for a few seconds.

Right: firing rate of a dorsal CA1 hippocampal neuron corresponding to each position.

Adapted from O'Keefe and Dostrovsky, 1971.

Muller and collaborators later introduced a specific way to represent place cell's activity under the form of a **firing map**. It indicates the firing rate of a cell as a function of the position of the animal in an environment, and allows to highlight the spatially-selective firing of place cells (Muller et al., 1987). Fig. 39 shows the firing maps of four dorsal hippocampus cells in a random foraging protocol, i.e., as the rat was chasing food pellets randomly scattered in a circular arena. The place that triggers maximal firing from the place cell has been termed **place field**.




Fig. 39: Place cells firing maps.

Firing rate maps of 4 different cells recorded from the dorsal hippocampus of rats (CA1 field). The environment was a 76-cm diameter circular arena with a white cue card positioned on the wall as indicated by the black line. To represent the activity of the cell, the environment is divided in square bins and the color of each bins indicates the median firing rate of the cell when the animal was in the corresponding position. Colors range from yellow (no firing) to purple (highest firing) with intermediary rates in the following order: orange, red, green and blue. White stands for unvisited bins.

Adapted from Muller and Kubie, 1987.

In sharp contrast with neurons of the primary visual cortex and neurons from the motor cortex, which are tied either to a particular stimulus or to a particular action (Hubel and Wiesel, 1959; Georgopoulos et al., 1986), the spatially tuned activity of place cells is seemingly independent from sensory or motor correlates. Rather, it could be seen as a multisensory and allothetic representation of space (see Sec. 3.1.2 p. 32 and Sec. 3.1.4 p. 35 for definitions of these notions). Place cells were postulated to be the neural support of the cognitive map, as defined by O'Keefe and Nadel (1978).

Cells with similar properties were subsequently discovered in other mammal species, such as mice, bats, monkeys and humans (McHugh et al., 1996; Rolls, 1999; Ekstrom et al., 2003; Ulanovsky and Moss, 2007). Place cells were central to our work; a detailed review of their functional properties will be found in Chapter 5 (p. 81).

Researchers later discovered other cells that expressed spatial firing, namely, head-direction cells, grid cells and boundary cells, which all have the property of representing spatial parameters in an allocentric reference frame, and belong to structures generally connected to the hippocampus. Together, these structures are thought to form the spatial circuitry of the brain (Moser et al., 2008).

4.2.2 Cortical and subcortical networks with head direction cells

Head direction cells were first found in the dorsal presubiculum by Ranck (in 1984) and subsequently described by Taube and colleagues (Taube et al., 1990). The most striking property of these cells is their activity tuned to a specific direction in space, independently from the position of the animal. The firing map of a head direction cell is shown in Fig. 40B, either when the head direction of the animal is not considered (in the middle), or when a separate map is created for specific head directions (peripheral plots). In this case, the firing is maximal for a head direction around 45°. The directional tuning of head direction cells is yet more visible when the firing rate of the cell is expressed as a function of the head direction (Fig. 40C). Thus, head direction cells can be said to encode an allocentric direction.




Fig. 40: Head-direction cells from the dorsal presubiculum.

A: Position of the postsubiculum where the cells were found in a schematic coronal section of rat brain.

Retspl: retrosplenial cortex; Postsub: postsubiculum; Subic: subiculum; DG: dentate gyrus.

B: Composite (in the middle) and direction-specific maps for a given head direction cell. The colour code ranges from yellow (no spiking) to purple (max firing) with orange, red, green and blue as intermediary rates.

C: Firing rate of three cells as a function of rat head direction. Adapted from Taube et al., 1990.

The activity of head-direction cells relies on allothetic cues, as demonstrated by the fact that when a prominent cue in the environment is rotated, the preferred directions of these neurons rotate accordingly (Knierim et al., 1995; Sargolini et al., 2006). Active locomotion also influences the response of head-direction cells, which will fire more for active than passive rotational movements (Zugaro et al., 2001). If a rat is disoriented and then provided with a visual cue, the whole population of head direction cells seems to rapidly and coherently reorient (Zugaro et al., 2003). Head direction-cells maintain their activity in the dark, indicating that self-motion cues are sufficient to maintain their selectivity (Taube et al., 1990). For instance, optic flow information was recently shown to drive the activity of head direction cells in the anterodorsal thalamus (Arleo et al., 2013).

Head direction tuned signals were also observed in many structures: the entorhinal cortex (Sargolini et al., 2006), the lateral dorsal thalamus (Mizumori and Williams, 1993), the anterior dorsal thalamus (Taube, 1995), the lateral mammillary nuclei (Stackman and Taube, 1998), the retrosplenial cortex (Chen et al., 1994), the striatum (Wiener, 1993), the cingulum (Leutgeb et al., 2000), and more recently, the nucleus reuniens (Jankowski et al., 2014). Interestingly, a majority of the structures where head-direction signals were found belong to the Papez circuit (Fig. 37, above). They are particularly abundant in the anterodorsal thalamic nucleus (Taube, 2007). Slight differences between head direction cells from different areas have been described (reviewed in Taube, 1998, 2007).

Functional interactions between place and head direction cells seem to occur in both directions. Hippocampal lesions impair the directional signal from the dorsal presubiculum and the anterodorsal

thalamic nucleus (Golob and Taube, 1999). Conversely, following a lesion of the dorsal presubiculum, the place cell activity becomes unstable. This is also the case, though to a lesser extent, in rats with a lesion of the anterodorsal thalamic nucleus (Calton et al., 2003). Other studies have shown that lesions to the mammillary bodies impair head direction signals in the subiculum, post subiculum, medial entorhinal cortex and pre- and para- subiculum, but they do not alter the firing of hippocampal place cells (Sharp and Koester, 2008). Thus, only a subpart of the head direction system seems to functionally interact with the hippocampus (see Wiener and Taube, 2005, for a review on head direction cells).

4.2.3 Cortical networks with boundary cells

The most recently discovered class of spatio-selective cells is a class of cells whose activity is associated with the borders of an environment, termed ‘boundary vector cells’ (Barry et al., 2006, Lever et al., 2009), ‘putative boundary cells’ (Savelli et al., 2008) or ‘border cells’ (Solstad et al., 2008). The firing of boundary vector cells depends on the rat’s location with respect to environmental boundaries, independently of the rats’ heading direction (Lever et al., 2009). Putative boundary cells or border cells fire when the animal is close to one or several local boundaries of the environment, also independently from the rat’s head direction (Boccaro et al., 2010, see Fig. 41). The border cells or putative boundary cells were postulated to be a subclass of boundary vector cells with short distance tuning (Lever et al, 2009). The boundary-specific firing occurs regardless of the position of the boundary in the environment (see the second inset from cell 2 in Fig. 41B) and is expressed for similar borders in different environments. Similarly to head-direction cells, boundary cells maintain their firing in darkness conditions and their field will rotate following the rotation of polarising visual cues (Solstad et al., 2008; Lever et al., 2009).




Fig. 41: Border cell from the entorhinal cortex.

A: Firing map of a border cell. The activity of the cell follows the wall when the environment is stretched.

B: Firing map of the same border cell (same orientation of the field as in A with respect to distal cues). When a wall (white part in the middle inset) is inserted in the environment, this triggers firing of the cell. Adapted from Solstad et al., 2008

Boundary vector cells can be found in the subiculum (Barry et al., 2006; Lever et al., 2009) and border cells in the medial entorhinal cortex (Savelli et al., 2008; Solstad et al., 2008; recent recordings in the juvenile rat: Bjerknes et al., 2014) and the pre- and parasubiculum (Solstad et al., 2008; Boccara et al., 2010).

In contrast to the classic view of the subiculum, which is seen as an output structure of the hippocampal formation, Lever and colleagues (2009) argue that subicular boundary cells would provide input to hippocampal CA1 cells. This would be possible through the subiculum-entorhinal-hippocampus circuit (Kloosterman et al., 2004). Such a signal could be used to provide place cells with information about the borders – possibly, the geometry – of the environment.

4.2.4 The entorhinal cortex and grid cells

Grid cells are another type of spatially-tuned neuron with very specific properties (Fyhn et al., 2004; Hafting et al., 2005). These cells express a very regular spatial firing pattern that tessellates the environment (see Fig. 42), that is to say, each grid cell shows multiple peaks of activity (or place fields) positioned regularly with the same distance from each other.




Fig. 42: Grid cells from the medial entorhinal cortex.

Example of four different simultaneously recorded cells in a foraging task in a square environment. On the left of each inset, the spikes are represented in red, superimposed on the rat's path, in black. The corresponding firing rate map is on the right. Numbers indicate the peak rate for each map.

Adapted from Fyhn et al., 2004.

Grid cells usually express a direction-independent firing, but a subset of them are 'conjunctive' grid cells, i.e., they combine a spatial and a directional selectivity (Sargolini et al., 2006). This means that their maximal firing can only be observed when the animal is in specific places and has a specific head direction. Similarly to head direction cells and border cells, the firing of grid cells is stable across successive visits to the same environment (Hafting et al., 2005). When recorded in different environments, grid cells do not change their intrinsic characteristics such as the distance between their place fields or the size of the fields (Fyhn et al., 2007). However, the grid is not completely rigid. Following reshaping of a familiar environment (expansion or contraction), the grid pattern will also expand or contract (Barry et al., 2007). Moreover, upon entry in a new environment, the grid is expanded and progressively comes back to normal upon several exposures (Barry et al., 2012). Interestingly, the scale of the 'grid' of different cells, i.e., the distance between adjacent firing fields, increases along the dorsoventral axis of the medial entorhinal cortex (see Fig. 43; Fyhn et al., 2004; Hafting et al., 2005; Brun et al., 2008). This increase in grid scale is discretised: cells are organised in anatomical clusters along the dorsoventral axis that share similar grid sizes and functional properties (Stensola et al., 2012).




Fig. 43: Anatomical gradient of grid size along the dorsoventral axis.

Sagittal section of rat brain with hippocampal formation subfields (DG, CA3, CA1), the postrhinal cortex (POR) and the medial entorhinal cortex (MEC) where grid cells can be found. The colour code is as above. The distance separating fields increases along the dorsoventral axis
From McNaughton et al., 2006.

Grid cells are located in the dorsolateral medial entorhinal cortex, which is directly connected to the hippocampus proper. More precisely, the layer II of medial entorhinal cortex mainly holds grid cells whereas the deeper layers contain grid cells, head direction cells and conjunctive cells. Grid cells can also be found in the pre- and parasubiculum (Boccarda et al., 2010), which are also connected to the hippocampus. Combining electrophysiological recordings with optogenetic techniques, Zhang and collaborators (2013) aimed at identifying which types of cells from the entorhinal cortex were projecting to the hippocampus. Intriguingly, they found that not only ‘spatial’ cells (grid-cells, head direction cells and boundary cells) projected to the hippocampus, but also irregularly spatial cells and non spatial cells. Lesions to the entorhinal cortex impair the activity of place cells (Brun et al., 2008; Van Cauter et al., 2008). More precisely, Brun and collaborators lesioned the layer III of the medial entorhinal cortex, which is the one that projects towards CA1 (see Fig. 35 above). Logically, this lesion affected the spatial activity of CA1 but not CA3 cells. However, in these tasks, the entorhinal lesions did not completely abolish the spatially-selective signal of hippocampal cells, which means that place cells probably receive complementing or possibly redundant spatial information from other structures of the brain spatial circuitry.

Because of the regularity of their firing, and the fact that grid cells seem to encode both information about distance and direction, these cells are thought to be the neural support of path integration. This hypothesis is supported by results from lesion studies addressing the possible role of the entorhinal cortex in path integration (Parron and Save, 2004; Jacob et al., 2014). However, the above-mentioned results (for example, rescaling of the grid according to experience or novelty) moderate this hypothesis. It was recently proposed that the medial entorhinal cortex (MEC) processes both idiothetic and allothetic information but conveys different signals to the hippocampal formation depending on the currently available information (Poucet et al., 2014). The proposed circuit is presented in Fig. 44. In the dark, idiothetic information is predominant. This type of information would mainly be processed by the dorsal MEC and then sent to the hippocampal formation to update the map-based representation of space. On the contrary, in light conditions, the main flow of information would be allothetic. It would be processed by the ventral MEC (with place-like cells) and then to the hippocampal formation to update the place cell activity. In return, the hippocampus would also update the grid cell representation.




Fig. 44: Two-state model of navigation.

Information flow in the entorhino-hippocampal circuit under either light (left) or dark (right) conditions.

In light conditions, the active circuit is the red one; in dark conditions, it is the green one.

dMEC: dorsal medial entorhinal cortex
vMEC: ventral medial entorhinal cortex
From Poucet et al., 2014.

To conclude, the entorhinal cortex can provide the hippocampus with at least two possible types of information: spatial information, built from either idiothetic or allothetic sensory sources depending on environmental conditions, and novelty information (with the expansion of fields upon entry in a new environment, Barry et al., 2012). For recent reviews on the links between place and grid cells, the reader is referred to Bush et al., 2014 and Poucet et al., 2014.

4.2.5 Other structures with spatially tuned activity

Spatially-selective cells were also found in the subiculum (Barnes et al., 1990; Sharp and Green, 1994), in the dorsal presubiculum (Taube, 1995), in the parasubiculum, in the medial and lateral entorhinal cortices (MEC and LEC) and in the postrhinal (Burwell and Hafeman, 2003) and perirhinal cortices (Burwell et al., 1998). A subset of these cells is shown in Fig. 45. The properties of their place fields differ, in particular, with respect to their **spatial information**, a quantitative measure of the amount of information about position provided by a spike from a given cell. While CA1 cells have high levels of spatial information, lower values are found for MEC and parasubiculum cells, and an even lower spatial precision is expressed by LEC and perirhinal cells (Hargreaves et al., 2005).




Fig. 45: Spatial maps of cells with spatial tuning from different structures of the hippocampal area.

Three example cells recorded in different areas: CA1, the medial entorhinal cortex (MEC), the lateral entorhinal cortex (LEC), the parasubiculum and the perirhinal cortex. Note that the third cell recorded in the MEC seems to present the regular pattern characteristic of grid cells. From top to bottom, spatial information content decreases. Adapted from Hargreaves et al., 2005.

Hargreaves and colleagues (2005) interpret these differences in spatial information by proposing that the LEC would convey non-spatial signals (about objects, for example) from the perirhinal cortex to the hippocampus. This information would then be combined with spatial information arriving from the MEC. In this way, the hippocampus could express object-place representations. This hypothesis is supported by lesion studies showing that *i*) the perirhinal cortex (which projects to the LEC) is necessary to detect objects-related novelty and *ii*) the postrhinal cortex (projecting to the MEC) would be involved in novelty related to object-space configurations. However, lesions of the postrhinal cortex do not seem to have much effect on place cell firing (Nerad et al., 2009).

In summary, the hippocampus interacts with other structures of the spatial processing network mostly through the entorhinal cortex (Fig. 46). The medial entorhinal cortex provides allothetic information (geometry via border cells, orientation via head direction cells) and idiothetic information (path integration estimation of position via the grid cells). Depending on the environmental conditions, the relative importance of allothetic versus idiothetic input and output would vary (see Fig. 44). The lateral entorhinal cortex provides the hippocampus with information about objects and possibly supplementary allothetic and idiothetic spatial information with odours (and odour flow).




Fig. 46: The hippocampus within a spatial processing network.

Adapted from Brandon et al., 2014.

The discovery of place cells put the emphasis on the hippocampus as a central structure for spatial processing (O’Keefe and Dostrovsky, 1971; O’Keefe and Nadel, 1978). The discovery of other cells with striking spatially-related correlates rebalanced the roles. It is still tempting to see the hippocampus as the structure integrating spatial signals of different kinds in order to produce the spatial representation that is ultimately used for navigation. However, this view is probably too schematic (Bush et al., 2014). More could be said on the properties of head direction, grid and border cells, but as they are not the center of our study, these cells will not be detailed any further. The interested reader might find more (for example) in Bush et al., 2014 and Hartley et al., 2014.

4.3 The hippocampus within a decision-making network

Aside from the structures involved in spatial processing, the hippocampus also directly communicates with a subset of structures known for their role in processing decision-making parameters. We will focus on the interactions of the hippocampus with the main structures of this ‘decision-making’ network, namely, the striatum, the amygdala, the ventral tegmental area and the prefrontal cortex. Note that our use of decision-making is in its broad sense and encompasses regions linked to value processing, action selection as well as planning and decision.

4.3.1 The hippocampus and the striatum

The **striatum** is divided into a dorsal and a ventral region. The dorsal striatum can be separated in a dorsolateral and a dorsomedial domain. The ventral striatum encompasses the nucleus accumbens and the olfactory tubercle. Finally, the nucleus accumbens is commonly dissociated in a ‘core’ and a ‘shell’ parts.

To clarify the connectivity explained above, the subiculum and CA1 project to the most ventral parts of the striatum (medial, ventral and rostral shell), as well as to the immediately adjacent parts of the core (Voorn et al., 2004). Similarly to its connectivity with the entorhinal cortex, the connectivity between the hippocampus and the nucleus accumbens is topographically organised (Voorn et al.,

2004; Strange et al., 2014). Dorsal hippocampus connections target the lateral part of the nucleus accumbens, while ventral hippocampus connections connect the medial nucleus accumbens, as shown in Fig. 47.




Fig. 47: Hippocampal to nucleus accumbens connectivity.

Dorsal-to-ventral connectivity from the hippocampus to lateral-to-medial regions of the nucleus accumbens.

ac: anterior commissure

Adapted from Voorn et al., 2004.

The view of the striatal function stated previously is that it would be involved in procedural memory, response strategy and habitual decision-making. The picture is actually more complex as the different striatal regions appear to have different contributions. The dorsolateral striatum is generally thought to be involved in habitual behaviour while the dorsomedial striatum would support action-outcome behaviour (Balleine and O'Doherty, 2010). The ventral striatum, to which the hippocampus sends projections, would also play a role in goal-directed behaviour (Johnson et al., 2007), in particular, linking spatial with value information (Humphries and Prescott, 2010) or evaluating the possible outcome of actions (van der Meer et al., 2010). Lesions of the ventral striatum impair learning (Sutherland and Rodriguez, 1989), short-term (Ferretti et al., 2007) or long-term (Ferretti et al., 2010) memory in spatial tasks (for a review, see Rinaldi et al., 2012). Moreover, the core and shell subregions of the nucleus accumbens have different functions (Everitt and Robbins, 2005). In particular, the core would be involved in goal-directed behaviour (but not in outcome evaluation *per se*; Corbit et al., 2001). The striatum of rats, as seen in chapter II, holds neurons that represent different types of value, most often in the ventral striatum (Lavoie and Mizumori, 1994; Kim et al., 2009; Roesch et al., 2009; Howe et al., 2013; Bissonette et al., 2013). The ventral striatum also holds neurons tuned to motivation-related information (Lansink et al., 2008; Bissonette et al., 2013).

In summary, the ventral striatum / nucleus accumbens is thought to be a major actor in selectively linking values to specific goals (Mannella et al., 2013). In this framework, the hippocampus could either provide the striatum with purely spatial information (Lansink et al., 2009; van der Meer and Redish, 2011), or with saliency of specific places that would be mainly linked to the novelty of stimuli (Mannella et al., 2013).

4.3.2 The hippocampus and the amygdala

The **amygdala** is a structure of the limbic system located just next to the most temporal part of the hippocampus. It is composed of several structures and nuclei presented in Fig. 48. Relevant for our work is the existence of reciprocal connections between CA1 or the subiculum and the amygdala. Fewer projections also arrive to CA3 (Pitkänen et al., 2000). These connections are largely confined to the ventral two-thirds of the hippocampus (Pitkänen et al., 2000, Strange et al., 2014). Once more, this connectivity is topographically organised (see Fig. 48).




Fig. 48: Topography of hippocampal to amygdala connections.

Intermediary to ventral connectivity from the hippocampus to the lateral-to-medial amygdala.

IC: internal capsule
 CeA: central nucleus of the amygdala
 LA: lateral amygdala
 BLA: basolateral amygdala
 BLP: posterior basolateral nucleus of the amygdala
 BLV: ventral basolateral nucleus of the amygdala
 BMA: basomedial nucleus of the amygdala
 ME: medial nucleus of the amygdala
 PMCo: Posteromedial cortical nucleus of the amygdala
 PLCo: Posterolateral cortical nucleus of the amygdala
 Adapted from Strange et al., 2014, originally adapted from Kishi et al., 2006.

According to the multiple memory systems seen in Chapter 1 (White and McDonald, 2002), the amygdala is thought to be the central structure of emotional memory. More precisely, it would support Pavlovian conditioning, i.e., associations between a neutral stimulus and one with emotionally salient properties, either rewarding or aversive. It is not implied in simply assessing the hedonic properties of a reward, or in the storage of preferences, but its role arises when value information must be linked with another parameter and stored for future recall. As an example, in a task where the presence or absence of food was indicated by the type of food previously encountered (either oat bake or Froot Loops), amygdala-lesioned rats were impaired in post-acquisition performance, but not in simple preference tests (Kesner and Williams, 1995). The authors provide other evidence that the amygdala would be involved in remembering the magnitude of a reward in tasks where it is needed for successful performance. However, the amygdala does not seem involved when the presence of food only indicates the appropriateness of an action. Importantly, rats with hippocampal lesions were not impaired in the above-mentioned task, indicating that the hippocampus might not be critical for the memory of magnitude of reinforcement. Similarly, other studies point to a role of the amygdala in processing the motivationally significant attributes of a reward (Blundell et al., 2001; Gilbert et al., 2003) or attention-related emotional processing (Meck and MacDonald, 2007). Neurons from the basolateral amygdala of rats were shown to fire at the presentation of a reward-predicting cue (Tye et al., 2008) and following extinction of an expected reinforcement (Tye et al., 2010). Overall, the amygdala seems to be involved in a variety of tasks involving emotions and emotion-related learning (such as

fear conditioning) but also reward-related learning as it has a major role in the evaluation of reward value – when relevant to the task (Balleine and Killcross, 2006).

The amygdala can modulate other memory systems such as the hippocampus. In particular, it would modulate memory consolidation (McGaugh, 2004). Moreover, the connections between the amygdala and the hippocampus could explain the role given to the ventral hippocampus in emotional processes and in particular fear conditioning (Strange et al., 2014). This connectivity seems functional as shown for example, by the results of Terada and colleagues (2013). They simultaneously recorded electrical activity from the basolateral amygdala and the CA1 field of the hippocampus of rats⁷. The rats were trained in a discrimination task with varying probabilities of reward delivery, which probably requires goal-directed behaviour since performance relies on proper evaluation of reward. In rats expecting a highly probable reward, there was a marked increase in the coherence of oscillations between the two structures. This would reflect an influence of the amygdala over the hippocampus centred on the communication of information about reward expectation. However, this does not necessarily mean that the hippocampus could be the structure where place and value are integrated, as more evidence points towards the striatum for this role (Pratt and Mizumori, 1998; Mannella et al., 2013). Indeed, projections both from the hippocampus and the basolateral amygdala converge onto similar regions of the striatum (Pratt and Mizumori, 1998).

4.3.3 The hippocampus and the VTA

The ventral tegmental area (VTA) belongs to the ventral midbrain and contains dopaminergic neurons. Dopamine was originally seen as a mere reward signal. More recent views state that it would be a modulator signal involved in learning and adaptive behaviour (Schultz, 2007; Roesch et al., 2007; Shohamy and Adcock, 2010; Lisman et al., 2011; see also the special issue about its role in behavioural flexibility, Beeler et al., 2014). Interestingly, VTA neurons were shown to discharge in relation with the rat's velocity in addition to reward – but this result could also be linked to motivation (Puryear et al., 2010). Dopaminergic neurons from the VTA innervate a wide range of brain structures. We saw that they directly project to the hippocampal formation, mainly to CA1 and the ventral subiculum, but VTA projections also exist to the dentate gyrus and CA3 (Gasbarri et al., 1994). Interestingly, the VTA projects both to the dorsal and the ventral hippocampus (Scatton et al., 1980). There does not appear to be any direct projection from the hippocampal formation to the VTA (Luo et al., 2011). However, Luo and collaborators found an indirect, functional connection, arising from CA3, which would connect the lateral septum and whose neurons connect in turn the VTA. A stimulation of dorsal CA3 triggered a response of a majority of VTA neurons. The authors proposed that this pathway could be involved in the processing of context-reward associations. More precisely, CA3 would provide information about the general context to the VTA. In addition, an indirect connection possibly also emerges from the ventral subiculum since a stimulation of this area has effects in the VTA. In that case, the circuitry remains to be determined but could involve the medial prefrontal cortex (Legault et al., 2000).

⁷ Probably the dorsal hippocampus, but that was not explicitly specified.

Lisman and Grace (2005) exposed a theory stating how the hippocampus and the VTA could communicate. They proposed that the hippocampus would serve as novelty detector, sending a signal to the VTA under novelty conditions (Fig. 49). Upon reception of this signal, dopaminergic neurons would release dopamine in the CA1 field, with the effect of triggering or facilitating learning.




Fig. 49: The hippocampus - VTA loop theory.

According to this theory, the hippocampal formation, the dentate gyrus and CA3 recall memories which are sent to CA1. Comparing these memories with direct cortical input from layer 3 of the entorhinal cortex yields novelty signals. These are sent to basal forebrain structures through the subiculum and reach VTA dopamine cells. Direct connections from the VTA to the hippocampus close the loop.

PPTg: Pedunculopontine tegmentum, an excitatory input to the VTA, receives information from the prefrontal cortex (PFC) and limbic structures (of which, the hypothalamus and the amygdala) and could provide affect-related information.

From Lisman and Grace, 2005.

The theory of Lisman and Grace insists on the novelty-related role of the hippocampus, in agreement with the view according to which the hippocampus would provide value of an item in the form of its novelty (Mannella et al., 2013). This type of value is to be later combined with other information, for example, goal-related motivation and salience information, to reflect the importance of the new information. Lisman and Grace suggest that the striatum would be the ideal structure to combine novelty signals and goal-dependent motivational signals, stating that the goal-directed information would come from the prefrontal cortex. In this framework, the prefrontal cortex would send goal-related information directly to the nucleus accumbens.

In line with this theory, Bethus and collaborators (2010) showed that the dopamine afferents from the VTA to the hippocampus were involved in memory. More precisely, an inactivation of the hippocampal dopaminergic receptors modulated the persistence of memory but did not impair either initial encoding or memory retrieval. This is also in accord with the idea that VTA afferents would enhance the long-term memory of new information. Coherently, an inactivation of the VTA impairs LTP in the hippocampus of anaesthetised rats (Ghanbarian and Motamedi, 2013). In the same study, a VTA inactivation also impaired memory consolidation after learning but only when performed immediately after learning, not 20 minutes after. In another study, a disruption of the VTA prior to testing in a working memory spatial task (in the radial maze) impaired performance and

disrupted the activity of place cells in CA1 but not in CA3 (Martig and Mizumori, 2011). In line with the whole theory, Fujisawa and Buzsáki (2011) demonstrated that the VTA, the hippocampus and the prefrontal cortex could interact during a working memory task via local field potentials synchronisation.

4.3.4 The hippocampus and the prefrontal cortex

We saw previously that the ventral CA1 and dorsal subiculum both send direct, unidirectional projections to the ventromedial prefrontal cortex. Moreover, the subiculum projects to the **nucleus reuniens**, which belongs to the midline thalamic nuclei and which, among other targets, projects back to the hippocampus. The nucleus reuniens also receives strong input from the medial prefrontal cortex (Vertes, 2004). Thus, a double circuit connects the hippocampus and the prefrontal cortex: directly through the CA1 and subicular connections, and indirectly through the nucleus reuniens. The overall connectivity is much stronger between the nucleus reuniens and the ventral hippocampus than with the dorsal hippocampus. Some of the nucleus reuniens neurons even show collateral projections to both the medial prefrontal cortex and the hippocampus (Hoover and Vertes, 2012).

In addition, the prefrontal cortex projects to the nucleus accumbens, the amygdala and the VTA, among others (Vertes, 2006). Interestingly, some of the hippocampal neurons that project to the amygdala also project to the medial prefrontal cortex (Ishikawa and Nakamura, 2006). The main circuits between the hippocampus and the medial prefrontal cortex are summarised in Fig. 50.




Fig. 50: Hippocampal-medial prefrontal connectivity.

Only direct connections are indicated. Note that many structures are not included here, for example, the entorhinal cortex which also receives input from the VTA, the nucleus reuniens and the prefrontal cortex and connects to the amygdala and the nucleus accumbens.

RE: Nucleus Reuniens
 Hip: Hippocampus
 mPFC: medial Prefrontal Cortex
 BLA: Basolateral Amygdala
 NACC: Nucleus Accumbens
 VTA: Ventral Tegmental Area

The nucleus reuniens is believed to play a role in hippocampal-dependent learning and memory (Loureiro et al., 2012; Xu and Südhof, 2013), possibly only when this learning is restricted to tasks requiring both the hippocampus and the prefrontal cortex (Hembrook et al., 2012; for review, see Cassel et al., 2013).

4.3.4.1 *The hippocampus and the medial prefrontal cortex*

The prefrontal cortex is thought to be involved in working memory, decision-making, strategy switching, planning; in a nutshell, executive functions (Fuster, 1997; Churchwell et al., 2010). More precisely, the medial prefrontal cortex (mPFC) is involved in selective attention, visceromotor control, decision-making and goal-directed behaviour (Vertes, 2006; Coutureau et al., 2012 – see Martinet et al., 2011 for a modelling approach). It has been shown to interact with the hippocampus in various ways (Vertes, 2006; Battaglia et al., 2011). In particular, these interactions would support memory consolidation and might be under modulation of dopamine (Benchenane et al., 2010). The interactions between these two structures could also play a role in working memory (Wang and Cai, 2006), above all if long retention delays are involved (Churchwell and Kesner, 2011). Finally, a role of these interactions in spatial navigation was proposed (Churchwell et al., 2010), supported by results showing alterations of hippocampal place cell firing following mPFC lesions (Kyd and Bilkey, 2003, 2005). As stated previously, the prefrontal cortex, the VTA and the hippocampus might interact at least via synchronisation of rhythmic activity (Fujisawa and Buzsáki, 2011).

Several studies from our laboratory helped to define the role of the medial prefrontal cortex in spatial cognition. First, lesions of the mPFC of rats that needed to update the value of a spatial goal impaired the long-term retention of the new (reduced) goal value (De Saint Blanquat et al., 2013). Interestingly, lesioned animals more quickly updated the new value during the task, which could be seen as a lack of competition from the previously stored value. In an electrophysiological study, De Saint Blanquat and colleagues (2010) recorded the activity of neurons from the medial prefrontal cortex of rats performing a radial arm maze task in its reference memory configuration (i.e., only a subset of the arms were baited). During choice periods in the central stem of the maze, specific subsets of cells would fire according to the status of the arm that would be subsequently visited (either already visited or not-yet-visited). Importantly, these cells did not specifically fire according to the baited or non-baited status of the arms. Thus, the value of goals does not seem to be encoded in the activity of mPFC neurons. Then, perhaps the effects observed by De Saint Blanquat and collaborators (2013) concerning the updating and the consolidation of the new value could result from a modulating effect of the mPFC on other structures involved in value coding instead of storage of the value by mPFC neurons.

Of particular interest to our work is a study in which Hok and colleagues (2005) evidenced the existence of goal-related activity in the medial prefrontal cortex using the **continuous navigation task** (cf. Fig. 51 A). In this task, a rat is trained to locate an invisible goal zone using its position with respect to a cue card on the wall. Whenever the rat stays at least for 2 seconds in this specific place, a food pellet is released from a dispenser located above. Pellets first fell on a landing zone then bounce to eventually stop in a random place. The rat can then eat the pellet and come back to the goal zone to trigger the release of a new pellet. One advantage of this task is that it enables to dissociate between two types of episodes: goal-directed navigation, when the rat navigates towards the goal zone, and foraging when it searches for the pellet. Another advantage is that the location of goal and reward are dissociated, enabling to selectively study the neural correlates of one or the other. The activity of medial prefrontal cortex neurons with respect to the position of rats is

presented in Fig. 51 B. Around 25 % of the recorded neurons expressed a spatially selective firing, although a bit less selective than hippocampal place cells. This spatially-selective activity had several characteristics of interest:

- i. It was exhibited under different forms but usually with a specific temporal profile, as shown by the results from a subsequent study (Burton et al., 2009), presented in Fig. 51 D. Overall, the mean population activity shows an increase of firing after entry in the goal zone and is maximal just before the pellet dispenser activation.
- ii. The ‘place fields’ of these cells were present regardless of the behavioural phase: there was no difference in firing between goal-directed and foraging episodes.
- iii. The **centroids** (i.e., centre of mass of the place fields) of these cells were not uniformly distributed. Most of them were located around the trigger zone and the landing zone (cf. Fig. 51 C), which are both task-relevant places. It is worth mentioning that in a purely foraging task, mPFC neurons do not express any spatially-selective activity, contrary to hippocampal place cells (Poucet, 1997).




Fig. 51: Goal-related activity in the medial prefrontal cortex.

A: Experimental paradigm. The rat arriving from the position indicated by the black dot waits two seconds in the trigger zone to activate the pellet dispenser. A pellet is subsequently released in the landing zone but stops in a random place. In this example, the rat first goes to the landing zone (goal-directed path, solid line), then to the final position of the pellet (foraging path, dotted line). The black curved line indicates the position of a cue card on the wall of the arena.

B: Firing rate maps of six representative cells with ‘fields’ either in the trigger zone or the landing zone. Cells could also demonstrate fields in other zones. Colour code ranges from yellow (no firing) to purple (maximum firing). Orange, red, green and blue indicate intermediary values in increasing order.

C: Spatial distribution of field centroids.

D: Temporal profile of goal-related firing. Mean peri-event time histogram (PETH) of activity before, during (gray overlay) and after the goal waiting period, cumulated over all trials for the cell population. The bin width is 250 ms. Error bars indicate standard error.

E: Temporal profile of goal-related firing in hippocampal rats. The parameters are the same as above.

A, B and C from Hok et al., 2005 ; D and E from Burton et al., 2009.

Of much interest are the results from a subsequent experiment in the same paradigm, in which trained rats were lesioned in the intermediary – to – ventral hippocampus (Burton et al., 2009). The lesions had three major effects. First, lesioned rats did not perform as well as controls. They showed thigmotaxis behaviour (i.e., they had a higher tendency to follow the walls of the arena) and were less willing to wait at the goal location, which could be related to impulsivity. Second, the number of medial prefrontal cells expressing goal-related activity was significantly reduced in lesioned rats compared to controls. This was the case even for a cued version of the task. Finally, for those trials where the rats were sufficiently motivated to trigger the reward, the temporal profile of activity expressed by control rats was abolished in hippocampal rats (see Fig. 51 D and E). The authors concluded that the hippocampus and the mPFC are probably involved in either impulse control or reward expectation and that the goal-related activity expressed by mPFC cells is dependent on the hippocampus.

4.3.4.2 *The hippocampus and the orbitofrontal cortex*

Another structure of the prefrontal cortex, the orbitofrontal cortex (OFC), houses neurons that seem to represent many parameters involved in decision-making such as expected outcomes (van Duuren et al., 2009), spatial goal correlates (Feierstein et al., 2006), reward magnitude correlates (van Duuren et al., 2007, 2008), neural correlates of the ‘regret’ associated to a decision (Steiner and Redish, 2014), cost (in time or effort) and confidence in a given decision (reviewed in Wallis, 2012). The localisation (and possible homology) of the OFC in the rat compared to primates still seems to be a matter of debate (Preuss, 1995; Uylings et al., 2003; Wise, 2008). The rat OFC is believed to receive fibres from the ventral hippocampus (Ishikawa and Nakamura, 2006). Moreover, it receives projections from the subiculum and shares numerous reciprocal connections with the parahippocampal area (reviewed in Ramus et al., 2007); it also receives inputs from the nucleus reuniens and from the medial prefrontal cortex, the two of which receive input from the hippocampus (Reep et al., 1996). To our knowledge, only a few studies addressed the possible interactions between the hippocampus and the orbitofrontal cortex in rats (Vafaei and Rashidy-Pour, 2004; Ramus et al., 2007). In particular, Ramus and collaborators (2007) trained rats in an olfactory memory task and found that a subset of neurons from the orbitofrontal cortex fired in anticipation of specific odours. Very interestingly, the authors showed that a lesion of the hippocampus abolished this anticipatory firing. They propose that the OFC and hippocampus interact during the storage of long-term memories and that the OFC would be the ultimate repository of memories pre-processed by the hippocampus, following the systems consolidation theory.

4.4 Conclusion

The hippocampus appears to be involved in two main circuits. Through its connections with the parahippocampal region and Papez circuit, the hippocampus falls within a well-known spatial processing circuit (Moser et al., 2008). Through its interactions with the decision-related regions (medial and orbitofrontal prefrontal cortex), the reward-related areas (VTA and amygdala), and the action-related structures (striatum), the hippocampus can communicate with a decision-making network. A proposition of architecture for decision-making was recently made by Verschure and

collaborators (Verschure, 2012; Verschure et al., 2014) and we believe it summarises well the current view concerning the possible contribution of the hippocampus to the decision-making circuit. The neural structures involved in this circuit are summarised in Fig. 52. Specific roles are assigned to these structures according to the parameters they process, which can relate to the self, the world, the task or the action to be performed. The evaluation of the self state (internal variables such as motivation) is performed via the collaboration of the hypothalamus (which evaluates drives), the **ventral tegmental area** (which computes values) and the **amygdala** (which associates values to stimulus). The world state is assessed by the **hippocampus** via integration of signals from sensory cortices. The task space (e.g., rules, constraints, goals) is defined by the prefrontal cortex, mainly the **medial prefrontal cortex** and the **orbitofrontal cortex**. Finally, the **striatum** enables action selection, eventually ending up in carrying out the action through signals sent to motor cortices. The subset of these structures involved in a given situation depends on the decision-making system at stake. Indeed, the architecture is also organised in four ‘layers’ with increasing levels of flexibility: somatic, reactive, adaptive and contextual (presented in Fig. 53). Classical conditioning, for example, would require only the involvement of the somatic and reactive layer, but requires neither a highly abstract representation of the world nor of the task space. On the contrary, during goal-directed behaviour, the contextual layer takes precedence over the other layers, relying on the information they provide and on information stored in long-term or short-term memory to produce a behaviour adapted to the situation. Central to this framework is the notion of goals, which are defined in the contextual layer as an ‘amalgamations of sensory, affective and action states, stored in different memory systems and defined on the basis of the interaction of the agent with its varying and often conflicting needs with its dynamic environment’. We will see in the next – and last – introductory chapter that the nature of the hippocampal contribution to the goal representation, as of today, requires clarification.




Fig. 52 (above): Schematic view of the brain structures involved in decision-making.

Brain architecture representing neural substrates of goal-directed behaviour in relation with the processing of parameters presented in Fig. 53. The arrows represent anatomical connections. Not all connections and structures are represented. In this view, the hippocampus is involved in processing the organism’s world state space parameters at the contextual level.

VTA: Ventral Tegmental Area
 mPFC: medial Prefrontal Cortex
 OFC: Orbitofrontal Cortex
 From Verschure et al., 2014.




Fig. 53 (left): Main parameters involved in decision-making.

These parameters are organised according to two dimensions: *i)* according to the nature of information, in world, self, task or action parameters; *ii)* according to the complexity of processing, in somatic, reactive, adaptive and contextual layers. Information from the world and the self is processed through the system and integrated in the layer involved in decision-making, eventually giving rise to action. Autobiographical memory, the only type of memory that we did not mention, basically concerns a combination of episodic and semantic memories about oneself (Bluck, 2003).

Adapted from Verschure et al., 2014.

Chapter 5 – Hippocampal place cells: beyond spatial memory

According to O'Keefe (1979),

*“a **place cell** is a cell which constructs the notion of a place in an environment by connecting together several multisensory inputs each of which can be perceived when the animal is in a particular part of an environment.”*

This last introductory chapter will focus on hippocampal place cells with the first goal to expose to what extent they could ‘construct the notion of space’, i.e., represent space. Second, we will attempt to link their activity with notions reviewed in the previous chapters, such as spatial memory and spatial decision-making. The anatomical overview performed in the previous chapter indicates that the hippocampus could communicate with structures from the decision-making circuit. We will see that electrophysiological studies indicate that the location-specific firing of place cells could be influenced, under certain conditions, by parameters of the decision-making domain such as the presence of a goal at a specific location. Such an encoding of goal locations by the place cell population could be the bridge through which spatial and decision-making structures communicate.

5.1 Place cells as an allocentric representation of space

5.1.1 Anatomical localisation of place cells

Place cells are pyramidal cells from the CA1 and CA3 fields of the hippocampus. The dentate gyrus also contains cells with similar firing properties that are not pyramidal but can be considered place cells depending on the authors (Hartley et al., 2014; Poucet et al., 2014). Place cells can be found along the entire dorsoventral axis of the hippocampus (Jung et al. 1994; Poucet et al, 1994). Interestingly, the size of place fields (the location in the environment which triggers firing) increases along this axis (Fig. 54, Jung et al., 1994; Kjelstrup et al., 2008), which parallels the increase in grid cells scale that was evidenced along the dorsoventral axis of the entorhinal cortex (see Fig. 43, p.66). As an example, in the 18-m long corridor used by Kjelstrup and colleagues, the diameter of place fields of cells from the ventral hippocampus could reach 8 to 10 meters. Most place cells studies generally record in the dorsal hippocampus, since lesion studies showed that the dorsal hippocampus is more strongly involved in spatial processing (Moser et al., 1993).




Fig. 54: Dorsal and ventral place fields.

Left: example cells recorded from the dorsal hippocampus in a foraging task in a 68 x 68 cm environment. Firing rates range from blue (no firing) to red (maximum firing).

Right: cells from the ventral hippocampus recorded in the same conditions.

Adapted from Jung et al., 1994.

Contrary to primary sensory or motor areas in the brain, there is no topologic relationship between the position of the pyramidal cells in the hippocampus and the proximity of their respective place fields in an environment (Wilson & McNaughton, 1993; Redish et al, 2001; Thompson and Best, 1990).

5.1.2 Place fields and the coding of space

In the relatively small laboratory environments commonly used, most place cells express a unique place field. This was originally one of the main characteristics of place cells. However, in bigger and more realistic environments, the majority of place cells express several place fields, as presented in Fig. 55 (Fenton et al., 2008; Park et al., 2011). These cells were recorded in a random foraging task, which consists in sending food pellets in random places in the arena so that the animal covers homogeneously the environment. Contrary to grid cells, when cells express multiple place fields, they are irregularly spaced. The field size also tends to increase in bigger environments (Kjelstrup et al., 2008; Park et al., 2011).

The population activity of recorded place cells allows to accurately (with an error smaller than 5 cm) estimate the current position of the animal's head in the environment, whether using single-field or multiple-field cells (Wilson and McNaughton, 1993; Jensen and Lisman, 2000; Fenton et al., 2008). The existence of multiple fields indicates that the neural code subserved by place in these cells is more a sparse distributed than a local code (see 1.2.2, p.11). The spatial firing is independent from the direction of the rat in open field environments (Muller et al., 1994; O'Keefe, 1979). Thus, place cell firing could be said to represent the position of the animal in an allocentric reference frame. However, in paradigms that constrain the movements of rats such as linear tracks or maze arms, place fields are usually directional; namely, their firing will be different in the inward and outward run on the arm of a radial maze (McNaughton et al., 1983; O'Keefe and Recce, 1993; Muller et al., 1994).




Fig. 55: CA1 and CA3 place cells in small and large environments.

Firing maps of two sets of five simultaneously recorded place cell from either dorsal CA1 or CA3. They were recorded during random foraging in a 76-cm diameter cylinder and a 1.8 x 1.4 rectangular box.

The colour of pixels indicates the firing rate for a given position (ranging from dark blue to red). White pixels indicate unvisited bins. The numbers stand for the lowest value of firing for red bins.

From Park et al., 2011.

In a given environment, place fields usually cover the whole environment in a rather homogeneous manner with a tendency to concentrate near the borders (Muller et al 1987; Hetherington and Shapiro, 1997).

5.1.3 Place fields stability

If no changes are made to the recording environment, place fields location will remain highly stable upon successive exposures to this environment (Thompson and Best, 1990). In this study, the authors monitored the activity of CA1 place cells for several months and showed that the location of their place field was very similar across days (Fig. 56). Pharmacological and genetic studies suggest that the long-term (but not short-term) stability of the spatial firing require long-term potentiation (Kentros et al., 1998; Renaudineau, 2009).




Fig. 56: Stability of a place field.

This unit was recorded on a radial maze. Bigger dots indicate detected place fields. The numbers indicate the days since the beginning of the recording. Adapted from Thompson and Best, 1990.

5.1.4 Place cells and multisensory integration

A number of studies manipulated different types of cues, whether idiothetic or allothetic, to try to understand ‘why place cells fire where they fire’ (O’Keefe and Conway, 1978). Similarly to behavioural studies assessing the relative importance of these different cues (see chapter III, p. 31), electrophysiological experiments tend towards the conclusion that place cell firing is mostly controlled by allothetic cues, visual information in particular, but can be supplemented by idiothetic cues when necessary. In any case, their signal seems to arise from the multisensory integration of spatial information.

5.1.4.1 Control by allothetic cues

The first experiment testing how place cells were influenced by vision was performed by O’Keefe and Conway (1978) who recorded place cells in a T-maze. The maze was surrounded by distal cues. Coherent rotation of the distal cues when the rat was removed from the environment yielded coherent rotation of the place fields. This result was reproduced with proximal cues (a cue card on the arena wall) by Muller and Kubie (1987) in a circular arena while rats were foraging for food pellet (Fig. 57).




Fig. 57: Place cells are controlled by visual cues.

Example of a dorsal CA1 cell recorded in a circular arena surrounded by curtains. The firing rate of this cell is indicated by colours ranging from yellow (no firing) to purple (maximum firing) with intermediary colours being orange, red, green and blue. The black curved line indicates the position of a cue card polarising the environment. The cue card was rotated of 90° between session A and B and rotated back to its original position between B and C. Note how the place field rotates following cue rotations. From Muller & Kubie, 1987.

Importantly, O’Keefe and Conway (1978) showed that a subset of the cells recorded in an environment surrounded with distal cues still fired in the same location after removal of some of the

cues. Thus, place cells are not mere responses to specific cues but rather integrate information from multiple cues. Interestingly, when a cue card is the only item polarizing a circular environment, if it is removed, most cells will still fire but they will rotate by an unpredictable angle (Muller and Kubie, 1987).

The geometry of the environment can also influence place cell firing. When recorded in a cylindrical environment, if this environment is expanded, place fields will also expand while maintaining their location and shape (Muller and Kubie, 1987). Moreover, place fields follow the stretching of the walls in square or rectangular environments (O'Keefe and Burgess, 1996). This influence of geometry could reflect an input from boundary cells (see Sec. 4.2.3, p.64). Context, as defined by non-spatial cues such as the odour or the colour of an environment, can also influence place cell activity (Jeffery and Anderson, 2003), showing that olfactory information is used by place cells.

Both distal and proximal cues can influence place cell firing (see Knierim and Hamilton, 2011 for review). This means that place fields will 'move' their preferred location when both types of cues are displaced, but generally in coherence with behavioural results. When a rat cannot use a configuration of items to locate itself, the activity of place cells will not be altered by a rotation of this set of items (Cressant et al., 1997). Moreover, the relative importance of cues used by place cells does not seem to be 'hardwired' but rather dependent on the context (Shapiro et al., 1997; Renaudineau et al., 2007) and the reliability of cues (Jeffery, 1998; Jeffery and O'Keefe, 1999; Knierim et al., 1995), in line with the flexibility of behaviour demonstrated in control rats (Poucet et al., 2011). Place fields can also be affected by the insertion or displacement of objects or barriers (Muller and Kubie, 1987; Rivard et al., 2004; Lenck-Santini et al., 2005). The olfactory cues that can influence place cell firing can be of two types: either experimenter-controlled cues, or self-deposited cues. Experimentally controlled cues positioned in a geometrical (square) configuration increased the precision of the spatial coding of place cells in rats foraging in darkness. Moreover, rotation of the odour cues triggered a consistent rotation of place fields (Zhang and Manahan-Vaughan, 2013). However, self-deposited olfactory cues were shown not to be sufficient to stabilise the place cell representation of mice exploring an environment (Aikath et al., 2014).

5.1.4.2 *Influence of idiothetic cues*

When the amount of available allothetic cues is reduced, e.g., in darkness and without prominent olfactory cues, place cells were shown to maintain their positional firing provided that the light was switched off after the exposure to the environment (Quirk et al., 1990). However, when a rat is put in the environment under darkness conditions, the position of most place fields is modified. Place fields in the dark and without olfactory cues show a progressive shift of their field (Save et al., 2000), possibly reflecting the error accumulated in the absence of allothetic cues, and were shown to be less stable (Markus et al., 1994). Another study showed that an inactivation of the vestibular system completely disrupted the firing of place cells (Stackman et al., 2002). In a study where place cells were recorded with a decreasing amount of allothetic sensory cues, half of the place cells stopped firing in the condition where visual and olfactory cues were removed (Poucet et al., 2000).

Save and collaborators recorded place cells from blind rats in an environment where objects enabled to disambiguate location (Save et al., 1998). Their place cells had similar fields to those of normal rats, albeit with a slightly reduced mean firing rate. Interestingly, in this task, no cell was active until the rat had contacted one of the objects, and all of the recorded cells were active once contact had been made with all available objects. Thus, tactile information also seems to be used by place cells in combination with idiothetic information. Sharp and colleagues (1995) recorded place cells while modifying either proximal cues (walls of the environment) or vestibular cues (by rotating the floor). In the dark, if the floor rotation was sufficiently fast to be detectable by the rats, place fields were stable. However, slow rotations caused place fields to drift. Finally, a few experiments recorded place cells in virtual reality, either in mice with their head immobilised (Harvey et al., 2009; Dombeck et al., 2010; Chen et al., 2013) or in rats with their body immobilised (Ravassard et al., 2013). Clear and spatially focused place fields can be found in virtual reality conditions. Hence, visual cues (either from distal landmarks or from optic flow) are sufficient to generate a spatially-selective activity from place cells. Ravassard and collaborators more specifically addressed the difference between virtual reality and ‘real world’ place fields. They found that fewer cells were active in virtual reality. The fields in both environments were directional, which is common in linear tracks. However, in virtual reality, cells coded for distance with respect to the departure point rather than absolute position. This is actually surprising, as one would expect that distance information would rather be provided by the self-motion system than by vision. Yet, optic flow was also available in virtual reality and maybe its influence over the place cell activity was increased with respect to other, absent, self motion information. The authors indeed propose that competitive and cooperative interactions of different sensory sources are engaged for the control of place cell activity.

Overall, quite similar to results from behavioural studies indicating that animals can take into account several types of cues to navigate (see Sec. 3.1.2, p. 32), the place cell representation seems to incorporate and combine information from many types of cues, both idiothetic and allothetic. Idiothetic cues by themselves can be taken into account to maintain and update the place cell signal for short periods of time. However, these cues need to be used in combination with allothetic cues to guarantee a long-term stable representation. The hippocampus is well positioned to perform multisensory integration since it receives multimodal information from its entorhinal inputs (see Fig. 46, p.69).

5.1.5 Place cell firing variability

Place fields are generally not directional. However, experience can make them directional if the recording environment is linear (Muller et al., 1994; Markus et al., 1995; McNaughton et al., 1983). Moreover, within a few traversals of a linear track, place fields were shown to shift backwards (Mehta et al., 1997, 2000). This phenomenon was interpreted as Hebbian plasticity that would strengthen synapses connecting place cells with close fields; thus, the firing of a given cell, after such strengthening, would more easily trigger the firing of the subsequent cell in the sequence. We will see that this mechanism is probably reinforced by the theta sequences mechanism (to be explained in Fig. 62, p. 93). Interestingly, in other conditions, place fields can shift forward with experience (e.g., in a continuous alternation task, Lee et al., 2006).

Another kind of firing variability at the single cell level is the **overdispersion** phenomenon (Olypher et al., 2002). This term is used to describe the variability of firing of a place cell for similar trajectories through the place field (contrary to the above-mentioned variability that concerns trajectories of opposite direction). Indeed, each time a rat passes through the place field of a cell, this cell does not necessarily fire, and when it does, there is a large variability in the number of spikes emitted (Fenton and Muller, 1998). The overdispersion of place cell firing could be seen as noise but might actually reflect the processing of different types of information, or a change in the focus of attention, as we will see throughout this chapter.

5.2 Population activity of place cells

In a given environment, the activity of a given place cell, through the temporal stability of its firing, its spatial selectivity, and its multisensory nature, could be said to represent an allocentric place (O'Keefe and Nadel, 1978). At the population level, the subset of place cells active in a given environment, in combination with the location of their fields, appears to be a unique code for this specific environment.

5.2.1 Place cells are not alone

The population of pyramidal cells in the hippocampus can be split up in place cells and silent cells. **Silent cells** are hippocampal pyramidal cells that fire few or no spikes in an environment (Epsztein et al., 2011). They are generally not mentioned but they are important for our work, as will be seen in our experimental results (see Sec. 8.2, p.156).

In any environment, approximately 40% of hippocampal pyramidal cells will show location-specific firing and will be categorised as place cells (Wilson and McNaughton, 1993; Guzowski et al., 1999). The remaining 60% are silent cells. These proportions can vary with the size of an environment since the number of cells expressing a field increases with the size of the environment (Park et al., 2011). Thompson and Best (1989) were the first to study silent cells, although these cells had been mentioned previously by Muller and colleagues (1987). The study of silent cells was triggered by the observation that brain slices staining following electrophysiological recordings often evidenced many cell bodies around the electrode tip; yet, only a subset of neurons expressed spontaneous firing during the recordings (Thompson and Best, 1989). Moreover, the authors often recorded cells spontaneously active in the anaesthetised animal that did not show any activity once the animal was awake, although these cells spontaneously fired again during sleep or anaesthesia (Thompson and Best, 1989). Thus, the authors isolated cells during anaesthesia and recorded them in the freely-moving rat (Thompson and Best, 1989). Over 63% of the hippocampal pyramidal cells they recorded remained silent during spatial exploration, which correspond to the ratios reported above. Silent cells were originally found to have the same electrophysiological characteristics under anaesthesia than place cells (Thompson and Best, 1989). However, a slight difference appears when the animal is emerging from anaesthesia: cells that would later become place cells more easily fired action potentials in bursts (Epsztein et al., 2011). The authors evaluated the firing threshold of both place and silent cells and found that during navigation, place cells had lower thresholds (i.e., they would

fire more easily). Lee, Lin and collaborators (2012) later clarified the parameters that make place cells fire and not silent cells. They performed intracellular recordings of silent cells while a rat was moving around an oval track. The example from one cell is shown in Fig. 58. At first, the cell was silent and this was the case for 10 laps (left panel), until the experimenter injected a small depolarization current into the cell. Then, a firing activity with spatial characteristics similar to those of a place field emerged (right panel). This activity was present as long as the injection was continued and ended together with the injection.




Fig. 58: From silence to place.

Upper panels: spikes from the recorded CA1 pyramidal cells (red) superimposed on rat's trajectory (gray).

Bottom panels: firing rate map of the recorded cell.

Left: data collected before injection of a constant 83 pA current into the soma.

Right: data collected after the beginning of the injection of the current.

Adapted from Lee, Lin et al., 2012.

Although this is but one example, and a single 'place field' only appeared on only two cases out of the ten recorded cells, this study shed light on the possible determinants of place cell firing. As stated by the authors, this supports the idea that most CA1 pyramidal cells receive spatially tuned inputs but these inputs do not trigger firing in silent cells because they are not sufficient to reach the firing threshold. This is in line with the above-mentioned results from Epsztein and collaborators (2011) who showed that silent cells had higher spiking thresholds. Thus, the role of CA1 would be *"not only to create spatial tuning, but to choose which subset of neurons should be active within a specific environment"* (Lee, Lin et al., 2012).

In the hippocampus, other cell types coexist with pyramidal cells, namely, glial cells and interneurons. At least 21 different types of inhibitory interneurons have been identified, which have various firing properties and project to pyramidal cells in different ways (Klausberger and Somogyi, 2008). Much could be said on them; the interested reader is referred towards a subset of recent studies (Wilent and Nitz, 2007; Gloveli, 2010; Dupret et al., 2013). Glial cells were also shown to interact with neural cells in the brain. In particular, astrocytes can interact at the level of the synapse and modulate synaptic strength (Araque et al., 1999, reviewed in Perea et al., 2009).

5.2.2 The importance of exploration

Exploration is thought to be used to build cognitive maps (O'Keefe and Nadel, 1978). Depending on the cells, the spatial firing of place cells can appear immediately in a new environment or after several minutes of exploration (Hill, 1978; Wilson and McNaughton, 1993). With time, the spatial precision of the firing increases: Wilson & McNaughton report that the error in position estimation in a new environment using the active cells diminishes with time, for a subset of rats. Interestingly, place cells recorded in the dark were recently found to accumulate errors with time. However, a subset of them seemed to recalibrate using border information (Zhang et al., 2014). It was recently demonstrated that head-scanning events, in a new environment, were reliably correlated with subsequent creation of a place field (Monaco et al., 2014). This means that the occurrence of a head-scanning event at a given lap on a circular track predicted the emergence of a new place field at that position during subsequent laps. Interestingly, hippocampal interneurons might play a role in the creation of place fields in a new environment. Wilson and McNaughton (1993) noticed that the interneuron population markedly and abruptly decreased its firing when the rat was exposed to a new environment. This effect was also reported by Monaco and colleagues (2014): in their experiment, the scanning event predicting the creation of a place field was marked by a transient drop of interneuron activity. This drop was proposed to underlie a release of the interneuron inhibition on place cell population, which would allow the creation of the new 'map' of the environment by facilitating synaptic plasticity (Wilson and McNaughton, 1993).

5.2.3 Remapping

The activity of place cells seems highly stable in conditions of environmental stability. However, modifications in the external environment can alter their firing, more or less importantly. In two different environments, the place cells from the first environment can either remain active (in a similar location or a seemingly unpredictable one) or stop firing, other cells can start firing, and the firing rates can also change. This is called remapping (Muller and Kubie, 1987). Remapping can be of different kinds (See Poucet et al., 2000; Colgin et al., 2008 for reviews):

i) **global remapping**: all place fields will change position or disappear while other fields from other cells can appear (Leutgeb et al., 2005); changing the colour of a cue card on the wall of an arena (in the absence of the rat) is sufficient to trigger global remapping (Bostock et al., 1991). Changing the shape of an environment also provokes global remapping if the rat was previously trained to dissociate the two environments (Muller and Kubie, 1987; Wills et al., 2005). Also, a change of odour can be sufficient to trigger global remapping (Anderson and Jeffery, 2003). An example of global remapping generated by recording in identical arenas located in two different rooms is presented in Fig. 59.

ii) **rate remapping**: modification of the firing rate of the active fields without altering their preferred position (Leutgeb et al., 2005); rate remapping can be triggered when the colour of the recording apparatus is modified without changing distal cues (Leutgeb et al., 2005). By contrast, if the rat is recorded in the same box but in two different rooms, this will trigger global remapping (Leutgeb et al., 2005).

iii) **partial remapping**: a subset of cells remaps while the rest remains stable (Anderson et al., 2006).

iv) **local** remapping: similarly to partial remapping, only a subset of the fields remaps, but this change occurs in the close vicinity of the item that caused remapping. This can occur when objects are repositioned (but not substituted) within a well-known environment (Lenck-Santini et al., 2005). Interestingly, hippocampal lesions alter the exploratory response following object repositioning but not object substitution (Save et al., 1992). Local remapping can also occur when the topology of the environment is modified, for example by addition of barriers or creation of shortcuts (Muller and Kubie, 1987; Alvernhe et al., 2008, 2011).

Overall, the parameters whose modification can trigger remapping can be classified in two types: spatial (geometry, landmarks, topology) or contextual (colour, odour) (Anderson and Jeffery, 2003; Jeffery and Anderson, 2003).




Fig. 59: Global remapping.

Rate maps of a set of simultaneously recorded CA3 place cells in the same box placed in two different environments. Note that 16 silent cells were also detected in one or the other environment but are not shown. Adapted from Leutgeb et al., 2005.

In specific conditions, when the rat is brought from a familiar environment to a new environment, the setting up of the final activity of cells will take more time (Bostock et al., 1991; Lever et al., 2002a). This duration can differ between cells. Remapping and LTP seems tightly linked as the long-term stability of newly developed fields is altered by pharmacological blockade of LTP (Kentros et al., 1998; Renaudineau et al., 2009; Rotenberg et al., 1996, 2000). A common feature of remapping is that the whole population of recorded cell do not generally behave coherently. This could be how the same population of cells codes different environments while keeping track of their similarities.

The remapping phenomenon underlies a specificity of the hippocampal place cell population, which is its heterogeneity. This is also exemplified in the fact that a subset of cells can be tied to a specific subset of cues. Overall, we can interpret this as reflecting the plasticity of the hippocampal system, coherently with its role in flexible and adaptive behaviour.

5.2.4 Pattern completion and pattern separation

Marr (1971) was the first to suggest that different parts of the hippocampal formation could support different memory mechanisms. Namely, the dentate gyrus would perform pattern separation while CA3 would enable pattern completion (Lever et al., 2002b; Colgin et al., 2008). Indeed, cells in the dentate gyrus are far more numerous than those in the entorhinal cortex or in CA3, and CA3 cells present numerous collaterals that could support an attractor mechanism (Leutgeb et al., 2005; Poucet and Save, 2005, Wills et al. 2005). The fact that a new representation at the population level takes more time to be set when two environments are similar (Lever et al., 2002a) could underlie the creation of a new basin of attraction resulting from pattern separation between the stored environment and the new one (Poucet and Save, 2005) – see Sec. 1.2.5.2, p. 19, and Fig. 11 p. 20 for a reminder of pattern completion, pattern separation and attractor networks. Concerning the differences between CA3 and CA1, Guzowski and collaborators (2004) proposed that CA3 would react to environmental changes in a nonlinear fashion while CA1 represented changes in a more linear way. The idea that CA3 would nonlinearly transform its inputs is in line with its putative attractor-like nature. On the contrary, CA1 would have more heterogeneous response patterns and would be more influenced by external (i.e., mainly entorhinal) inputs.

5.2.5 Replay and sequences

5.2.5.1 Replay

Replay is an interesting phenomenon that exemplifies the “representational” nature of place cells. During replay, sequences of place cells that were previously activated in reality are compressed and repeated in the same order. Replay was first evidenced during sleep (Louie and Wilson, 2001; Lee and Wilson, 2002) and later during the awake state (Foster and Wilson, 2006). Replay events can be termed forward or backwards, depending on the order with which cells fire. Both types are presented in Fig. 60.




Fig. 60: Forward and reverse replay.

The activity of 13 simultaneously recorded cells from the dorsal hippocampus of a rat is indicated either before (left), during (middle) or after (right panel) a single run on a linear track with water reward at the end. The left and right panels are magnifications of the activity presented in the middle panel. The velocity of the rat is shown in the lower panel. From Carr et al., 2011, originally adapted from Diba and Buzsáki, 2007.

Replay often occurs during **sharp wave / ripples**, which are local field potential events composed of large negative waves and rapid oscillations (Buzsáki, 1986). They can take place during sleep or quiet wakefulness. Replay could be involved in the formation or consolidation of memory by allowing the repetition of experienced sequences on a fast timescale (Girardeau et al., 2009; Ego-Stengel and Wilson, 2010; see O’Neill et al., 2010; Girardeau and Zugaro, 2011 for reviews).

5.2.5.2 Phase precession and theta sequences

A prominent marker of hippocampal activity, at the extracellular level, is the existence of **theta rhythm**. It is a fluctuation of local field potential at a characteristic frequency of 7-9 Hz, correlated with type I movements (Vanderwolf et al., 1975), i.e., locomotion, orienting, rearing or exploratory sniffing behaviours. Theta rhythm is also strongly present during REM sleep (Winson, 1972). Hippocampal theta rhythm seems to depend on the medial septal area as it is eliminated following lesions or inactivations of this area (for review, see Buzsáki et al., 1983). Theta rhythm and place cell activity are tightly linked, as evidenced by the phase precession mechanism.

Phase precession refers to a shift in synchronisation of place cell firing with respect to theta: when a rat enters the place field of a given place cell, this cell will spike at late phases of the theta cycle, and spiking shifts to earlier theta phases as the rat moves forward through the place field (O’Keefe and Recce, 1993; Skaggs et al., 1996; see Fig. 61).




Fig. 61: The phase precession phenomenon.

Firing of a CA3 place cell in relation with the phase of theta cycle. The **first line** indicates spikes of the cell as a function of time while the rat was running through the place field. Note that the cell fires six separate bursts. The **second line** indicates the phase of each spike as a horizontal line. Note the decreasing phase for each burst. The **last line** represents the hippocampal theta activity. Note how each burst of the cell is emitted in a given theta cycle and with decreasing phase as a function of time.

Adapted from O’Keefe and Recce, 1993.

The firing phase appears to code for the distance travelled through the field and provides additional spatial information compared to place cells firing rate alone (Jensen and Lisman, 2000). Interestingly, it depends on experience: the phase precession effect is stronger once the animal is familiar with the environment (Mehta et al., 2002). Note that Wilson and McNaughton (1993) had also found that the precision of decoding from firing rate alone improved with exploration. First discovered in linear tracks, phase precession is also expressed in open field environments (Skaggs et al., 1996; Huxter et al., 2008; Jeewajee et al., 2014), by place cells, dentate gyrus cells, grid cells from layer II of the entorhinal cortex (Kjelstrup et al., 2008; Jeewajee et al., 2014) and others such as striatal cells (Malhotra et al., 2012). Interestingly, phase precession is also expressed in virtual reality (Ravassard

et al., 2013). This phenomenon is an example of temporal coding, which demonstrates that both rate and temporal coding coexist in the same brain structure.

One of the consequences of phase precession, at the population level, is the compressed representation of place sequences within a given theta cycle (see Fig. 62). When a rat moves forward, a subset of his place cells (those which code for the current location) will fire, in the order of the places they code for. Because close place fields overlap, and thanks to the phase precession phenomenon, the firing of these multiple cells will remain, within a given theta cycle, in the order of the sequence of places crossed. Such **theta sequences** are repeated for each theta cycle.




Fig. 62: Theta sequences.

Upper panel: Representation of the firing fields of four place cells with the rat trajectory as a black arrow.

Lower panel: Theta rhythm and cell firing. If each cell shows phase precession, the spikes of the four cells will be emitted as ordered sequences along the path of the rat and also in the same order within each theta cycle, under the form of compressed sequences, i.e., theta sequences.

Adapted from Skaggs et al., 1996.

Theta sequences had been predicted by Skaggs and collaborators (1996) and were demonstrated by Foster and Wilson (2007). One of the functionalities of such repeated sequences could be to reinforce the connections between cells 'representing' adjacent places by the way of synaptic plasticity, probably spike-timing dependent plasticity. Indeed, the time scale of theta sequences is compatible with STDP (Gupta et al., 2012). Gupta and collaborators showed that sequences depend on rat's speed: when the animal accelerates, the sequences go forward, activating cells that represent places farther away. On the contrary, when the rat decelerates, the sequences begin farther behind the rat.

In conclusion, the combined population activity of hippocampal place cells and silent cells could be said to encode environments, on top of the activity of single place cells that encodes location within an environment. The multiple properties of place cells both at the single cell and the population level converge towards the idea that they might be the instantiation of a neural representation of space. However, the experiments we illustrated were usually done in foraging conditions, i.e., in tasks where the subject does not necessarily need to locate itself or, in any case, that do not require intensive spatial processing. From lesion studies, we know that the hippocampus is preferentially involved when flexibility is needed and where a goal must be located with respect to distal cues. Several studies addressed the existence of a link between place cell activity and behaviour.

5.3 Place cells and goal-directed behaviour

If place cells represent the memory of space, disturbing their activity should impair spatial memory-based behaviour. Coherently with results from hippocampal lesion studies, links between place cell activity and behaviour should depend on the strategy used: modifications of the place cell firing should not impair response strategy while it should impair map-based strategies.

5.3.1 Place cells and attention

A few studies demonstrated that the level of attention required for a task, which is generally linked to the level of spatial processing required, has an effect on place cell activity. For example, Zinyuk and collaborators (2000) recorded place cells from rats performing two different tasks in the same environment: either random foraging or continuous place navigation. A constant rotation of the arena disrupted in an unpredictable way the place fields of most rats recorded in the foraging task. By contrast, their place fields were more localised and coherent in the navigation task, and their location was generally more predictable, being either influenced by distal cues or a combination of both distal and proximal cues. Fenton and collaborators (2010) also recorded rats in different foraging or navigation conditions in the same apparatus. They concluded that place cells could process spatial information from different subset of cues. Moreover, in the goal-directed task they used, the overdispersion of place cells was reduced. The authors proposed that in baseline conditions, the place cell representation would alternate between different reference frames, whereas in a goal-directed task, attention was more focalised on a specific reference frame, and so would be the place cell activity (Fenton et al., 2010). The effect of attention was also assessed in mice by Kentros and colleagues (2004). They recorded place cells from the dorsal CA1 of mice in four different conditions: *i*) no explicit task, *ii*) random foraging, *iii*) random foraging in a new environment and *iv*) a spatial task which required locating an unmarked goal zone. Remarkably, the long-term stability of place fields increased with the level of attention required from the task (Fig. 63).




Fig. 63: Long-term field stability is a function of attention.

Four groups of mice were used, one per condition. Both the fields from the new environment and the spatial task were significantly more stable than those from the no task group. Moreover, fields from the animals recorded in the spatial task were significantly more stable than those recorded in the three other conditions. The comparisons were made between recordings separated by at least 6h.

From Kentros et al., 2004.

In this task, some of the mice were not performing as well as others. Interestingly, place cells from these animals were dramatically less stable than good performers. Kentros and collaborators conclude in these words:

“not only do place cells look like memory traces, they act like them as well: both seem to require attention for their storage and faithful recall”.

5.3.2 Place cells and navigational decisions

The first experiment addressing a direct link between the location of place fields with respect to distal cues and the navigational decisions of rats was performed by O'Keefe and Speakman (1987). In a cross maze task testing reference memory, rats had to go to the goal arm by using its position with respect to distal cues. Once trained, the rats were able to reach the goal arm even if the cue card was removed once it had seen it. In conditions where the cue was absent upon entry of rats on the maze, there was a marked decrease in performance. However, in these conditions, the place where the rat searched for the goal was still at the same relative location with respect to the active place fields, compared to the control condition. Later, Lenck-Santini and colleagues (2001a, 2002) performed a series of studies where the position of place fields was coherent with the navigation decisions of rats, either in a Y-maze or a continuous navigation task (see p.75 for the protocol of this task). This was the case when rats performed correctly but also when they made mistakes in goal localisation. Interestingly, in the 2002 study, one of the rats always searched for the goal in the proper location independently from the position of his place fields.

In divergence with these results, Jeffery and collaborators (2003) used colour-induced remapping to link the activity of place cells with performance. They trained rats in a hippocampal-dependent task to locate a goal using distal cues. After training, the colour of the recording box was changed and the performance of rats was assessed in this 'new' environment. Although rats' performance remained above chance, the population of place cells showed global remapping. Thus, in this experiment, place field locations were not consistent with performance. This is similar to the one rat from Lenck-Santini and colleagues' study (2002) which performed independently from the position of its place fields. At least three interpretations of these results are possible: *i*) hippocampal place cells are actually not used for navigation; *ii*) hippocampal place cells were not used for navigation under the testing conditions; a response or guidance strategy could be used by the rats at that stage of learning, as they were always placed at the same departure position in the environment (Poucet et al., 2004); or *iii*) some regularities in remapping were not detected but could still enable proper navigation; this could perhaps underlie the existence of a more 'global' map outside of the hippocampus (see Jeffery, 2008)⁸. In a follow-up experiment, Anderson and collaborators (2006) used the continuous navigation task to address the same issue. Similarly to the water maze task, the optimal strategy in the continuous navigation task is to use a place strategy. After training, a change in colour of the apparatus was performed so as to trigger remapping. However, only partial remapping was observed, and the performance of rats was not affected. This was interpreted by the fact that the activity pattern of place cell could simultaneously be sufficiently stable to represent the relationships between relevant spatial cues used for navigation and sufficiently flexible to detect environmental novelty.

Kubie and colleagues (2007) attempted to clarify the extent to which place cells could be linked to behaviour and especially to locating a goal. First, they recorded CA1 and CA3 place cells from rats

⁸ For example, subicular cells with spatial firing were shown to remain unchanged following manipulations that caused CA1 to remap (Sharp, 1997, 2006)

performing a random foraging task in a circular environment (Fenton et al., 2000a). The environment was polarised by two cue cards attached on the walls. In specific conditions, the cards were rotated and this triggered a displacement of the place fields. The authors proposed a model that could accurately predict the displacement of the fields as a function of the displacement of the cue cards (Fenton et al., 2000b). In a subsequent experiment, they trained rats in a continuous navigation task in the same circular environment polarised by the same cue cards (Kubie et al., 2007). In this task, the estimated location of the goal, for the rat, is indicated by the place where it stops to trigger the pellet dispenser. Similar manipulations of the two cue cards were performed, which caused displacements of the location where the rats stopped, i.e., their estimated goal location. The main point of this study is that the model used to predict the displacement of place fields also accurately predicted the displacement of estimated goal location shown by rats' behaviour.

Overall, in tasks which favour the use of a place strategy, place fields' positions seem to be coherent with the putative estimation of the rat position (or, of its goal) as indicated by its behaviour. The conditions where such connection is observed parallel the conditions where hippocampal lesions impair performance. These results evidence a possible role of the place cell 'map' in supporting the estimation of position of a subject. However, these clues are nothing more than correlations.

5.3.3 Place cells can represent past and future positions

5.3.3.1 *Prospective or retrospective coding*

Several studies showed that hippocampal neurons could fire differently in the same place according to 'experience', that is to say, what has just happened but also what is just about to happen (Ferbinteanu and Shapiro, 2000; Frank et al., 2000; Wood et al., 2000; Ainge et al., 2007); however, this was not observed by Lenck-Santini and colleagues (2001b)). In a continuous alternation task (see Fig. 32, p.52), Wood and collaborators (2000) recorded place cells in the central stem that fired differentially depending on the turn that would be performed next (or on the turn previously made, as both were highly correlated). This was interpreted as memory for specific episodes (a left turn or a right turn trial). In a follow-up study, Ainge and collaborators (2007) reported that a lesion of the entire hippocampus did not impair performance in this task. Only when adding a delay in the central stem would the task become hippocampal-dependent. The differential activity of place cells was observed in both the delay and non-delay version of the task. The authors argue that during the delay, this differential activity was caused by inputs from the medial prefrontal cortex, where such delay or experience-dependent firing was observed (Jung et al., 1998; Baeg et al., 2003; see also Hok et al., 2005). In the non-delay version, the trial-dependent activity in the hippocampus could be caused by striatal inputs. The difference between the delay and non-delay version would be one of strategy: place strategy in the former, response in the latter (Ainge et al., 2007).

5.3.3.2 *Forward sweeps*

Memory and mental representations allow memorised information to be used outside the place or context of its initial presentation (Johnson and Redish, 2007). Thus, if place cells were only active when the exact stimuli that first triggered their firing were present, they could not be qualified of

support of spatial memory. However, the fact that they can be active when only a subset of environmental cues are present (e.g., with a reduced number of distal cues, or in the dark) and their remarkable stability over time in the absence of significant changes in the environment suggests that they could support spatial memory. An important feature of memory and episodic memory in particular is that the memory of a place can be recalled from any other place. Such recall can be used, for example, to perform trajectory planning or to assess possible options. Johnson and Redish (2007) showed that the hippocampal population could transiently represent possible routes ahead of the current location of the rat while it was pausing at choice points in a decision-making task. To demonstrate this, they trained rats in a multiple-T-maze choice task (see Fig. 64). This maze is composed of a central stem containing multiple intersections. At the final intersection, rats could choose between turning left or right. Only one side of the maze was rewarded for a given day, and the rat was prevented to go backwards. Thus, the final choice was said to be a high-cost choice.




Fig. 64: The multiple-T maze task.

The task consisted of four choice points (in the middle) that end with a high-cost choice to be made between going either left or right. In a given day, only reward sites from one side were active. Rats were blocked from going backwards. From Johnson and Redish, 2007.

The authors used a decoding algorithm to continuously estimate the position represented by the population activity of place cells. This estimation of position accurately represented the current position of the rat during navigation. However, when the rat paused at the high-cost choice point, the representation transiently swept forward, as if evaluating the two possible paths (Fig. 65). These ‘non-local’ representations tended to be more present during vicarious-trial-and-error (explained in Sec. 2.6, p.29). The forward sweep direction was not correlated with the final decision of the animal, indicating that it would rather be an evaluation mechanism than a decision-related one. The authors related these results to those of other studies where prospective firing of place cells was found (e.g., in the central stem of a radial maze: Wood et al., 2000; Ainge et al., 2007). At a decision point where one of the possible routes is more probable to be taken, the forward sweeps would be biased towards one option and could be interpreted as differential firing according to future decision. The discovery of forward probes led to the hypothesis that the hippocampus could be involved in decision-making by providing the ‘decision-making’ circuit with projection information related to a search process, that is to say, a virtual exploration of possible trajectories (Johnson et al., 2007).




Fig. 65: Forward sweeps by the place cell population.

At the high-cost choice point (white circle), the estimation of position from the CA3 place cell population transiently ‘sweeps forward’. The estimated position is indicated for 40 ms bins. The representation intensity is shown in colours indicating the probability of position estimation (red: high probability; blue: low probability). Adapted from Johnson and Redish, 2007.

Events such as replay or forward probes are qualified of ‘**non-local**’, in the sense that place cells will fire out of their place field, possibly to ‘represent’ their preferred position (Johnson and Redish, 2007; Pfeiffer and Foster, 2013).

Yet more recently, a direct correlation between the non-local activity of place cells and navigation decisions was demonstrated by Pfeiffer and Foster (2013). The authors recorded CA1 place cells from rats that had to elaborate paths to a goal location in an open arena. When the rat was pausing prior to initiating navigation, the population of CA1 cells generated brief sequences that encoded **spatial trajectories** from the current place of the rat towards the goal location. These sequences occurred during sharp-wave-ripples events and were found to be functionally similar to replay events. They were neither related to the previously completed journey nor to the current spatial view of the rat. Moreover, the task design controlled that new paths had to be elaborated in a proportion of trials, and sequences were also observed before initiation of these entirely new trajectories. Importantly, the sequences events strongly matched the future path chosen by the rat. Thus, place cell sequences seem to concretely instantiate trajectory planning, in support for a flexible use of the ‘cognitive map’⁹.

5.3.4 Manipulation of the ‘memory engram’

Although the above-mentioned studies add more evidence in favour of a use of the hippocampal place representation for localisation and even for trajectory planning, a necessary step to demonstrate that place cells are used by the subject to locate itself is to evidence a causal link between the activity of place cells and some aspect of behaviour related to “where the rat thinks it is”. A few studies attempted to use such a causal manipulation of spatial memory. Ramirez and collaborators (2013) combined immediate-early genes with optogenetics in mice using a fear

⁹ For better visualisation of the place cells forward probes and sequences, the reader is referred to the online supplementary videos of both of these studies (Johnson and Redish, 2007; Pfeiffer and Foster, 2013).

conditioned protocol. Mice usually respond to a fear-conditioned environment by freezing, easily providing a way to measure the memory for contextual fear. The authors first used *c-fos* (an immediate-early-gene, see p. 15) to label the subset of cells from the dentate gyrus which would get activated upon exposure to a given environment. They combined this *c-fos* approach with optogenetics such that only the *c-fos* labelled cells could be later activated by light. Mice were fear-conditioned in a new environment while the subset of cells previously labelled was activated by light. Subsequent exposure to the first environment (which corresponded to the labelled cells) triggered freezing, although mice had never been exposed to the fear-conditioning protocol in that environment. Thus, it was as if the memory of fear had been associated to the first environment by reactivation of the memory ‘engram’ (i.e., the physical instantiation of the memory in neural networks). By contrast, the similar procedure applied to CA1 cells did not trigger such ‘false memory’. The authors suggest that this could be due to the use of a population rate by the dentate gyrus while CA1 would rather rely on a temporal code. In another study from the same team, using the same technique, Redondo and collaborators (2014) managed to reverse the valence of the memory associated to a place upon reactivation of dentate gyrus cells. Interestingly, a similar procedure with amygdala cells did not change the memorised valence (see Takeuchi and Morris, 2014 for a commentary on this study).

In summary, although the above-mentioned studies relied on emotional conditioning and not goal-directed behaviour, they still provide strong evidence of a direct, causal link between the activation of a population of dentate gyrus cells in a given environment and the memory for this specific environment (see Spiers and Bendor, 2014 for a review on manipulation or enhancement of memory). The results reviewed in this section strongly support the idea that the activity of place cells is used during goal-directed navigation and perhaps even contributes to spatial-based decision-making.

5.4 Place cells and decision-making parameters

So far, we emphasised the spatial parameters that controlled hippocampal pyramidal cells’ firing. Even for non-local events such as replay or forward probes, the information thought to be encoded in place cell firing concerns space. In this last section, we will focus on studies addressing the influence of non spatial parameters, in particular those involved in decision-making such as goal or reward value.

5.4.1 Behavioural and task – related correlates of place cells

First, place cell firing is positively modulated by speed, up to a given threshold (McNaughton et al., 1983). Conversely, there is a dramatic decrease in place cells’ discharge when an animal is tightly restrained (if the rat learned that locomotion was impossible; Foster et al., 1989). Pyramidal cells can also express various other possible correlates often combined with space (e.g., Eichenbaum et al., 1987; Wiener et al., 1989; Hampson et al., 1999; Komorowski et al., 2009; Kennedy and Shapiro, 2009; MacDonald et al., 2011; Kraus et al., 2013; reviewed in Eichenbaum et al., 1999; Eichenbaum and Cohen, 2014). Generally, pyramidal cells will specifically fire for items or events that are relevant

to the current task. O’Keefe (1976) reported ‘misplace cells’ that would exhibit increased firing after removal or replacement of items. Wiener and colleagues (1989) recorded hippocampal pyramidal cells in a place task and found that almost 70% of them had movement (speed, direction and/or angular) tuning. Most of the time, this behavioural tuning was observed only when the rat moved within the place field of the concerned cell. The authors concluded that a ‘place’ cell cannot entirely be described solely from its place tuning, stating that:

‘even when the animal was in the place field, the firing rate of some cells was no higher than the out-of-field rate if the animal was moving at a non-optimal speed or in a non-optimal direction or turning angle.’

In a last example, Kennedy and Shapiro (2009) recorded pyramidal cells using different paradigms (foraging or goal-directed task) and different motivational states (hungry or thirsty). In the goal-directed task, rats needed to discriminate the goal using their motivational state (for example, if they were food deprived, they had to go left). Place cells were influenced by the motivational state of the rat specifically in the goal-directed task but not in the random foraging task. Interestingly, this modulation was independent from the actual place of the goal (which was irrelevant for proper performance).

5.4.2 Influence of the goal location on place cell activity

The location of a goal (whether cued or not) in an otherwise homogeneous environment has been shown to influence place cells’ firing in a variety of ways. However, unknown parameters seem to play a role, since a subset of studies also did not evidence any control of the firing of place cells by the goal location.

5.4.2.1 Goal-related activity of place cells

Following the discovery of the spatio-selective activity of medial prefrontal cortex cells at the goal location (reviewed in the previous chapter, p 75), and observations of data from previous studies (Lenck-Santini et al., 2002), Hok and collaborators (2007) analysed the activity of dorsal hippocampal place cells from CA1 in the continuous navigation task. The protocol is the same as the one used for the mPFC study (Lenck-Santini et al., 2001a; Hok et al., 2005) where the rat must locate a hidden goal zone to trigger the release of reward. Most place cells recorded in this task expressed **goal-related activity**, that is to say, an out-of-field firing localised at the goal (Fig. 66 A). In contrast to the activity expressed by medial prefrontal cells, the hippocampal goal-related activity was specific to navigation episodes (Fig. 66 B). Still in contrast with prefrontal cells, no overrepresentation of the goal was detected in the distribution of centroids of place cells (Fig. 66 C).




Fig. 66: Goal-related activity of CA1 place cells in the continuous navigation task.

A: Firing rate maps of four CA1 place cells during a 16 min session. The black circle indicates the uncued goal location. The colour map ranges from yellow to purple with, in order, intermediate rates indicated by orange, red, green and blue.

B: Separation of firing between navigation and foraging episodes. The navigation map regroups firing activity emitted 4s before pellet release. Foraging firing consists of the remaining activity where low speed firing (emitted when speed was inferior to the session mean speed) has been removed.

C: Spatial distribution of place fields centroids. The grey overlay indicates goal location.

From Hok et al., 2007.

Importantly, this goal-related firing was not correlated with the occurrence of sharp waves / ripples at the goal (which are tightly linked to replay, as seen p.91). A specific temporal profile of the goal-related activity was observed in the case of the place continuous navigation task (uncued goal zone), but this profile was shifted forward in time in a cued version of the task (see Fig. 67).




Fig. 67: Temporal profile of the goal-related activity of place cells.

Peri-event time histograms for the population of recorded place cells aligned to the entry into the goal zone (0s). The peak activity is shifted forward in the case of a cued goal where a circular disk is positioned at the goal location. From Hok et al., 2007a.

The following remark can be found in Hok and collaborators' study, suggesting that silent cells might play a role in goal-related activity:

“of the total sample (174 pyramidal cells), 22 cells were “silent”; they discharged only a few spikes in a session and did not have a firing field according to our criteria. The sporadic firing of these silent cells often appeared to be concentrated in the target zone, but their properties were not further characterized.”

We saw previously that an inactivation of the hippocampus impaired the spatially-selective activity expressed at the goal location by medial prefrontal cortex (mPFC) cells (Burton et al., 2009). The authors suggested that the hippocampal goal-related activity provided spatial information to the mPFC which could in return use this information to plan trajectories. Hok and collaborators (2013) subsequently tested whether an inactivation of the mPFC would have any effect on rats' behaviour or on the activity of place cells. Interestingly, the lesion did not impair rats' performance, nor did it impair the goal-related activity of place cells. However, a major effect of this lesion was seen on place cells' overdispersion, which was significantly reduced, possibly indicating a decrease in the flexibility of the place representation. Overall, these results indicate that hippocampal place cells do not need mPFC information to generate their goal-related firing. The authors proposed that this signal could underlie a temporal estimation, since the duration of rats' stay at the goal zone is important for feedback on the goal location. This hypothesis was tested, and rejected, in a follow-up study using partial extinction procedures and a cued version of the task (Hok et al., 2007b). The hypothesis put forward by the authors is that the goal-related activity would be a spatial signal indicating that the rat is at the proper location. In the literature, in addition to be cited as a goal-related signal, this activity has been assigned several roles, among which a time-related activity (Hirel et al., 2013), an anticipatory signal (Erdem and Hasselmo, 2012), a 'nonspatial' event (Itskov et al., 2012) and even a reward-related signal (Terada et al., 2013). Overall, the role and the source of this firing are still, to date, unknown. Because this 'goal-related' signal might be of major importance for goal-directed navigation and might underlie a transfer of information either from the spatial to the decision-making circuit, or the other way round, we decided that we would look more closely into the possible mechanisms subserving this signal. Thus, the study of goal-related activity is a major part of the experimental contributions that we present in the thesis.

5.4.2.2 *Other types of goal-related signal*

Aside from the above-mentioned study, goal correlates in place cell activity can be found or not depending on the tasks. First, Speakman and O'Keefe (1990) recorded place cells in a cross maze. A change in the baited arm did not visibly alter the location of place fields. Hölscher and collaborators (2003), in a radial-arm maze, did not find any overrepresentation of goal locations (i.e., the end of arms), nor did they find cells that would fire exclusively in baited arms. Rather, they found reward-related signals, as will be reviewed in the next section. Siegel and collaborators (2008) used a similar task in a square arena and did not evidence any goal-related firing either (but a shorter (1s) delay was required in the goal zone). In addition, in all the instances of the place preference task mentioned in this chapter, none of them – apart from studies from Hok and collaborators (2007a,

2007b, 2013) – mentioned goal-related activity, although they sometimes explicitly searched for it (Lenck-Santini et al., 2002; Kentros et al., 2004; Anderson et al., 2006; Kubie et al., 2007 – see also Jeffery et al., 2003). Note however than in Lenck-Santini and collaborators' (2002) study, goal-related activity was visible on the firing rate maps, and smoothing of the maps on other studies might prevent it to be seen.

Other studies did evidence forms of goal-related activity, either as a shift of place fields or as the appearance of new place fields at the goal location. In a series of tasks, Eichenbaum and colleagues (1987) recorded categories of hippocampal pyramidal cells which had other than purely spatial correlates. These were termed “goal-approach” cells. According to the authors,

“the best description of goal-approach cells is that they fire during an act of orientation toward a target of attention, regardless of its immediate egocentric perspective or of the movements required to obtain it.”

They were further described in a follow-up study (Wiener et al., 1989) where two tasks were performed in the same apparatus, to compare the possible effect of the task on pyramidal cells' activity: a spatial task and an odour discrimination task. In the place task, there was no evidence of specific concentration of the fields around the goal or the locations of interest. In the odour discrimination task, several types of cells were found, among which the above-mentioned ‘goal-approach’ cells which were tuned to either approach to the port area or approach to the cup containing water reward. Their main field was not necessarily located around the port; interestingly, they fired significantly more when in the main field in the spatial task than when the rat was performing the port approach behaviour.

Hollup and colleagues (2001b) managed to record place cells in an annular version of the water maze. Similarly to the original paradigm, rats had to reach an uncued platform to get out of the pool. However, in this version, the hidden platform was not available from the beginning and rats had to make at least one lap around the annulus before the platform would be made available (but still hidden). This task was shown to be hippocampal-dependent (Hollup et al., 2001a). The authors recorded dorsal CA1 cells of rats performing this task. They found that the number of pyramidal cells whose firing fields were in the platform segment was significantly larger than in the other segments. Fields were also more numerous in the segment preceding the platform segment than in others, non-platform segments. A subset of example place cells recorded in this task is presented in Fig. 68 A. Also, whenever the rat reached the platform, this firing was reduced for a subset of the units (see Fig. 68 B). The authors noted that theta rhythm was still present before and after climbing on the platform, with a slight shift towards lower frequencies, which was also the case in Hok and colleagues' study (2007a).




Fig. 68: Place fields in the annular water maze task.

A: Firing activity (red) superimposed on rat's trajectory for seven CA1 place cells. The platform location is indicated in green (bottom-left area of the annulus).

B: Firing before (**top**) and after (**bottom**) climbing on the platform.

Numbers indicate the peak firing rate and the location of peak activity.

From Hollup et al., 2001b.

In the same annular water maze task, Fyhn and collaborators (2002) tested the effect of a dislocation of the platform on the activity of CA1 pyramidal cells. Interestingly, they found that some cells that were previously silent at the goal location would fire when the platform was encountered at the new location for the first time. This was accompanied by a reduction in the firing of simultaneously recorded interneurons. The new activity was maintained for tens of seconds before it decayed. The authors do not interpret the activity of these neurons as a goal signal but rather as a signal transiently encoding new features of the environment, following a mismatch between expected and encountered items (similar to the misplace cells of O'Keefe (1976)).

Kobayashi and collaborators (2003) recorded CA1 and CA3 place cells from the dorsal hippocampus of rats in a task where they had to shuttle between two unmarked reward locations. As a reward, rats received intracranial stimulation of the medial forebrain bundle at the goal locations (after a 1s delay). A proportion of the recorded cells (10 / 14) modified their firing as rats learned the goal-oriented task, seemingly moving their activity towards one of the two goal locations (but not both). These patterns of activity were not seen if the rats were randomly foraging, except in cases where they behaved in a similar manner to the shuttling task. The authors concluded that the change in activity only seen in shuttling behaviour could support Eichenbaum and collaborators' (1999) view insofar as place cells, in this case, represent specific actions occurring at a specific location.

The study of Pfeiffer and Foster (2013, mentioned p.98) evidenced the existence of sequences of place cells that might underlie trajectory planning. They found that the goal location was overrepresented in the trajectory events that were initiated from a random place when planning to reach the goal. However, they did not mention specific activity of the cells at the goal location.

To conclude, the discrepancy between several studies appears difficult to solve. In a subset of studies, a displacement of the centroids (centre of mass) of place fields towards the goal location is observed, while it is not the case in other studies. Moreover, a proportion of studies evidence an increase of activity at the goal location, whereas others do not. A tentative explanation could be that tasks involving stereotyped trajectories will eventually trigger a displacement of the centroids of place fields (probably through theta sequences and plasticity phenomena). This could be the case in Kobayashi et al.'s (2003) study and Hollup et al.'s (2001b) study. Perhaps such trajectories underlie a procedural strategy to reach the goal. Conversely, in the continuous navigation task, each trajectory to the goal is virtually different, which is an indicator of the use of a place strategy, and no such overrepresentation of the centroids was evidenced. It would be interesting to know if this was also the case in the study from Pfeiffer and Foster (2013), since new trajectories had to be elaborated to reach the goal. Unfortunately, they do not report any analysis of this sort. Concerning the increase in firing rate of place cells at the goal location, this could be the trace of an increased spatial processing, since it appears in the continuous navigation task where a rat has to locate a hidden goal (Hok et al., 2007a, b, 2013), but not in a radial arm maze where, once the arm is chosen, local cues are sufficient to guide navigation (Hölscher et al., 2003). In any case, this variability in the results indicates the need for clarification of the parameters that give rise to the goal-related activity. A possible explanation could be that the goal-related activity indicates a combination of place and value information and somehow instantiates a form of reward or goal value expectation.

5.4.3 Reward – related correlates of place cell activity ?

To our knowledge, studies recording hippocampal activity do not mention place cell firing in reaction to the mere consumption of reward. For example, Hölscher and collaborators (2003) state that “the firing activity of the cells was low and did not increase after retrieval of a pellet”. However, changes in the reward contingency of a task reduce the overdispersion of place cells (Wikenheiser and Redish, 2011).

A few published studies questioned whether hippocampal cells would encode reward-related parameters. First, Hölscher and collaborators (2003) used an eight-arm radial maze while recording place cells and found that the firing rate of CA1 pyramidal cells was seemingly higher in baited than non-baited arms. However, in this task the reward was all or none and the differential activity could either indicate the baited versus non-baited status, the expected reward magnitude, or even a change in attention. To our knowledge, only one study addressed the question of reward magnitude coding in the hippocampus. In this study, Tabuchi and colleagues (2003) trained rats in a cross maze task to visit arms in a specific order, which corresponded to the order of decreasing amount of delivered reward at the end of each arm. They did not evidence any coding of reward magnitude in CA1 pyramidal cells. They did, however, observe a few neurons with ‘position-selective reward site activity’, i.e., neurons that would combine place and reward information and fire at a specific reward site. A few recent studies specifically addressed the encoding of reward probability in dorsal CA1 cells (Lee, Ghim et al., 2012; Terada et al., 2013). Lee, Ghim and collaborators (2012) combined modelling and experimental approaches in a continuous choice task and found subsets of CA1 cells whose main field activity correlated with action value. Terada and colleagues (2013) performed dual recordings

both in the amygdala and the hippocampus in a light-side discrimination task, again in dorsal CA1. They report the existence of one hippocampal cell that encoded reward expectation.

In conclusion, if the question of goal coding in the hippocampus is controversial, the question of reward coding also raises discrepancies in the literature. In any case, it seems that information about certain forms of value can be found in the place field of CA1 hippocampal cells. Could the out-of-field, goal-related activity of place cells also be a signal indicating goal or reward value?

5.5 Conclusion

From the electrophysiological results presented in this chapter, a few general conclusions on the role of the hippocampus and place cells in spatial cognition can be drawn. First, the involvement of the hippocampus in processing space appears to be clear. It holds neurons whose combined activity represents with a striking precision the allocentric position of the animal in space as well as the environment where it evolves. The activity of these neurons is controlled by elements of the environment that are relevant to spatial orientation and navigation, such as geometry, barriers or shortcuts. In addition to coding for space, place cells also appear to participate in memorising it. The population of place cells displays many properties that characterise a memory system: their activity can remain stable for months in a given environment; they can store different patterns of activity in different, although similar, environments (pattern separation); they are able to retrieve the memory for an environment even if it is not exactly similar to the stored memory (pattern completion); through replay, they could participate in memory consolidation. But the hippocampus is not only a passive repository of the memory of places. This memory of space can be flexibly used to direct behaviour in conditions that require a certain amount of spatial processing. Electrophysiological studies, when performed in animals engaged in foraging, exploring, or even inflexible navigation strategies, evidence a heterogeneous population of place cells whose global conduct in the face of environmental modifications appears unpredictable or, in any case, highly variable. However, when an animal is engaged in goal-directed behaviour, place cells display a coherent activity pattern at the individual (less overdispersion) as well as at the population level, tightly correlated with navigation decisions. These results fit in well with those from lesion studies demonstrating that the hippocampus is essential only when accurate and flexible localisation of a goal is required. Furthermore, when flexible decisions are required, place cells engage in phenomena such as forward probes and even sequence planning that reflect not only the prospective evaluation of possible options but also the final decision taken by the animal. Finally, the hippocampal population expresses a particular activity related to spatial goals whose role is still to be elucidated. Delineating the role of the hippocampus in spatial goal-directed behaviour might require a combination of approaches drawn both from the spatial and the decision-making domains. Such an approach could help clarifying to what extent the spatial representation in the hippocampus could be used for flexible decision-making. In particular, the goal-related activity expressed by place cells in a task encouraging the use of a place strategy could be the link through which spatial and motivational aspects combine to perform goal-directed navigation.

**PART TWO:
EXPERIMENTAL
CONTRIBUTIONS**

Objectives and working hypotheses

The following pages present the experimental work realised during the thesis. The objectives of our experimental work were two-pronged. First, from a behavioural point of view, we wanted to develop a task combining flexible spatial cognition with decision-making. Second, we aimed at clarifying the nature of the goal-related signal of place cells. Namely, we asked whether it was more of spatial or motivational nature. To gain insights into this issue, we defined four objectives. The first two combine behavioural and electrophysiological aspects:

- i. Assess the behaviour of rats and the electrophysiological activity of hippocampal cells in a task allowing to freely choose between **several goal locations**. From a behavioural point of view, a free choice should engage flexible decision-making from the rats, and the availability of several goal locations should require complex spatial processing. From an electrophysiological point of view, assessing the goal-related activity at multiple goal locations could help determining whether this signal is of spatial or motivational nature: a motivational signal would be identical at both zones, since it would reflect expectancy of reward, or increased attention. A spatial signal should differentiate the two goal locations.
- ii. Assess rats' behaviour and their hippocampal cells' activity in a task that requires the **estimation of goal value** for proper performance. From a behavioural point of view, this should test for the goal-directed nature of behaviour, depending on whether or not rats would adapt their behaviour as a function of the consequences of their actions. From an electrophysiological point of view, if the goal-related activity is of motivational nature, it would be more likely to vary as a function of goal value, as is the case in the VTA or in the amygdala. A spatial signal would not express such variations.

We also had two supplementary objectives that only concerned electrophysiological parameters:

- iii. Record neurons from CA3. If the goal-related signal is generated in the hippocampus, it should be visible in CA3. If it merely reflects motivational inputs, it is most likely to be visible only in CA1, which receives more subcortical inputs. In particular, CA3 does not receive input from the amygdala and few projections from the VTA.
- iv. Record hippocampal silent cells. Hok and collaborators (2007a) mentioned that a subset of cells had very low firing while being more active at the goal zone. These cells could possibly have a role in the goal-related activity.

In a first series of experiments, we aimed at defining a behavioural paradigm that corresponded to the above-mentioned constraints. In addition to using a spatial task, we included a decision-making component in the form of modifications of goal value. Chapter 6 and 7 present exclusively behavioural results. Since the experiments presented in chapter 6 did not yield any exploitable result, they will be briefly overviewed. Chapter 7 will go into more details about the main experiment of this thesis, the two-goal navigation task, and will expose the methods as well as the behavioural results obtained from this task. Chapter 8 exposes the electrophysiological results obtained from the same experiment.

Chapter 6 – Introducing outcome devaluation in spatial cognition

6.1 Objectives

The aim of this first series of experiments was to elaborate a protocol combining spatial processing with decision-making where the value given by rats to specific goals could be measured by a quantitative parameter. The final objective was to use electrophysiological recordings in the task. However, none of these protocols were eventually selected. We chose to rely on **satiety devaluation** procedures to attempt to selectively lower the value of the outcome, which is taken into account for the final estimation of goal value. Satiety devaluation (see p.7) consists in free-feeding a rat with a given reward for a certain duration so that the value of this specific reward is subsequently lowered. Satiety devaluation has rarely been used in spatial tasks, with the notable exception of the series of experiments addressing episodic-like memory in rats by Babb and Crystal (2006a) (mentioned in the introduction, see Fig. 4, p. 7). In the instrumental learning domain, it is thoroughly – and successfully – used to selectively lower the value of a reward in order to assess the goal-directed nature of a behaviour (Adams and Dickinson, 1981; see Sec. 2.2, p.24). The experiments had to match a set of constraints:

- i. at least two possible ‘spatial goals’ should be available in order to have a ‘control’ goal and a ‘modified’ goal.
- ii. the goals should be associated with different types of reward so that we could selectively modify the value of one goal by selectively decreasing the value of the associated reward.
- iii. the experiment should allow for measuring the preference of rats towards specific goals in order to gain insights into their estimation of goal value.
- iv. the experiment should be adaptable to electrophysiological recordings.

We elaborated seven different protocols using two apparatus: the eight-arm radial maze and the continuous alternation maze. These protocols never evidenced any consistent effect of the devaluation among rats and their results were not further exploited. Each of them will be presented very briefly.

6.2 Radial-arm maze tasks

As previously stated, the radial-arm maze task is a widely used test in spatial cognition. It allows measuring spatial memory, probably engages a place strategy from rats, and has already been combined successfully with outcome devaluation (Babb and Crystal, 2006a, see Fig. 4, p. 7). Moreover, place cells were shown to fire differently depending on the baited versus non-baited status of arms (Hölscher et al., 2003). However, to our knowledge, no attempt had previously been made to evaluate whether place cells would be modulated by goal or outcome value in a quantitative way. For all these reasons, we chose to rely on this paradigm to elaborate a protocol where the value of a goal could selectively be modified. General methods apply to each of our radial maze

experiments. Namely, the experiments consisted in teaching the eight-arm radial maze to rats, where specific rewards (of different types) were associated to specific arms in the maze. Rats had to learn to collect the rewards without going back to previously visited arms. Working memory (visits to already visited arms) and reference memory (visits to non-baited arms) were recorded. The record of the order of arm visits was also kept. The radial arm maze used is presented in Fig. 69.



Fig. 69: Eight-arm radial maze.

Left: Radial-arm maze used. Each arm is visually similar and holds a possibly baited cup at its end.

Right: Long-Evans rat performing the radial maze task. Visual cues surround the maze and provide spatial information. The aim of the rat is to visit only baited cups and not going back to previously visited arms.

For all experiments, once the rats had reached a performance criterion, we applied a satiety devaluation procedure to one of the two rewards. We then assessed the post-devaluation performance on the maze and compared it to pre-devaluation performance. Post-devaluation test was always performed in extinction, i.e., in the absence of reward, unless otherwise specified. Two parameters were used to evaluate the preference of rats: in the maze, the order with which arms were visited; in their home cage, the amount of food of each type they consumed as well as the order with which rewards were consumed. We expected that rats would first visit the arms associated to the non-devalued food and reject (or lower their consumption of) the devalued food in the home cage test. The specific parameters (type of food used, parameters of the devaluation, configuration of rewards on the maze) are summed up in the Appendix I (p. 200).

Ten naive male Long-Evans rats were used for this series of experiment and the following one. They weighed between 230 and 250 g at their arrival from a commercial supplier (Janvier, Le Genest-St-Isles, France). They were handled for 10 minutes per day and food deprived to 90% of their free-feeding weight before training started.

6.2.1 Experiment number one: working memory

In this first experiment, the task was used in its working memory configuration, i.e., all arms were baited and errors were recorded when rats visited again the previously visited arms. Two different types of food pellets (A and B) were available at specific positions. The reward distribution was constant throughout the experiment but was rotated between rats. We expected that rats would learn to associate a type of reward to a specific arm and that devaluating one type would affect the

order of visits of the arms, namely, that rats would first visit the arms associated to the non-devalued reward. Rats were trained on this task for four trials a day until they reached a criterion of one error or less in four consecutive trials.

Satiety devaluation was performed once the rats had reached the criterion, between the 3rd and the 4th daily trials. Rats were given access to ten grams of one of the reward types for one hour. The test following devaluation was performed in **extinction** so that rats could not update their internal estimation of outcome value using feedback from food consumption; consequently, their behaviour should reflect their internal estimation of outcome value.

After the 4th trial in the maze, a **preference test** was carried out in rat's home cages to determine if the devaluation was effective. Rats were given access to the two types of food and both the order in which they were consumed and the amount of food consumed were recorded.

The results obtained during this preliminary experiment did not show any significant effect of devaluation on the order of arm visits (in terms of the percentage of visits to one or the other rewards for the first four arms visited; data not shown). Moreover, the preference tests carried out in the cage did not show either a significant rejection of the devalued food nor an inversion of preference following the devaluation. Note that the preference was assessed by evaluating which type of food had been entirely consumed first. A second experiment was therefore carried out with modifications on the devaluation protocol and on the distribution of rewards.

6.2.2 Experiment number two: reference memory

The second protocol was similar to the first one but we wanted to test rats' reference memory, in order to be sure that they would pay attention to the distribution of rewards to specific arms. The reference memory version of the task consists in selectively baiting a subset of arms. With training, rats should first visit the arms that are always baited. We used two types of food reward that were distributed in six arms (3 occurrences of each food type). Thus, two of the arms were not baited. The distribution of rewards was the same for a given rat along the experiment but the pattern was rotated between rats. Another modification of the protocol, in this experiment, was that food was provided *ad libitum* during the 1-hour satiety devaluation process. Moreover, the post-devaluation test in the radial maze was performed with the rewards for half of the rats (and in extinction for the other half), because we observed that the absence of reward seemed to disturb rats' behaviour.

Again, the results of this second experiment showed no significant effect of devaluation on the order of visits in the maze. However, results from the subset of rats tested with the rewards tended towards a decrease of the number of arms baited with the devalued food in the four first visits. The devalued food was not completely rejected in the preference test performed after the test in the maze, in rats' home cages. However, rats less readily consumed the devalued food: out of the 60% of rats which preferred the devalued food before the devaluation, only 20% still showed a preference for the same food after the devaluation protocol.

We noted that reference memory errors seemed to reach a plateau after twenty trials (on average, 1 error per trial). This could either mean that rats did not remember the distribution of reward in the maze, or that they did not organise the order of their visits with respect to the baited versus non baited status of arms. We hypothesized that they did not pay enough attention to rewards located at the end of the arms because the cost of an error was too small.

6.2.3 Experiment number three: cost increase

This protocol was similar to the previous one, with the exception of 9-cm high **barriers** that were inserted at the beginning of each arm (see Fig. 70). We postulated that this would increase the cost of arm visits and that rats would ‘pay more attention’ to their choices. Moreover, the criterion set prior to performing devaluation was more selective: no more than one reference and one working memory error had to be done in six consecutive trials.



Fig. 70: Radial arm maze with barriers.

Barriers of equal size and shape were added at the entry of each arm. The picture shows an intermediary setup; in the final one, barriers were all painted in grey so that they could not be visually distinguished.

The insertion of barriers appeared to have the expected effect: rats learned to visit the non-baited arms last and the number of reference memory error steadily declined to 0.4 errors (on average) after 36 tests. In addition, while the devaluation procedure did not seem to trigger the total rejection of the devalued food, it still resulted in a decrease in consumption. Indeed, in the preference test performed in the rats’ home cages, 60% of the animals preferred food B (grain-based food pellets) after the devaluation of food A (‘purified’ food pellets) while all 10 rats preferred food A after the devaluation of B. The initial preference of rats for food A appeared to be in conflict with the devaluation of this food.

Concerning the test performed in the maze, even with a more selective learning criterion, there was still no significant effect of devaluation on the order of rats’ visits.

6.2.4 Experiment number four: limited visits

A new protocol was developed to encourage rats to reflect their preferences in the order they adopt to visit the arms. For half of the daily trials in the radial maze, randomly selected, the rat was removed from the maze when returning in the central stem following its fourth visit. Furthermore,

food types were changed to get a smaller initial preference bias and the distribution of rewards was simplified by spatially regrouping rewards of similar type.

Once more, the results we obtained in this paradigm were not those expected. Satiety devaluation still did not affect the order of visits of rats (either concerning the percentage of arms visited after the first four visits or concerning the rank of the visited arms).

In conclusion, the radial arm maze paradigm did not allow us to observe an effect of satiety devaluation that would be selective for a reward type. Thus, we shifted to a simpler spatial task, where tests could be repeated in order to have a significant behavioural effect that might be less sensitive to the extinction phenomenon. Moreover, since rats' preferences were always directed towards the same type of reward, we performed a thorough evaluation of their preferences to select equally-favoured rewards.

6.2.5 Reward selection

Five different types of reward were used for this experiment. Rewards were home made using condensed sweet milk diluted in water at 10%. Five different flavours were used to create five types of reward. Consumption tests with all five rewards available were performed twice a day for two days. The time spent smelling and eating each reward, together with the quantity of reward consumed during the test, were monitored. These parameters were used to perform choice tests using different pairs of rewards. Finally, a pair was selected using rewards that seemed of equivalent valence for the rats (namely, strawberry and vanilla flavours). A satiety devaluation test was performed on the selected pair, which showed a significant (Wilcoxon signed-rank test, $p = 0.041$) decrease of consumption for the devalued food (performed on $n = 10$ rats). Thus, we selected this pair of rewards for the remaining devaluation experiments.

6.3 Continuous alternation tasks

A second series of experiments was performed using the continuous T-maze paradigm (Wood et al. 2000; see Fig. 71). In this task, rats are released in the central stem of the T-maze and must proceed forward towards the intersection. They can choose between turning left or right. A spatial goal is located at the end of each of the two arms. Reward is delivered at the goal location provided the previously visited goal location was different from the current one. Upon consumption of reward, the rat must go in the side return arm and a new trial is initiated. Rats were prevented from going backwards. In order to get the maximum amount of reward, rats need to continuously alternate between left and right turns (i.e., left or right goal visits). Note that rats have a natural tendency for alternation (e.g., see Baker and Franken, 1968), i.e., exploring new possibilities rather than revisiting previously visited places, which is exploited in the continuous alternation task. In our variant of the task, a given reward type was always associated to a given goal for a given rat and this was counterbalanced between rats. The rewards used were those chosen in the reward selection experiment (Sec. 6.2.5). Three variations of this protocol were used, still with the aim of being able to observe, behaviourally, a change in rats' spatial preferences following devaluation. For the first two,

the rats used were the same as those used in the radial maze experiments. Naive rats were used for the last of these experiments.



Fig. 71: Continuous T-maze.

A rat released in the central stem has the choice to go either left or right at the intersection. Two cups are available at the end of both the left and right arm of the T-maze. The rat must then go through the side return arm in order to begin a new trial from the central stem.

6.3.1 Experiment number five: free alternation and 2:1 rule

For this first protocol in the continuous T-maze, the alternation rule was modified. In order to be able to estimate rats' preferences for a given reward (thus, a given associated goal), we let them choose between left and right visits. Thus, their preferences could be estimated via the number of visits they made to each goal. However, a rule was implemented so that they would remain in a goal-directed state: rats were only rewarded at the second successive visit of the same goal. Two tests were performed per day, each of which lasted twenty minutes. Two training stages were implemented. During the first stage, rats were rewarded regardless of the previously visited goal. In a second stage, they had to execute two laps on the same side of the maze in order to obtain the reward. A subset of the ten rats previously trained in the radial arm maze experiments were trained in this paradigm.

During learning, we noticed that rats that had learned the two-lap rule mainly focused on one side of the apparatus. However, we needed rats not to have any initial preference so that the effect of devaluation would be visible regardless of the location of the goal associated to the devalued reward. For this reason, we discontinued the experiment before the devaluation test and elaborated a new protocol.

6.3.2 Experiment number six: continuous alternation

We designed a new paradigm in the continuous T-maze where we decided to facilitate alternation in rats by preventing them to access an arm when it had previously been visited three times in a row. This was done using the same plastic board used to prevent them from going backwards. In addition, rats were habituated to extinction periods of 5 min for one out of two successive tests. When rats had reached an average performance of at least 1.5 visits per minute on four consecutive tests, we proceeded to devaluation. After devaluation, rats were tested back in the maze, in extinction conditions, and the number of visits they made to each goal was recorded. The rats used for this experiment were different from those trained in the previous alternation task, but still belonged to the ten rats trained in the radial arm maze task.

The results from this experiment did not show any significant effect of the devaluation on the number of visits. However, we did observe a tendency to visit less the goal associated to the devalued reward. Overall, we noted that the devaluation greatly reduced rats' performance, regardless of the position of the goal. This was probably due to a decrease in overall motivation due to satiety combined with the absence of reward on the post-devaluation test.

Furthermore, in this experiment, the preference tests performed in rats' home cages after the task did not evidence any rejection of the devalued reward nor a difference in preference between the devalued and the other reward.

One of the few parameters that we had not yet manipulated in this series of studies was the subjects themselves. Indeed, the rats used as subjects were always the same regardless of the experiment performed. Naïve rats were used for the next (and last) devaluation experiment.

6.3.3 Experiment number seven: continuous alternation

Six naive Long-Evans rats were used for this experiment. They were food-deprived to as to reach 90% of their original weight. The paradigm was similar to the previous experiment (continuous alternation) with two modifications: first, the reward was released only when rats alternate; second, rats were progressively habituated to extinction.

The habituation protocol was used once rats had reached a plateau performance. In practice, this occurred starting from the 14th session of training. From there, extinction periods were introduced for one out of two consecutive sessions, the duration of which was adapted to each rat so that it was the longest possible while rats still performed the task. The mean number of alternations (i.e., left + right laps combined) per 20 min session is presented on Fig. 72 for the duration of the experiment.




Fig. 72: Performance in the alternation task.

The mean number of alternations over six rats is presented (error bars indicate the s.e.m.). The beginning of the extinction habituation protocol is indicated by a red line.

Once a plateau performance was reached, satiety devaluation was performed. Both the test on the continuous T-maze before and after devaluation consisted in a given number of alternations in extinction, then, rewarded trials. The number of extinctions was adapted to each rat and corresponded to the one used during training. Overall performance was drastically reduced following devaluation, but not selectively for one of the goals. Note, however, that rats had been trained to alternate. Thus, we analysed the speed in the central arm, separately for left or right turns, to assess whether speed was selectively decreased when the rat aimed at the devalued goal. This was indeed the case for two out of six rats (data not shown). Interestingly, the home cage preference test subsequently performed demonstrated a selective reduction of preference for the devalued food only for those two same rats.

6.4 Conclusion

In these series of experiments, we aimed at adapting the satiety devaluation procedure to spatial cognition paradigms. The general conclusion that can be drawn is that obtaining reproducible effects of food devaluation on consumption tests was difficult, and observing a selective reduction of preference for the spatial goals associated to a devalued food was impossible. Several factors can explain these difficulties:

- i. An inefficient devaluation procedure;
- ii. No learning of the place – reward type association;
- iii. Lack of selectivity of the devaluation;
- iv. Decrease of motivation caused by the combination of devaluation and extinction test;
- v. High inter-individual variability of devaluation effects on subjects, i.e., need for more subjects to observe a statistically significant effect of devaluation;
- vi. Need for better tuning of devaluation parameters to find the balance between the stability of behaviour (so that rats keep performing in spatial tasks) and its plasticity (so that they adapt their behaviour to reflect the changes in goal value).

Other studies – albeit quite rare - report selective effects of outcome devaluation in spatial tasks. As an example, in the study from Babb and Crystal (2006a; see Fig. 4, p. 7), rats selectively diminished their visits to the arm in a radial maze that was associated with a devalued reward. However, in this study, rats needed to pay attention to the distribution of reward types in the maze, because different types of rewards might not be available depending on the experimental conditions. Moreover, the training protocol was rather intense and complex. This experiment was reproduced by the same authors (Babb and Crystal, 2006b) and others (Naqshbandi et al., 2007). More recently, Alcaraz and colleagues (2014) tested the effects of satiety devaluation in a cross maze. Two different types of reward (grain-based versus sucrose food pellets) were associated to two different arms. Satiety devaluation performed on rats that had been trained in this experiment resulted, on average, on a preference for the arm containing the non-devalued food. We note, however, that on all the above-mentioned experiments, the preference for the devalued food is not that obvious. Effects on performance are generally subtle and such a paradigm might not be appropriate for electrophysiological recordings, in which a small number of rats is used, and in which there is a need for reproducible behaviour from one rat to the other.

Chapter 7 – Goal-directed behaviour in the two-goal navigation task

Aiming at the objectives defined in the foreword (p. 109), we designed a new task adapted from the continuous navigation task (Rossier et al., 2000; Hok et al. 2007). The continuous navigation task has several advantages for studying both goal-directed behaviour and the electrophysiological properties of hippocampal neurons. As a dry version of the water maze task, it is assumed to require a place strategy from the rat. This is confirmed by the fact that rats trained in this task significantly decrease their performance following intermediary and ventral hippocampal lesions (Burton et al., 2009). In contrast to the water maze, it is a continuous task, thereby allowing for multiple sampling of the goal location. It consists of two alternating episodes (goal-oriented navigation and foraging) which allow neural activity to be sampled in the same environment under two different levels of spatial processing. Importantly, it enables to dissociate the goal location from reward consumption. Finally, goal-related activity in the hippocampus was previously observed in this task (Hok et al. 2007a, 2007b, 2013).

Based on our working hypotheses, the rationale to design the new task was twofold. First, to investigate the spatial nature of hippocampal goal-related activity, we introduced two uncued goal zones, so as to evaluate whether neural responses would occur at the two goal locations or they would rather be spatially selective. Second, to study the possible value component of hippocampal goal-related activity, we modified the relative value of two goals by changing the corresponding reward magnitudes, i.e. the amount of food pellets provided whenever each of the two goal locations was activated. Moreover, we allowed rats to freely choose between the two goal zones, such that their behaviour could (indirectly) reflect the value that they had attributed to each goal. A free choice has also other advantages: it should promote an adaptive behaviour, and allow the neural activity at the two goal zones to be compared within-session. This new task was named '**two-goal navigation task**'. For the sake of clarity, we decided to report the behavioural results and the electrophysiological findings obtained with this new task in two separated chapters. The present chapter focuses on the behavioural observations. It first exposes the material and methods we used, and then it presents the results along their interpretation. A discussion focusing on these behavioural results closes the chapter.

7.1 The two-goal navigation task: material and methods

7.1.1 General paradigm of the two-goal navigation task

The two-goal navigation task allows rats to choose between two goal locations in a continuous manner. It can be broken down in four distinct behavioural phases that are continuously repeated:

- i. **Choice** phase: the rat chooses the goal to navigate to.
- ii. **Navigation** phase: the rat moves to the chosen location.

- iii. **Delay** phase: the rat waits at the goal for 2 seconds. If it waits at the proper place, reward is delivered from an overhead dispenser in the form of one food pellet. Importantly, the food pellet bounces at a random place on the floor and stops at another random place.
- iv. **Foraging** phase: the rat searches for the food pellet and consumes it.

These phases alternate in a continuous way (Fig. 73). The protocol described corresponds to the ‘reference’ condition of the task where both goal zones trigger the release of one pellet. The reward magnitude associated to one goal zone can be modified in value-changing conditions, as will be seen later during the explanation of the full experimental protocol (Sec. 7.1.3.3, p.125).




Fig. 73: Behavioural phases of the two-goal continuous navigation task.

The two dashed circles indicate the (unmarked) goal zones. Adapted from Hok et al, 2007b.

7.1.2 Material

7.1.2.1 Subjects

Ten naive male Long-Evans rats, weighting between 230 and 250 g at their arrival from a commercial supplier (Janvier, Le Genest-St-Isles, France), were housed in the laboratory's colony room in compliance with French and European ethical rules, under a 12h/12h light-dark cycle beginning at 7 A.M. and with a temperature of $20 \pm 2^\circ\text{C}$. The procedures were approved by the local ethical committee (authorization no A81212) and the experiments were performed in accordance with European (European Community Council Directive 86/609/EEC) and national guidelines (Council Directive 87848 and permission no 13.24 to E. Save). Rats were weighted and handled each weekday for 10 minutes each until the training began. They were housed in pairs during the training phase. Seven of them were implanted with recording electrodes, after which they were housed individually. Because one rat died following surgery, we implanted an additional rat with electrodes. Rats were

provided with *ad libitum* water and were food-deprived all along the experiment so that their weight remained between 90 and 95 % of their free-feeding weight¹⁰. The daily diet of the rats was standard chow pellets. The food used as a reward during the experiment was 20 mg food pellets (A/P formula, TestDiet®, St Louis, MO, USA).

7.1.2.2 Experimental apparatus

The experiment took place in a 76-cm diameter arena surrounded with black curtains and located at the centre of a square room. It had a grey painted wooden floor and 50-cm high black metal walls. A polarising white cue covering 100° of arc was painted on a portion of the wall (see Fig. 74). The room was dimly lit with the spots oriented in a manner providing indirect and uniform lighting of the arena. The cues used on specific sessions to mark the position of the goals were two 20-cm of diameter gray plastic disks covered with a contrasting linoleum layer and temporarily attached to the floor with ‘blu-tac’. They were only used during training and few recording sessions.

The two computer-defined goal zones, 20 cm of diameter each, were equidistant from the cue card, in a symmetrical manner, so that none of the goals would be easier to locate (see Fig. 75).




Fig. 74: Experimental arena.

Left: Arena from above. **Right:** Arena from above with the goal cues.

¹⁰ This deprivation procedure might seem unusual as animals are generally more deprived so as to reach 80-85% of their free-feeding body weight. We did so in order to maintain behavioural spontaneity and try to guarantee a certain amount of exploratory behaviour (in the exploration/exploitation sense, see p.28; Inglis et al, 2001).




Fig. 75: Constraints for goals positioning.

Segments of same length are indicated by similar symbols.

A pellet dispenser (Med Associates, St Albans, VT, USA), a radio, two cameras and four light spots were attached to a metallic structure two meters above the floor along with a turning commutator. A small plastic device (with one input tube and four output tubes in each cardinal direction) was added to the pellet dispenser's output in order to randomise the pellet landing zone in the environment. The electrophysiological cable was connected to the commutator and its weight was counterbalanced by a pulley and weight system, allowing a rat to freely move when connected to the system. A headstage was attached to the other end of the cable. It had a red light emitting diode (LED) and one operational amplifier (TLC 2272, Texas Instruments) for each channel. When plugged into the rat's microdrive (see p.127 for building microdrive methods), the LED was on the back of the device, above the rat's neck (Fig. 76).




Fig. 76: Headstage connected to a rat.

During the experiment, the radio set was switched on to a FM broadcast station in order to maintain a homogeneous auditory background and mask incidental auditory cues.

7.1.2.1 Task monitoring and data collection

The task was monitored from the adjacent room. Thus, the experimenter was not interfering with the animals' behaviour during the task. Note, however, that two events would trigger an intervention from the experimenter. First, whenever there was rat's urine on the floor – a potential olfactory cue and source of interference with positional tracking – the floor was quickly cleaned with a sponge with the rat remaining in the arena. Note that these events became sparse following a few weeks of

training. Second, if the electrophysiological cable got twisted, the experimenter turned the pulley system to disentangle it. In both cases, precautions were taken so as to enter and exit the curtain zone from a different place every time, in order for the cue card to remain the only polarising cue of the environment.

A diagram of the connections between the cables coming from the experiment room and the apparatus used in the monitoring room is presented in Fig. 77. The video signal from one of the cameras was used by the experimenter to monitor behaviour and pellet consumption. The video signal from the other camera was split up and sent to two computers. One of the computers (Videotrack, Viewpoint, France)¹¹ monitored the rat's position and automatically triggered the pellet dispenser whenever the appropriate conditions were met. The other computer received filtered and amplified (Neuralynx, Bozeman, MT, USA) electrophysiological signals in addition to the video signal. The neural data was digitalized by a dedicated program (Sciworks, Datawave, Loveland, CO, USA). The corresponding parameters are shown in Table 1. The input from one of the channels was used to record Local Field Potentials (LFPs)¹². A communication between the two computers was set up so that the goal-monitoring computer sent event flags to the neural data computer **whenever the rat was detected in a goal zone for at least 2 s**. The latter also received event flags **whenever a food pellet was consumed by the rat**, which were entered by the experimenter.

Throughout the thesis, we will use the term **goal activation** to refer to a goal visit of sufficient duration to release an event-flag (and the activation of the pellet dispenser on rewarded conditions).




Fig. 77: Monitoring room equipment.

¹¹ The Videotrack program was custom made for this experiment by the ViewPoint company as a result of an interaction with one of their engineers.

¹² Note that LFP data was not used in the present study.

	Gain	Low cut filter	High cut filter	Sampling rate
Unit activity channels	10,000 x	600 Hz	6000 Hz	32,000 Hz
LFP channel	1000 x	1 Hz	475 Hz	1024 Hz

Table 1: Signal processing parameters.

7.1.3 Procedures of the experiment

The entire experiment lasted for about nine months, seven of which dedicated to the electrophysiological recordings. The successive phases of the experiment are indicated on Fig. 78.



Fig. 78: Phases of the experiment.

7.1.3.1 Behavioural training

The training phase lasted for approximately 6 weeks and was broken down in progressive steps. First, rats separately learned to activate the goal locations with the help of a 20-cm of diameter cue located at the goal position during the first out of two daily 16-minutes sessions. The two goal positions were first learned separately, alternating from one day to the next. During this step, the delay necessary to activate a given goal was increased from 0 to 2 seconds by steps of 0.5 seconds every time rats reached a performance threshold of 2 visits per minute (in practice, this happened approximately every other day). Then, the two goal locations were presented simultaneously, first with the cues, then without them. Because of the free choice given to the rats to activate either goal, a challenge of this task was to ensure that animals would not focus on one goal only. Some rats indeed expressed a marked preference for one goal during training. To make sure that this preference was not linked to any exterior element (e.g., the position of the door in the experimental chamber), we performed several training sessions where the arena was rotated with respect to the room. This also ensured that rats would rely on the cue card to locate the goal positions. As this was not sufficient to totally remove the preference, a specific balancing protocol was used which involved the eventual extinction of the preferred goal for 8 min periods, meaning that an activation of this goal would not provide any pellet (see Appendix I, p. 200).

7.1.3.2 Surgery and retraining

Once all the rats had reached a performance threshold of 1 visit per goal and per minute with no significant preference between the two goals, they were implanted under general anaesthesia conditions with a bundle of 4 recording **tetrodes** in the **right dorsal hippocampus**. Five rats were implanted above the **CA1** field (antero-posterior: -3.8 mm and medio-lateral: -3 mm with reference

to Bregma, dorso-ventral: -1.5 relative to brain surface). Two other rats were implanted above the **CA3** field with the same A.P. and M.L. coordinates and -2.5 mm dorso-ventral relative to brain surface (Paxinos and Watson, 2005). The microdrive building and surgery procedures will be detailed afterwards (p. 127).

After a recovery period of 1 to 2 weeks, the rats were screened daily for neural activity and the electrodes were lowered (of approximately 22.5 μm) if no useful signals were detected. At the same time, the rats were trained again on the task. The balancing protocol was used again to further consolidate the absence of any marked preference for one of the two goals (see Appendix I, p.200 for the retraining protocol).

7.1.3.3 Electrophysiological recordings protocol

Once implanted and re-trained in the task, rats underwent the same protocol for five or six days a week:

- i. a **screening** session of variable duration ($10 < t < 30$ min);
- ii. a **recording** sequence of sessions whenever the neural signals were estimated to be of interest.

Every morning, all rats' home cages were transported to the monitoring room. Because the experiment usually lasted for the whole day, the order with which rats were tested was changed each day in order to even out the impact of possible circadian variations of motivation on each individual. The rat to be tested was first brought to the experimental room in its home cage. Once the monitoring equipment was ready, the rat was connected to the recording cable and left in the arena from a random position so that the only polarising cue would be the white cue card and not its entry position into the arena. No attempts to disorient rats were made. The task monitoring apparatus was remotely started. Neural signals were screened and positive amplitude detection thresholds were set for each channel while the rat was performing the task in its reference session (explained in the next paragraph). Electrode references were chosen for each channel and kept identical between two consecutive days when we aimed at recording the same neurons.

Whenever neuronal activity considered of interest by the experimenter was detected during the screening session, the rat underwent the following **recording protocol**: 4 consecutive sessions of 16 minutes each where reference sessions alternate with 'value-changing' sessions. During **reference sessions**, both goals were equally rewarded with the release of one pellet. The value changing sessions could be of two types:

- i. **extinction sessions**: one of the two goal zones did not trigger the pellet dispenser anymore. We assumed this would lower the goal value.
- ii. **high value sessions**: one of the two goal zones triggered the successive release of 3 pellets. We assumed this would increase the goal value.

Two specific sequences of sessions were used:

- 1) the **extinction sequence**: alternation of reference and extinction sessions (see Fig. 79).




Fig. 79: Extinction sequence of sessions.

- 2) the **high value sequence**: alternation of reference and high-value sessions (see Fig. 80).




Fig. 80: High-value sequence of sessions.

In all conditions, a re-entry into a goal zone would only be detected when occurring at least five seconds after a previous entry into the same goal zone. Because of this **refractory period**, it was more rewarding for rats to alternate between activations of the two goals. Within a given sequence, the position of the value-changing goal was counterbalanced¹³. The event flags recorded for each goal activation specified the goal position and the reward condition. Importantly, the beginning of a new session was **not indicated** in any way to the rat. After each sequence of sessions, the rat was removed from the environment and the floor was wiped with water to remove (or, at least, mix) the possible olfactory traces.

Whenever putative neural signals were detected, we recorded them on two consecutive days in the two above-mentioned sequences of sessions. A third type of sequence could occasionally be used on a third day, if signals of high amplitude were simultaneously recorded from many cells. In this ‘cued’ sequence, the same cues as the ones used during training marked the two goal positions. Within a cued sequence, reference cued sessions alternated with high-value or extinction cued sessions. Once the signals were recorded for two consecutive days in the extinction and high-value conditions (or three days if we used also the cued protocol), electrodes were lowered by approximately 27 μm to try to find new neurons or get a better signal from the same ones. After a session in which no exploitable signals had been detected, the electrodes were generally lowered by 55 μm . If absolutely no signal of interest was detected, we lowered them by 110 μm instead.

The recording part of the experiment lasted for 7 months, as long as we had (or were expecting) neural signals from the hippocampus to record. The behavioural results presented in this chapter refer to those obtained during electrophysiological recordings. The electrophysiological results (and the specific methods) will be presented in Chapter 8 (p. 147).

¹³ Except once, as an attempt to even out the preference of a rat which always visited the same goal.


It is worth mentioning that the pellet dispenser activation emitted a slight sound that might have been used by the rats as a feedback on their performance. For the extinction condition, there was no such sound. For the high value condition, the three pellets were released sequentially with a 200 ms interval, which resulted in three dispenser sounds. Moreover, a degree of variability in reward contingency was caused by the fact that pellets would sometimes bounce out of the apparatus or get stuck in the pellet dispenser. The overall occurrence of these events was rare (evaluated as <5% of overall dispenser activations). We did not aim at reducing it because it habituated the rats to some uncertainty in pellet release, which we thought would promote perseveration in the extinction sessions (Capaldi, 1957).

7.1.4 Further procedures

7.1.4.1 Microdrive building

The microdrives used for neural activity recording were home made. The manufacturing process of a microdrive can be split up in several steps (Fig. 81):

- i. The electrodes, cut from formvar – insulated nichrome wires, were twisted by groups of four to make tetrodes. Four of these tetrodes were inserted into a 30-gauge stainless steel tubing (17 mm of length for CA1, 18 mm for CA3). The bared end of each wire was then connected to one of the pins of an 18-pins Mill-Max connector. The tubing was connected to one of the remaining pins. We tested the connectivity of the system along the way. The tubing served as the animal ground as well as a guide for the tetrodes.
- ii. Three screws (with a thread of 450 μm) inserted into plastic cuffs were positioned around the connector and the preparation was joined together with acrylic. Once implanted, these screws were used to precisely lower the electrodes into the brain.
- iii. Small connector parts were glued to the sides of the main connector – they were used during the recordings to physically (but not electrically) attach the cable to the microdrive (see Fig. 76).
- iv. The tetrodes were cut so that they exceeded the tubing length by approximately 1 mm. They were subsequently gold-plated so as to lower the impedance of each electrode between 200 – 400 k Ω . This step allowed testing again the conductivity of the system. Microdrives were selected for implantation if they did not include more than four channels with poor conductance (i.e., impedance after gold-plating higher than 600 k Ω).



Building steps (from Jacob, 2014)

Final result
(without the nylon cuffs)

Fig. 81: Microdrive building.

The microdrives could be built in advance but the final steps (cutting and gold-plating) were always performed the day preceding implantation surgery.

7.1.4.2 Surgery details

The aim of the surgery was to implant seven of the rats with the above-mentioned bundle of four tetrodes in the right dorsal hippocampus. First, the rat was deeply anaesthetised with an intraperitoneal injection of a homemade solution of Kétamine, 60 mg/kg (Imalgène 1000, Merial, France) and medetomidine (0.25 mg/kg, Domitor, Janssen, France). Once asleep, the rat was placed in a stereotaxic apparatus in flat skull position. A midline incision of the scalp exposed the skull, which was cleaned with Betadine. We placed five small anchor screws in the skull bone. A hole was drilled at the desired coordinates (with reference to Bregma: -3.8 AP, -3 ML). The dura was carefully removed from the exposed brain surface and the tip of the electrode bundle was positioned just above the brain surface. Then, the electrodes and the surrounding tubing were lowered inside the brain (1.5 DV for five of the rats, 2.5 for the remaining two). From there, sterile soft paraffin was inserted in the hole and around the tubing. The plastic cuffs of the electrode bundle were attached to the skull and the screws using several layers of dental cement. Once the whole system was solidified, the skin was sutured around the electrodes bundle. Finally, the rat was removed from the stereotaxic apparatus and subcutaneously injected with an antibiotic (Amoxicilline, 150 mg/kg) and an analgesic (Buprenorphine, 0.05 mg/kg). After surgery, rats were placed in a recovery room (22°C) for three days before being returned to the colony room. They were provided with *ad libitum* food and water during one week, after which the food deprivation schedule began again (along with the retraining protocol).

7.1.4.3 Histology

The histological study of the brains of the implanted rats was necessary to confirm the localisation of the recording electrodes. Thus, at the end of the experiment, we injected the rats with a lethal dose of dolethal (pentobarbital, i.p.). Once they were deeply anaesthetised, the positions of the electrodes were marked by passing anodal current through one of the wires of each tetrode (15 μ A for 30 s). The rats were then perfused transcardially, first with a saline solution (9%), then with a formalin solution (4%). Their brains were extracted and left in a 30% glucose solution for one or two days. Then, they were frozen with 'carboglace' and stored in a freezer (80°C). The frozen brains were cut coronally with a cryostat and the slices (30 μ m of width) closest to the implantation point were mounted, stained with cresyl violet and stored for further observation. They were examined afterwards under a light microscope to identify the electrodes traces and deduce the putative origin of the recorded cells.

7.1.4.4 Data handling

The pre-processing of recorded data from this experiment was done in two steps:

- i. Neural data, i.e., putative action potentials of undetermined origin, first needed to be **spike-sorted**. This step consisted in allocating each spike either to a putative neuron or to background "noise". It will be detailed in the next chapter (p. 147).
- ii. Once sorted, data files were read by a Matlab program (adapted from a program used in the team, courtesy of Francesca Sargolini) and converted to an open format, containing spike, position and event flag data, in order to be further analysed. All other analyses of behavioural as

well as neural data were done with programs specially made for the thesis using the **Python Programming Language** (version 2.7.2), unless otherwise specified.

7.1.5 Behavioural analysis methods

Only the methods used for behavioural data analysis will be presented here. The analysis methods for neural data collected in this experiment will be seen in the next chapter (Sec. 8.1, p. 147). Two types of data were used for behavioural analyses: rats' trajectories and the event flags recorded along the session.

7.1.5.1 Trajectory correction

The position tracking system was sometimes subject to spurious detections (for example because of reflection caused by the walls of the arena, by rat urine, or because the recording cable was hiding the LED). Position data was processed before being further analysed. The program doing this processing used several steps (inspired from programs used in the team, courtesy of Bruno Poucet and Francesca Sargolini):

- i. Positions were stretched in the x-axis to correct for camera deformation. A factor of 1.2 was used.
- ii. All positions detected outside the arena limits were removed and replaced by linearly interpolated data.
- iii. All positions where rat's speed was higher than a given threshold ($100 \text{ cm}\cdot\text{s}^{-1}$) were also removed and replaced by linearly interpolated data.
- iv. For a few trajectories (15 of them), there were still obviously spurious detections after this correction process. In this case, we manually indicated to the program the area where the bad detection happened. A second run of the program was performed where the positions within this zone were removed (between step ii and iii) and replaced with interpolated positions.

The corrected positions were used to build spatial occupancy maps (and firing activity maps, whose computation is detailed in Sec. 8.1.1.2, p.148).

7.1.5.2 Occupancy maps

A spatial occupancy map corresponds to the time spent by the animal in each part of the environment, clustered in spatial bins. We used square occupancy maps of 32x32 bins. Each bin was a square with sides of approximately 2.5 cm. To compute the spatial map from a raw trajectory, we gathered all position samples corresponding to each bin. The time spent in each bin then corresponded to the number of samples divided by the sampling frequency (50 Hz in our case). An example is shown in Fig. 82.




Fig. 82: From a trajectory to the occupancy map.

Trajectory of a rat recorded during a 16 min reference session in the two-goal navigation task.

7.1.5.3 Behavioural preference index

To evaluate the relative preference of rats for a given goal, we computed an index that we called **behavioural preference index**. It can assess either the value preference or the spatial (side) preference. The value preference index was computed as follows:

$$\text{Value pref} = \frac{\text{No. of highest value goal activations} - \text{No. of lowest value goal activations}}{\text{No. highest value goal activations} + \text{No. lowest value goal activations}}$$

This normalised index could be applied both to extinction or high value condition. Thus, the highest-value goal could be the one providing either 3 pellets (on high value sessions) or 1 pellet (on extinction ones). The lowest value goal could provide either 1 or no pellet (on high value or extinction sessions, respectively). This index ranges between -1 and 1. Negative values indicate a preference for the goal with the lowest value and positive values point out a preference for the most rewarded goal. Small values indicate a well-balanced distribution of goal visits. When computing the value preference index, sessions with the same value modification were grouped regardless of the side of the modification.

When dealing with reference conditions – in which the values of both goals are equivalent – or when addressing the question of spatial bias, we used the **side preference index**:

$$\text{Side pref} = \frac{\text{No. left goal activations} - \text{no. right goal activations}}{\text{No. left goal activations} + \text{no. right goal activations}}$$

High values of this index indicate a marked preference for the left goal, low values indicate a preference for the right one and values around zero reveal a well-balanced behaviour. When computing this index, sessions in which the position of the modified goal differed were kept separated.

7.1.5.4 Statistics

The general protocol for statistical analysis was the following: first, we performed a normality test on the data by testing skew and kurtosis¹⁴. The normality hypothesis was rejected whenever p was smaller than 0.05. In that case, a non-parametric test was used for further analyses. Non-parametric tests were also used whenever the size of the sample was smaller than 20 items. Otherwise, a parametric test was used. The tests used are summarised in Table 2. The given H_0 hypothesis was rejected whenever the p value of the test was lower than 0.05. The testing procedure was the same for behavioural and electrophysiological analyses, unless otherwise specified. The tests were always two-sided.

	One sample (compare to a value v)	Two samples (paired)	Two samples (independent)
Parametric test (normal data)	One-sample t-test <i>stats.ttest_1samp</i>	Student t-test for related samples <i>stats.ttest_rel</i>	Student t-test for independent samples <i>stats.ttest_ind</i>
Non-parametric test (not normal data or $n < 20$)	Wilcoxon signed-rank test comparing sample to a list of zeros <i>stats.wilcoxon</i>	Wilcoxon signed-rank test <i>stats.wilcoxon</i>	Mann-Whitney rank test <i>stats.mannwhitneyu</i>

Table 2: Summary of the statistical tests used.


The function used for each test is indicated in italics (Python Programming Language).

The one sample non-parametric test consisted in creating an ‘ideal’ set of reference values of n samples of zeros (to which we wanted to compare the experimental sample of n values) and applying the Wilcoxon signed-rank test to compare the original sample with this simulated sample.

The legend used in the presented graphs will always be the following: *** for $p < 0.001$, ** for $p < 0.01$, * for $p < 0.05$. Error bars will always indicate standard error of the mean (s.e.m.). Mean and s.e.m for all the graphs presented in this part were computed over sessions.


7.1.5.5 Terminology

The following terminology will be used to refer to specific sessions on the figures:




ref correspond to all first reference sessions. **ref (ext)** corresponds to S0: reference sessions initiating extinction sequences (see Sec. 7.1.3.3, p. 125). **ref (high)** refers to S4, reference sessions used at the beginning of high value sequences. We did not take into account the intermediary reference sessions for the analyses, as they were used as buffer sessions and might present individual variations resulting from the session that was experienced before.

¹⁴ Normality test used: *scipy.stats.normaltest*



ext refers to all sessions in which one of the goal did not provide reward any more (S1 and S3). Depending on the analysis, these sessions could either be combined or separated. In the latter case we use **ext-L** to refer to an extinction session where the left goal was extinguished and **ext-R** when the right goal was extinguished (as in the inset on the left).



high regroups all sessions with one of the goal providing 3 pellets (S5, S7). Similarly, **high-L** is used whenever the left goal was providing 3 pellets and **high-R** otherwise (as shown on the left).

As previously stated, a **goal activation** will always refer to a successful visit of the corresponding goal location, i.e., one which triggered the pellet dispenser activation – for a 1 or 3 pellets goal – or which would have done so – for an extinguished goal. A **trial** corresponds to a unique choice – navigation – delay – foraging sequence.

7.2 Behavioural results

This section presents the behavioural results obtained from the two-goal navigation task. Two main questions will be addressed: whether or not rats were able to learn the task, and whether or not they adapted their behaviour to modifications of the amount of reward provided by the goals.

7.2.1 General performance

7.2.1.1 Sample size

Seven implanted rats underwent the protocol described in the previous section (Sec. 7.1.3.3, p. 125) as long as neural signals of sufficient signal-to-noise ratio were detected during screening. A total of 224 sequences of 4 sessions each were recorded from six rats. The seventh rat did not provide any exploitable neural signals. Only those sequences for which the general spike-sorting quality was deemed sufficient were used for all analyses (presented in Table 3).

	Extinction	High value	Cued	Total
Rat 32	11 / 18	12 / 18	3 / 4	26 / 40
Rat 35	21 / 29	19 / 27	4 / 5	44 / 61
Rat 36	11 / 20	9 / 17	1 / 2	21 / 39
Rat 37	5 / 17	6 / 17	0 / 0	11 / 34
Rat 38	6 / 13	7 / 12	0 / 0	13 / 25
Rat 39	6 / 14	4 / 10	0 / 1	10 / 25
Total	60 / 111	57 / 101	8 / 12	125 / 224

Table 3: Number of sequences of sessions (used / all).

Number of sequences of 4 sessions experienced by each rat, organised by sequence type. The numbers correspond to the used sequences over the total number.

As previously stated (p. 125), the position of the goal whose value was changed was counterbalanced within a given sequence, with the exception of a unique sequence. Thus, out of the 60 extinction

sequences, 59 concerned the left goal and 61 the right goal. Concerning the high value sequences, they included 57 sessions with the left goal set at the high value and 57 sessions with the right goal set at the high value.

7.2.1.2 *General observations*

The first thing we noted was that rats were motivated to perform the two-goal navigation task, despite its duration (a recording session, screening included, lasted around 75 minutes) and the fact that they were not highly food-deprived. Qualitative observations of rats' behaviour during the foraging phase seemed to indicate that they first used auditory cues to orient themselves (straight movement towards the pellet's landing zone), then, once closer, probably used a combination of olfactory and haptic information (performing lateral scanning head movements¹⁵). Rats hardly ever left pellets uneaten between two goal visits. Navigation phases did not evidence any general stereotyped behaviour, apart from one of the rats which displayed a preferred trajectory to reach one of the two goals, but not in an exclusive way. At the goal zone, rearing could sometimes be observed in a subset of the rats. In two randomly selected sessions from the two rats which provided the most neurons (35 and 32), we observed rearing in 10% and 20% of goal visits, respectively. This behaviour seemed to mostly occur following missed attempts at activating a goal. We did not attempt to quantify it further.

7.2.1.3 *Spatial coverage of the environment*

To get a first glimpse into rats' behaviour, we used occupancy maps to assess the average time spent in each location of the arena (see Sec. 7.1.5.2, p. 129 for details). As rats were required to visit (and stay at) the goal location to trigger the release of reward, they were expected to spend more time there than anywhere else in the environment. Fig. 83 shows the cumulated occupancy maps averaged over all sessions and regrouped by session type. They show that rats indeed spent most of their time at the two goal locations. Also, rats adapted their behaviour to goal value modifications by visiting more the goal with the highest value. Finally, the preference seemed more evident in extinction sessions. Similar qualitative results were found during cued sessions (the occupancy maps are presented in Appendix III, p.203).

This descriptive approach does not indicate to what extent rats were successfully activating the goal, which implied waiting at the proper location for the correct amount of time. Thus, a series of analyses were performed relying on the event flags recorded during the experiment. As stated in the methods (Sec. 7.1.2.1, p.122), event flags were recorded automatically each time a goal was activated and manually at the time of reward consumption.

¹⁵ Also entitled « *vacuum cleaner*» behaviour (V. Hok, personal communication).




Fig. 83: Occupancy maps for all sessions (s).

The position of the arena and goals is always the same, with the cue card on the upper part.

7.2.1.4 Task performance

To assess whether rats learned the task rules properly, we used the **activation rate**, which refers to the number of times that rats activated any of the two goals per minute (both goals taken together). As shown in Fig. 84, the rate was significantly higher in the extinction condition compared to the

corresponding reference condition. Conversely, the goal activation rate was significantly lower in the high value condition compared to the corresponding reference.

An explanation for this phenomenon could be that foraging for three pellets takes more time than for one. Thus, if rats activated more the high value goal, they should spend more overall time foraging, hence dedicating less time to goal activations. There was also a significant increase of successful visits in the extinction condition. A similar explanation may hold: as there was no pellet to be retrieved when the extinguished goal was activated, rats could dedicate more time to goal visits. That being said, the average correct response rate, all conditions taken together, was 4.37 visits per minute. Classical performance in the one-goal continuous navigation task is about 2 correct responses per minute (Rossier et al., 2000; Hok et al. 2005, 2007). Thus, in the two-goal navigation task, rats did learn to activate the goal(s) properly.




Fig. 84: Behavioural activation rate.

Mean rate of goal activations per minute, averaged over sessions. Independent t-tests were performed.

Data from **ref (ext)** and **ext** was significantly different ($t = -2.05$, $p = 0.041$).

The values from **ref (high)** and **high** were also significantly different ($t = 8.13$, $p = 8.3 \cdot 10^{-14}$). See methods for further details (Sec. 7.1.5.4, p. 131).

7.2.1.5 Behavioural preference

To evaluate the rats' behavioural preference and a possible bias towards one of the two goal locations, we computed the **side preference index** for each session. The results were regrouped using the above-mentioned categories (see Sec. 7.1.5.5, p. 131) with respect to the side of the changing value goal (Fig. 85). First, data from both reference conditions did not significantly differ from a no-preference condition – although there was a tendency for the ref (high) condition. Second, all value-changing conditions strongly differed from a no-preference condition.




Fig. 85: Behavioural side preference index.

Side preference index averaged over all sessions organised by the corresponding conditions. The values for each condition were compared to a no-preference condition (see methods, p. 131).

Data from reference sessions (normally distributed) did not differ from a 0-mean population (t-tests; $t = -1.53$, $p = 0.13$ for **ref (ext)**; $t = -1.72$; $p = 0.09$ for **ref (high)**).

Data from value-changing conditions (non-normal) all differed from a no-preference sample (Wilcoxon, **ext-L**: $z = 0.0$, $p = 3.45 \cdot 10^{-11}$; **ext-R**: $z = 12$, $p = 2.91 \cdot 10^{-11}$; **high-L**: $z = 486$, $p = 0.002$; **high-R**: $z = 289.5$, $p = 8.86 \cdot 10^{-6}$).

These results show that the preferred ‘side’ (i.e., goal location) during value-changing conditions always matched the goal with the highest value. For example, in ‘ext-L’ sessions (where the left goal value did not provide any more reward), rats had a mean preference index of approximately -0.58, meaning that they directed most of their visits towards the right goal. Thus, changes in goal value had a significant impact on rats’ choices in value changing conditions while rats’ visits were well-balanced when both goals provided the same amount of reward. Visits tended to be slightly biased towards the right goal zone in reference conditions. Importantly, this did not prevent rats from preferably choosing the goal with the highest value in value-changing conditions. Thus, the influence of goal value was significantly stronger than any influence of goal position.

We computed the **value preference index** to assess the impact of goal value independently from the position of the goals (see Fig. 86). In both extinction and high-value sessions, the mean value preference index was significantly different from a no-preference condition.




Fig. 86: Behavioural value preference index.

Value preference index averaged over all value-changing sessions. The sessions from both sides are grouped together (see methods, Sec. 7.1.5.3, p. 130).

In both extinction and high value sessions, rats expressed a significant preference (Wilcoxon, $z = 28$, $p = 2 \cdot 10^{-54}$ and $z = 5.3$, $p = 4.8 \cdot 10^{-7}$, respectively) towards the goal with the highest value.

These results coupled with those addressing the side preference demonstrate that rats predominantly directed their visits towards the most rewarded goal, in extinction sessions as well as

during high-value sessions. As a consequence, **the goal value had a major impact on rats' choices while the spatial position of goals did not affect their behavioural preference.**

We did not aim at comparing the two types of value-changing sessions with each other. However, the preference pattern seemed much different between the extinction and high-value condition. Namely, the preference for the goal with the highest value seems more important in extinction sessions. We interpret this difference by a 'repulsive' effect from the extinguished goal whereas, in the high-value condition, the goal with the lowest value continued to provide one pellet. This could either maintain a certain amount of exploration (i.e., visits to both goals) or simply slow down learning. Also, assessing the contingency between the three released pellets and the last visited goal might be less straightforward than learning that a goal provides no longer a reward.

7.2.1.6 Reward consumption

Does adaptable behaviour translate into a higher profit at the end of the day? Considering that foraging for three pellets takes more time than foraging for one, we checked that rats' strategy – going more to the goal with the highest value – provided them with more reward. To check this, we computed the mean number of pellets eaten for each condition¹⁶ (see Fig. 87). During both the extinction and high value sessions, the number of pellets eaten differed from that of the corresponding reference condition. Thus, the strategies expressed by the rats were rewarding: they manage to get the best out of the high value sessions, and not to lose too much reward on the extinction sessions.




Fig. 87: Reward consumption.

Average count of reward events over all sessions of a given condition.

Mann-Whitney rank test for independent comparisons show a significant difference between the count for **ref (ext)** and **ext** ($U = 2137$, $p = 9 * 10^{-6}$) and also **between ref (high)** and **high** conditions ($U = 398$, $p = 0.95 * 10^{-21}$; see statistics, p.131, for further details).

Overall, this first series of results shows that rats successfully learned the two-goal navigation task. They could differentiate the two goal zones and appeared to assess a specific value to each, since

¹⁶ Note that, as this information was entered manually by the experimenter, a certain amount of pellets might have been missed. The measured numbers are likely to be underestimations of the actual ones.

their choices were always directed towards the goal with the highest value. Importantly, the position of the goal does not alter the choices of the rats.

7.2.2 Time course of behavioural preference

We carried out another series of analyses to evaluate the temporal evolution of the behavioural preference throughout a given session.


7.2.2.1 *Individual example of one session*

Fig. 88 illustrates an individual example of the temporal dynamics of one rat's behaviour. Two sequences of four sessions recorded on two consecutive days are presented. The side preference rate was computed per 1-min bins and plotted along with markers for each goal activation (left or right) and reward consumption events. One can note the well-balanced preference in all reference sessions, with a return to no preference in the intermediary reference sessions. This rat also adapted to value-changing conditions and appeared to treat extinction and high value conditions equally. Indeed, in both conditions it oriented its preference towards the most rewarded goal while still keeping a certain amount of 'exploratory' visits to the less rewarded goal. The behavioural preference towards the high-value goal took a few minutes to be established.

Of course, this is just an example and, overall, rats showed a variety of behaviours. Some of them seemed to lean towards a more exploitative behaviour (always visiting the same goal), whereas others were more explorative (alternating between goals). In rare occasions, exploitative animals could simply not notice at all the change in goal value as they just never visited the other goal but this happened rarely and even less frequently with time. Note that we rarely noticed individual examples of 'a-ha' or 'insight' moments, where a rat shifted its preference all at once (Durstewitz et al., 2010).

Fig. 88: Time course of choices along sessions – individual example.

Two sequences of 4 sessions each are presented, performed on two consecutive days. The first sequence is an extinction one. The second is a high value sequence. Four elements are presented for each session: *i)* the occurrence of **goal event flags** is presented as a function of time: dark blue dots for each activation of a goal that triggered the release of 1 pellet, red for extinguished goals, green for high-value (3 pellets) goals. The side of the activated goal is also represented in the y-position of the dot: +1 for left goal activations, -1 for right goal activations. *ii)* The **reward consumption** events are represented as crosses in the middle ($y=0$). *iii)* For each bin of time (1 min), the **side preference rate** is computed and plotted in cyan dots connected with a cyan line. *iv)* a least squared **polynomial fitting** of order 3 was performed on the preference rate data and is plotted as a dashed cyan line. This polynomial function was used as an attempt to detect inflection points in the preference, with no exploitable results.



To obtain a representative picture of the evolution of behavioural preference under the three conditions, we averaged the time course of value preference rate over all sessions (using 1-min bins, similarly to the individual example presented above). The resulting graph, presented in Fig. 89, allows three elements to be highlighted. First, the preference in the reference sessions remained stable across the duration of a session. Second, the setting up of the preference for the most rewarded goal occurred much faster under the extinction condition than in the high value condition. Third, the plateau performance was higher in the extinction condition¹⁷.




Fig. 89: Within-session time course of value preference index.

Each dot represents the value preference index (for extinction and high value conditions) or the side preference index (for the reference condition) computed on a 1-min time bin and averaged over all sessions of this condition. Data from ref (ext) and ref (high) are combined. The coloured overlays indicate the s.e.m. area.

To evaluate the inter-individual variability in preference, we computed the within-session time course of preference rate separately for each rat. The results are presented in Fig. 90. A general overview indicates that the above-mentioned time course of preference was reproducible among rats. However, clear inter-individual differences are evidenced. For example, one rat (32) displayed a strong bias for the right goal in the reference condition. However, this bias did not prevent its preference rate from evolving towards the goal with the highest value, in both extinction and high value conditions, even more efficiently than some of the other rats. The spatial bias of rat 32 may explain the tendency at the population level for a right goal preference (seen on Fig. 85, p.136). Notably, none of the other rats showed such a marked spatial bias in the reference condition. Another rat (39) seemed to have difficulties in differentiating the high-value goal from the other one, since its preference in the two different conditions tended to overlap. Nonetheless, this rat could differentiate the extinction condition from the reference condition.

Overall, the individual data highlighted that the two-goal navigation task leaves room for the expression of inter-individual differences. However, the above-mentioned general pattern of time course of preference (Fig. 89) was preserved at the individual level: the preference towards the most

¹⁷ Concerning the plateau performance, we cannot exclude that asymptotic performance would not reach a similar value had sessions lasted longer. However, this seems unlikely, given the specificities of each condition. The extinguished goal would always keep some kind of repulsive value (with a negative cost/ benefit ratio) compared to the one-pellet rewarded goal in the high value session.

rewarded goal was quickly set up for the extinction condition while taking longer to be established under high value conditions.




Fig. 90: Individual time course of preference along sessions.

The numbers indicate the numbers of sessions used in the analysis for each rat.

Red: ext; Green: high; Blue: ref

In order to further characterise the within-session time course of the behavioural preference, we computed the mean value preference index for time bins of eight minutes, thereby dividing each session in **two phases of equal duration** (Fig. 91). For the reference condition, there was no significant difference from one phase to the next. For both types of value-changing conditions, the

preference for the most rewarded goal significantly increased between the early and late phases. We concluded that splitting up the sessions in two parts of equal duration was sufficient to observe a significant within-session change in behavioural preference.




Fig. 91: Time course of preference using 8 min time bins.

The side preference index is shown for reference sessions and the value preference index for the value-changing sessions. Paired comparisons of the preference rate were done between the two parts of sessions.

There was no significant evolution of preference in the reference condition (Wilcoxon, $z = 3215$, $p = 0.99$).

On the contrary, the preference significantly increased with time for the extinction and high value sessions ($z = 322$, $p = 5 \cdot 10^{-18}$ and $z = 216$, $p = 3 \cdot 10^{-17}$, respectively).

In summary, this series of results on the temporal dynamics of preference demonstrate that rats effectively adapted their within session behaviour. They gradually directed their goal choices towards the most rewarded one during either extinction or high-value sessions, whereas their preference remained stable in reference sessions. Although rats could express inter-individual differences in their spatial preference towards one or the other goal, they consistently adapted their choices to modulations of goal value.

7.3 Discussion: behaviour in the two-goal navigation task

7.3.1 Summary of results

The aim of this experiment, from a behavioural point of view, was to design a task that would address both flexible spatial memory and goal-directed behaviour. Moreover, we wanted to assess whether or not rats were able to dissociate two different uncued goal zones, and whether or not they could succeed in associating a specific value to each goal location. Finally, we needed a behavioural measure of their estimate of goal value. To this end, we developed a continuous navigation task inspired from previous paradigms (Rossier et al., 2000; Hok et al., 2007) to which we applied two major modifications. In the original task, rats must locate a hidden goal and wait for two seconds at the proper location to 'activate' the goal, which triggers the release of reward. First, we implemented the possibility of choosing between two simultaneously available goal locations. Second, on specific sessions, we modulated the amount of reward released upon activation of one of the two available goal locations.

The results obtained from the two-goal navigation task can be summarised in the following way. First, rats demonstrated their knowledge of the basic rules of the task, which consisted in waiting a proper amount of time at the goal location. Also, all rats were able to memorise the two goal locations, as shown by the high frequency of goal activations per minute and the well-balanced distribution of their goal choices in the reference condition. In addition to spatial memory, the decision-making aspect of the task was also mastered by the rats, which readily adapted to modulations of goal value by orienting their choices towards the most rewarding goal. Interestingly, the preferences were not unilateral: a certain amount of ‘exploratory’ choices were often done, apparently more in the high-value condition than in the extinction condition. These results will now be discussed in relation to other experiments and to the concepts that were overviewed in Part I of this thesis.

7.3.2 Is the two-goal navigation task hippocampal-dependent?

A crucial question, since we are interested in the role of the hippocampus in goal-directed behaviour, is whether or not this brain structure plays a role when a rat is engaged in the two-goal navigation task. One could actually argue that the one-goal version of the task could be solved without relying on a hippocampal-dependent place strategy. Indeed, Burton and collaborators (2009) showed that rats with intermediary to ventral hippocampal lesions were still able to locate the goal zone precisely enough to trigger the release of reward, although their performance was significantly lower than controls. Moreover, Rossier and collaborators (2000) showed that rats could navigate towards the uncued goal zone in darkness, albeit less efficiently than in light. When the amount of available information about the environment is reduced, or when the hippocampus is not functional, it is possible that other strategies take control of behaviour. Indeed, we saw that the information necessary to the use of different strategies coexist in the brain (see Sec. 3.2.7, p. 51). In the one-goal navigation task, a path integration strategy could be used to shuttle between goal and reward locations. However, in the two above-mentioned studies, performance was much lower in hippocampal rats or in darkness compared to controls or light conditions, respectively. This result is likely to indicate that the optimal way to solve the task may involve the hippocampus and that a purely idiothetic-based strategy is not optimal. Consequently, we believe that the place navigation task relies on the hippocampus for most rats. It must be noted that a subset of rats, from either Hok and collaborators’ (2007) study or ours, occasionally expressed stereotyped trajectories. Interestingly, in our case, this was mostly the case for the one rat that showed a large bias for the right goal (rat 32, see Fig. 90, p.141). The spatial preference and the stereotyped strategies expressed by this rat seemed both to decrease as the rat became more familiar with the variations of reward magnitude (data not shown). This suggests two things. First, the changes in goal value might have promoted exploration of the two goal locations. This is in line with the studies about novelty-triggered exploratory behaviour, even though the two concepts of exploration are not exactly similar (see 3.2.1, p. 36). Second, the choice component of the task might have prevented the use of a path integration strategy. Indeed, we are not aware of studies in which rats could simultaneously keep track of two goal locations using path integration. Overall, we suggest that the introduction of a choice between two goals of changing value in the continuous navigation task is a supplementary reason for using flexible navigation and flexible decision-making. Because the hippocampus was

shown to be involved in the place strategy and in behavioural flexibility (see Sec. 3.2.6.3, p. 46), we believe that optimal performance in the two-goal navigation task relies on a properly functioning hippocampus.

7.3.3 Did rats demonstrate goal-directed behaviour?

In instrumental learning tasks, goal-directed behaviour is assessed by measuring the effects of outcome devaluation on performance (Balleine and Dickinson, 1991; see Sec. 2.2, p. 24). In these conditions, tests are usually performed in the absence of reward so that the animal can only rely on its internal representation of the outcome to drive its choices. Moreover, devaluation studies modify the ‘internal’ value of a specific outcome. In the present task, we did not manipulate the intrinsic reward value but rather the reward magnitude. Moreover, this manipulation relied on the short-term association of value to a goal. Thus, it is difficult to compare this task with instrumental learning tasks. That being said, we observed a clear and significant adaptation of rats to the new goal value in the course of a session. Because the goal was not directly associated to the reward, we believe that rats’ behaviour demonstrates that they were able to associate a value to each of the two goals and adapt their choices accordingly. Thus, most probably, they were expressing goal-directed behaviour.

Another argument for assessing flexible behaviour supported by cognitive maps is the demonstration of ‘insight’ by a subject (Tolman, 1948). Insight is reflected by radical transitions from low performance to high performance (Tolman and Honzik, 1930; Durstewitz et al., 2010)¹⁸. This type of task solving dynamics is also observable in humans (Bowden et al., 2005). Gallistel and collaborators (2004) highlighted that the incremental increase in performance often observed during learning could be a methodological artefact, due to averaging over subjects or even trials. In most paradigms, a sudden shift in performance (perhaps accompanied by a sudden insight about the solution) could actually be observed if one relied on the learning curves for individual subjects. Although this could have been the case in the present task, we seldom observed such ‘insight’ moments in rats’ behaviour, nor did we observe sudden transition in rats’ choices in individual data (see for example Fig. 88, p.139). However, the question requires further analysis, for example, searching for inflection points in the individual time course of preference. In any case, even though rats in this task probably did not show insight about the new goal value, we believe that modifying the amount of reward provided by the goals throughout the task allowed us to maintain the rats in an outcome-attentive state, which prevented them from switching from goal-directed to habitual behaviour with overtraining (see 3.2.7, p.51).

Overall, the results from the two-goal navigation task as well as our previous experiments with food devaluation (see Chapter 6, p.111) converge towards the idea that rats lean towards the strategy which requires the least amount of spatial processing. This is probably why food devaluation did not have any marked effect in the previous studies (because rats did not pay attention to the reward – place association) and this is why a specific training procedure, as well as manipulation of reward magnitude, was necessary to maintain rats in conditions of flexible behaviour.

¹⁸ One could even suggest that the ‘all-or-none’ nature of these insights or ‘a-ha’ moments could result from an attractor-like functioning of underlying neural networks.

7.3.4 Comparison with existing paradigms

A number of recently developed tasks similar to the one used in this work can be found in the literature. The multiple-T-maze task, for example, allows rats to continuously choose between two options (turning left or right, Johnson and Redish, 2007, see Fig. 64, p.97). However, the spatial processing demand of this task is quite limited. Similarly, Lee, Ghim and collaborators (2012), and others (e.g., Kim et al., 2009; Sul et al., 2011), used a two-armed bandit task adapted to a continuous T-maze where the two reward sites are associated with different probabilities of reward release. Hillman and Bilkey (2010) also used a continuous alternation maze where the cost-benefit ratio of the two options was manipulated by the introduction of barriers and modulations of the amount of reward (see also McHugh et al., 2008, in which delay, cost and reward magnitude were manipulated). In all of these tasks, rats were given the possibility for free choices. Nevertheless, the spatial component was reduced to a left-right choice, which is not so different from a lever-pressing task. On the other hand, spatial tasks that require flexible behaviour exist, starting from the Morris water maze (Morris, 1981), but also encompassing all types of mazes and paradigms that require a place strategy (3.2.6, p. 43; see also Sec. 5.3, p.94). Most of these spatial tasks usually rely on the memory for a unique goal location. A few paradigms involved the localisation of several goal locations (e.g., Kobayashi et al., 2003; Dupret et al., 2010) but these paradigms usually result in rats performing stereotyped trajectories, which might influence place cell's firing. Recently, Mc Kenzie and collaborators (2013) tested the ability of rats to remember several goal locations in a circular track. They demonstrated that rats were able to remember the position of at least 4 different goals, one of which was briefly cued (with light) at the beginning of trials. Overall, the two-goal navigation task is one out of few tasks that combine the localisation of a hidden goal in an open field environment with a free choice between two locations in a continuous manner. We believe that this task can bridge the gap between spatial cognition and decision-making studies in a more ethological (albeit, less controlled) manner than instrumental learning tasks. Thus, it probably addresses a more 'natural' functioning of the rodent' brain.

7.3.5 Finer analysis of natural decision-making

The two-goal navigation task is adapted to the study of neural correlates of decision-making through the manipulation of several of its parameters, such as the delay at the goal zone, the magnitude of reward, the reward contingency, the amount of spatial processing required (e.g., by changing the goal position or size), and finally the spatial 'cost', i.e. the distance between the current location and the goals. The last of these parameters is inevitably manipulated because, due to the random position of food pellets, the starting point of each navigation phase rarely happens to be at the same distance from the two goals. All these parameters are likely to play a role in the final selection of the goal. Using a theoretical model of behaviour to fit the actual choices expressed by rats could shed light on the relative importance of all of these parameters. Also, extending this model to account for electrophysiological data recorded during the task could help towards a better understanding of the underlying neural correlates substrates (similarly to the approach by Lee, Ghim and collaborators (2012)).

The two-goal navigation task could be used to discriminate several ‘behavioural profiles’, or ‘personality traits’ of individual rats (Pawlak et al., 2008). Individual levels of **exploration/exploitation ratio** could be part of these behavioural profiles (see Sec. 2.5, p. 29; Daw et al., 2006, Tamosiunaite et al., 2008). In the context of the two-goal task, this ratio could correspond to the percentage of choices when the rat alternated between goals over the percentage of choices when the same goal was chosen for two successive trials. Such a ratio would be most meaningful during value-changing sessions. Indeed, observations of rats’ behaviour during the task seemed to indicate that they might be categorized in two populations according to their tendency to ‘explore’ the less rewarded goal. Interestingly, rats that ‘exploited’ more were more likely to express stereotypical trajectories; this could underlie possible relationships between individual exploration/exploitation ratio and preferred strategies. The memory systems (e.g., implicit or explicit) preferably used by a rat could also be taken into account by the definition of these behavioural policy profiles. One might even propose that these profiles could be, to a certain extent, applied to human behaviour. Then, understanding the potential differences in the neural bases of these ‘personality traits’ could be of use in the neuropsychiatry domain (e.g., for models of anxiety or depression, Mällo et al., 2008, or schizophrenia, Floresco et al., 2009; Río et al., 2014) but also in neuroeconomics or sociology (see for example Daw et al., 2006; Franck et al., 2009; Jepma et al., 2010; Redish, 2013; Berger-Tal et al., 2014). This is maybe most developed in the large field of research that focuses on the neural bases of addiction, where researchers use instrumental learning to address the failures of the decision-making system (see Everitt and Robbins, 2005). Yet, the use of more ‘natural’ tasks (addressing natural abilities such as spatial behaviour) that allow behaviour to be precisely analysed, combined with recordings at the cellular level, could be a complementary approach towards a better understanding of the neural bases of personality traits in healthy individuals.

7.4 Conclusion

The two-goal navigation task appeared to have several advantages. First, rats were motivated and solved the task with good performance. It is a continuous task that allows data from a large number of trials to be collected throughout a session. It provides an accurate way to quantify rats’ preferences for a spatial goal and, indirectly, the value they assign to each goal. Furthermore, their preference appears to be modified on a short timescale, following changes in the magnitude of the reward associated to a goal.

This novel task combines the spatial demands of a navigation task and the adaptive behaviour demand inherent to a goal-directed choice task. Because flexible goal localisation and decision-making is required, we hypothesise that the hippocampus plays a role when animals are solving this task. Analysing the neural data collected during the two-goal navigation task should then help answering some of the questions raised in the introduction (see Objectives, p. 109); in particular, if the hippocampus is implied in the coding of goal value, we should find evidence for such a coding in the two-goal navigation task.

Chapter 8 – Hippocampal goal-related activity in the two-goal navigation task

This chapter is devoted to the results concerning electrophysiological data gathered in the two-goal navigation task. The experimental methods common to electrophysiology and behaviour were exposed in the previous chapter (7.1, p. 119). The processing and analysis methods specific to electrophysiological data will be exposed here, followed by the results which are separated in two sections. A global discussion of the overall electrophysiological results obtained in the two-goal navigation task will close this chapter.

8.1 Methods: preprocessing and statistical analysis of electrophysiological data

8.1.1 Pre-processing methods

8.1.1.1 *Spike-sorting*

Electrophysiological data files recorded along the experiment by the Sciworks computer were fed into a spike-sorting program: Offline Sorter™ (V2.8.8, Plexon). Clusters of spikes were delineated manually, relying mainly on three parameters: peak amplitude, peak-to-valley amplitude and amplitude at a chosen time ('slice'). The individual shape of each spike was also used to guide clustering. The distributions of inter-spike intervals (ISIs) were monitored to evaluate the quality of sorting: spikes with an ISI smaller than the refractory period of a neuron of 2 ms are considered to be either noise or interference from another unit (see, for example, Gerstner and Kistler, 2002). Thus, only clusters with less than 1% of spikes in this interval were selected for further analysis. An example of 4 putative pyramidal units discriminated on the same tetrode is shown in Fig. 92. The spike-sorting step was manually done at the end of the experiment. Whenever possible, similar clusters were applied to successive recordings so that the cluster corresponding to a given neuron would be the same across all the sessions in which that same neuron was detected.




Fig. 92: Four discriminated units in the spike sorting software.

The waveforms of the discriminated units, for each channel (electrode) of the tetrode, are presented in the upper left hand corner (coloured lines) along with the unclassified data (in gray). Below this are shown the mean waveform and the distribution of ISIs for each discriminated unit. Each spike's amplitude on the fourth channel as a function of its amplitude on the first one is represented on the right.

8.1.1.2 Firing rate maps

Firing rate maps were computed in three steps:

- i. Each spike was associated to the position of the rat closest in time to the time of spiking.
- ii. Spikes of a given unit were gathered in space bins using their computed position. We used the same bins as those used for occupancy maps (see Sec. 7.1.5.2, p.129), i.e., squared bins of 2.5 cm side each constituting a 32x32 grid. This resulted in a spike count map.
- iii. The spike count map was then divided by the occupancy map to obtain a **rate map**, which provides the mean number of spikes emitted per second at each location.

A sample rate map from a reference session of a cell recorded in CA1 is presented in Fig. 93. The same colour code will be used for all maps presented in the thesis: white pixels indicate unvisited locations, dark blue pixels indicate no firing and other colours represent different (discrete) levels of firing rate, from baseline to maximum, as indicated on the corresponding colour scale. All categories (and their associated colours) were covering the same range of firing rate, corresponding to one sixth

of the maximum rate. For example, in the firing map presented in Fig. 93 limits of the colour bins are always incremented of 5.13 Hz, i.e. 1/6 of 30.79 Hz. The firing rate maps were not smoothed.

The orientation of firing rate maps will always correspond to the view with the cue card on the top of the arena. Note that the goals are not horizontally aligned. This is due to the position of the arena with respect to the camera. In any case, the goals were equally distant from the cue card, as defined in the methods (see Fig. 75, p. 122).




Fig. 93: Firing rate map from a CA1 place cell in a reference session.
Left: Arena and goal positions. **Right:** Corresponding firing rate map.

8.1.1.3 Unit categorisation


Putative pyramidal neurons and interneurons were visually discriminated on the basis of their waveform parameters, their firing patterns (bursty or not), and their firing rate. Within the putative pyramidal neuron group, **place cells** and **silent cells** were differentiated by the experimenter according to their overall firing rate as well as the presence or absence of a visible place field in the arena. Fig. 94 presents a sample cell from each category, with its waveform and corresponding firing rate maps for a given sequence of sessions. Cells whose classification was ambiguous were discarded from the analysis. The total number of recorded cells of each type is presented in the results part (Table 4, p. 158).

Fig. 94: Waveforms and corresponding rate maps.

Left: waveforms of the signals recorded on each of the 4 channels of a tetrode during a 16-min session, categorised as emitted by the same source during spike sorting. In this example, the two middle channels (for cell A) and the first channel (for B and C) were noisy and not used during spike sorting.

Right: rate maps of the corresponding cell in 4 consecutive sessions. The waveforms shown correspond to the first (reference) session.

Cell A was categorized as a place cell, cell B as a silent cell and cell C as an interneuron. Cell A and B were simultaneously recorded. Note that, on these maps, the last bin of the colour scale was rounded down.



8.1.1.4 Pooling of similar units

Spike-sorting allows spikes to be clustered according to their putative origin (hopefully, a neuron). However, monitoring a given neuron in consecutive sessions might not be so obvious and even more challenging across days – due, for example, to the drifting of electrodes. Also, considering two neurons recorded in different sessions to be different whereas they actually are the same can lead to spurious influence of individual neurons on the final results. Thus, we used several parameters to label each cluster of spikes, that is to say, associate it with a neuron identifier. The classification was done by the experimenter using two types of information: first, waveform and cluster parameters (position, shape) from the spike sorting software, then, firing rate maps: if we suspected two clusters from different sessions to correspond to the same neuron given the appearance of their firing map, they were given the same label. For example, moving the electrodes forward could lead to different cluster properties but the spatial map would still look the same. When in doubt, we tried to be conservative and group cells together. Once units were grouped, they were considered to be the same neuron for all the analyses. Thus, whenever a given parameter (e.g., the firing rate) was computed, if several samples of a cell were available, the average over all of the samples within a given group of sessions was computed to obtain a **mean firing per session** for this neuron.

8.1.1.5 Place field detection and centroids

Numerous methods are employed in the literature to detect place fields. We defined a place field by at least nine contiguous (side-wise, but not corner-wise) pixels with associated firing larger than the grand mean rate, i.e., the total number of spikes divided by the total time. The grand mean rate was computed for each reference session (S0 and S4 sessions) and averaged over all sessions, across all pyramidal neurons (i.e., putative place cells and putative silent cells). An example of place field detection is presented in Fig. 95. Note that what seems to be goal-related activity (medium blue bins at the goal locations) was not detected by this method, certainly because of its very low firing rate and sparse activity.




Fig. 95: Place field detection.

Top: Rate maps corresponding to the four sessions of an extinction sequence.

Bottom: Detected field for each of the above sessions. The extent of the detected field is shown in red.

To estimate the distribution of place fields across the arena, we identified the **centroid** of each place field by computing the barycentre of the bins belonging to a given field weighted by the corresponding firing rate. The spatial distribution of centroids in the reference condition was used to determine whether or not there was a bias towards goal locations. Note that only the larger place field was used when multiple place fields were detected. Whenever several samples of the same neuron in the reference condition were available, one of them was randomly chosen.

8.1.2 Goal-related activity

Goal-related activity, in our work and previous studies (Hok et al., 2007a, 2007b, 2013), refers to **out-of-field firing emitted at the goal location during the 2s delay period** (see Fig. 73, p. 120, for description of the delay period). The detection of goal-related activity was done in the temporal domain using goal activation event flags. Goal-related activity can only be analysed if it is not ‘contaminated’ by the place field activity. For the rest of the analysis, only cells whose **place field did not cover, not even partially, any of the two goal zones** will be used. This classification relied on a visual inspection of the firing maps. To retrieve goal-related activity, all spikes emitted within a given time window around each goal activation event were gathered. For the analyses that concerned the goal-related activity, the time window was 2s before the event. For the peri-event time histogram (PETH) analyses, it ranged from -4 s to +2 s with respect to the goal event flag.

8.1.3 Peri-event time histograms

To evaluate the **temporal dynamics** of goal-related activity, PETHs time-locked to goal activation were computed. Spike times were aligned with the event flag of interest and bins of 100 ms each were used to build the histograms, i.e., the spike count per time bin. Then, each histogram was normalised by the spike density (i.e., the mean spike count per trial) so that we could average PETHs from different cells. PETHs were used to compare the temporal evolution of activity, whereas they did not allow the firing rates to be compared across different conditions.

8.1.4 Analysis of the reward consumption-related activity

We also analysed out-of-field neural activity related to reward consumption. As previously mentioned, food pellets could land anywhere in the experimental arena, possibly within the cell’s place field. To specifically evaluate a possible reward coding by out-of-field activity, our analysis program only selected reward events occurring when the rat had not been inside the cell’s place field(s) for the 2s time window preceding the event. A visual example of such a selection is presented on Fig. 96. The firing activity corresponding to the selected reward events was used to compute out-of-field reward-related PETHs.




Fig. 96: In-field reward events removing process.

Before/after representation of reward events (red dots) with the 2s long preceding trajectory (grey lines), for a 16-min session corresponding to the rate map presented on the left. Note how the trajectories (and the corresponding events) passing through the field (seen on the left trajectory inset) were removed after processing (right inset). The largest of the black circles represents the approximate limits of the arena and the smaller ones stand for the position of the goals.

8.1.5 Goal-related activity significance

The presence of goal-related activity on the firing rate maps of hippocampal could be due to the fact that the rats spent a significant amount of time at the goal (see behavioural occupancy maps, Fig. 83, p. 134). Note, however, that rate maps are normalised by time. Nonetheless, an important part of our work relied on determining analytically whether goal-related activity was higher than out-of-field activity. To do so, we evaluated the significance of goal-related firing for each goal location taken separately. We compared the mean number of spikes emitted during the delay preceding goal activation with the number of spikes emitted during the **same number of 2 s episodes randomly collected** out of the two goal zones and out of the place field(s) of the neuron. Whenever a sufficient amount of such episodes could not be gathered, we performed a random draw with replacement among existing episodes to get the proper number. Then, a two-sided Mann-Whitney U test was performed to compare the two series of spike count. Whenever the H_0 hypothesis was rejected at the $p < 0.05$ level and the number of spikes at the goal was higher than out-of-field collected spikes, the cell was considered to fire significantly for this specific goal location. This test allowed for evaluating the proportions of cells that expressed a **goal-related activity significantly higher than out-of-field firing**.

To compare proportions of cells significantly firing for one, both or none of the goal zones between different populations of cells, we used **Fisher's exact test** (performed in the Matlab© environment). A significant difference was assessed whenever the p statistic was smaller than 0.05. We used an exact test instead of a chi-square test (for example) because in some conditions the numbers of cells in a given category could be quite low (see results, p. 159).

8.1.6 Analysis of the value-related hippocampal activity

8.1.6.1 Firing preference index

One of the main questions of our work was whether the goal-related activity would underlie a coding of value, i.e., whether or not it would be modulated in a consistent way (positively or negatively) by a change in goal value. The task was designed so as to allow for within-session value preference, since one of the goals always provided the same amount of reward compared with the previous session (unchanged value) while the value of the other goal was increased or decreased (changing value). A coding of goal value could be absolute (i.e., only the goal with the changing value would trigger differential firing) or relative (both the goal with modified value and the other goal would result in a change in firing rate). For both scenarios, we reasoned that a coding of goal value would be visible on a value preference rate index, inspired from the one used for assessing behavioural preference (see p. 130). This firing preference index was computed as follows:

$$\text{Value pref} = \frac{\text{Firing rate at the highest value goal} - \text{Firing rate at the lowest value goal}}{\text{Firing rate at the highest value goal} + \text{Firing rate at the lowest value goal}}$$

This index ranges from -1 to 1. Negative values indicate a stronger firing at the goal zone with the **lowest value**. Conversely, positive values indicate a stronger firing at the goal with the **highest value**. Values (whether positive or negative) around zero indicate an absence of preference of firing for a given value. The value preference was computed for each neuron in a given condition. Whenever a given neuron had been recorded several times in the same session, the final value for this ‘neuron-condition’ pair was obtained by averaging the preference rate from over all its occurrences. In the end, we obtained a normalised value representing the firing preference for goal value for each neuron and in each condition, which allowed to get a general picture of the preference of the population of cells. Note that data were pooled with respect to goal value, not goal side: value preference index from the extinction condition, for example, was computed from ‘ext-L’ as well as well as ‘ext-R’ sessions (see terminology, Sec. 7.1.5.5, p. 131). Therefore, the value preference index removes the spatial aspect of the goals.

For the reference condition, which allowed us to assess the spatial bias of goal related activity, we used the side preference rate, computed as follows:

$$\text{Side pref} = \frac{\text{Goal firing rate at the left goal location} - \text{Goal firing rate at the right goal location}}{\text{Goal firing rate at the left goal location} + \text{Goal firing rate at the right goal location}}$$

The firing side preference index is positive if the cell preferentially fires for the **left goal** and negative if it fires more for the **right goal**.

Two types of statistical tests were performed on firing preference rate data. Either we wanted to assess whether or not the sample did not differ from a no preference population (i.e., sample of 0 values), or we wanted to compare (with paired tests) two samples obtained during the two parts of a

given session. The tests were the same as those used for behavioural preference and the same conditions (test for normality) were applied (see Sec. 7.1.5.4, p. 131).

8.1.6.2 *Cells included in the analyses*

Only a subset of the recorded cells was used for the analyses on the firing preference index, based on the following criteria:

- i. Since we wished to compare the activity at the two goal zones, we selected only the cells whose field did **not overlap** any of the two goal zones (as was the case for all analyses that concerned goal-related activity).
- ii. Since we were interested in the encoding of goal value by goal-related activity, we only took into account the subset of cells which exhibited a **significant goal-related signal** on at least one goal zone. Some of the analyses were also performed only on cells that expressed significant goal-related activity on the two goal zones.
- iii. Whenever the behaviour of a rat in a given session was strongly biased toward one goal location, there was a risk of spuriously biasing the firing preference rate in favour of the most visited zone. Indeed, at the single neuron level, the probability of firing appeared quite low for one trial (data not shown, but see Hok et al., 2007a). Moreover, rats that only experienced the outcome of a goal once or twice in the course of a session, probably did not properly update their estimation of goal value. Hence, we decided to only use data from sessions in which rats had visited the two goals a minimum amount of times. Setting such a **threshold on the number of goal activations** is bound to be somewhat arbitrary. According to behavioural data, the mean number of goal visits was the lowest for the high-value condition (3.8 visits/min, see Sec. 7.2.1.4, p. 134). This number corresponds to approximately 60 activations over a 16-min session. Would the visits be evenly distributed between the two goals, this would yield 30 activations of each goal for a session. We chose to take one third of this number as a threshold, i.e. to only take into account recorded sessions in which the rat visited **at least 10 times each goal**. For the analyses on subparts of sessions (i.e., first 8 min and second 8 min), we split this number in two and only took into account sessions where **at least 5 visits** were made to each goal in both parts of the session. We verified that the sessions complying with this threshold evidenced behavioural preference patterns similar to the ones observed without this threshold (both for entire sessions and half-sessions; data not shown).

8.1.6.3 *Box plots*

Box plots were used for presenting a subset of the results. In those cases, the box extends from the lower (Q1) to the upper (Q3) quartile values of the data with a red line indicating the median of the sample (Q2). The ‘whiskers’ (black horizontal lines) extend to the most extreme data points that are not considered as outliers: the lower one is placed at the minimum value above $Q1 - 1.5 \cdot IQR$, and the upper one is placed at the maximum value below $Q3 + 1.5 \cdot IQR$, where IQR stands for inner quartile range.

8.2 Results: hippocampal goal-related activity in reference sessions

From the behavioural results exposed in Chapter 7 (p. 119), we saw that the two goal navigation task could be well-adapted to answering to the questions we raised previously (see Objectives, p.109). Indeed, rats performed well in the task, they were able to discriminate two spatial goals, and to assign a value to each of them. In the present section, we present the analyses performed on electrophysiological data gathered during the **reference condition**. In particular, we addressed the issue of overrepresentation of goal locations, the proportion of pyramidal cells in the hippocampus that expressed goal-related activity, whether this goal-related activity was expressed by CA3 as well as CA1, whether silent cells play a role in the goal firing, and, finally, to what extent this goal-related activity contains a spatial component.

8.2.1 Histological results

In order to verify the anatomical origin of the recorded cells, we relied on light microscope observations of brain slices compared to stereotaxic atlas data (Paxinos and Watson, 2005). We also used a graphic representation of the number of cells recorded as a function of electrode depth (presented in Appendix IV, p. 204). This analysis confirmed that all rats had been implanted at the proper AP and ML coordinates, with the exception of one of them for which the implantation site was more anterior than expected (around -2.7 AP posterior to bregma, probably due to difficulties in locating the bregma on this rat). Hence, neurons from this rat were actually recorded in the lateral extremity of the hippocampus and were classified as CA3 cells. Note, however, that we cannot be entirely sure that some of these cells were not from the CA2 field. For the other rats implanted above CA1 (-1.5 mm DV with respect to dura), we managed to record spatially selective signals from two distinct layers of cells which we considered to be CA1 and CA3. In one of the two rats implanted above CA3 (-2.5 DV), there were also two distinct layers of cells. We considered that the first layer was CA1 and that we had actually implanted more dorsally than expected. It must be noted that we cannot be absolutely sure that some of the cells classified as CA3 were not dentate gyrus cells.

8.2.2 Number and types of cells

In total, we recorded 194 putative neural cells from the dorsal hippocampus of rats performing the task. 103 of them were classified as place cells, 41 as silent cells and 15 as interneurons. The remaining cells were discarded due to doubts on the quality of the spike sorting. A few examples of units categorised as **place cells**, recorded in the reference condition, are shown in Fig. 97. Note that each of these cells expresses a visible, localised place field, of variable size and shape depending on the cell. The firing rate maps of a subset of the recorded **silent cells** are shown in Fig. 98. They were qualified as 'silent' because they did not display any place field in the arena and they had an overall low firing rate (note the difference of scale between place and silent rate maps). The orientation of maps, as stated in the methods (see Fig. 93, p. 149) is similar for all cells, with the cue card on top.




Fig. 97: Place cells recorded in the reference session.




Fig. 98: Silent cells recorded in the reference session.

Table 4 presents the overall sample of recorded cell types and numbers and their distribution among rats. Only a subpopulation was used in most analyses, as mentioned in the methods (see pp.152 and 155). The number of cells used is always indicated in the corresponding results. Interestingly, silent cells could be found in CA1 as well as in CA3.

	Place Cells		Silent Cells		Interneurons		Unknown	
	CA1	CA3	CA1	CA3	CA1	CA3	CA1	CA3
Rat 32	-	31	-	6	-	3	-	15
Rat 35	35	7	22	2	1	3	8	1
Rat 36	10	1	5	-	3	1	6	-
Rat 37	3	2	1	-	2	-	1	-
Rat 38	-	6	-	4	-	2	-	2
Rat 39	4	4	-	1	-	-	1	1
Total	52	51	28	13	6	9	16	19

Table 4: Number and types of recorded units.

The **mean firing rate** (\pm standard deviation) for the whole population of putative pyramidal cells was 0.69 (\pm 0.07) spikes per second. This value was used for the detection of place fields (see methods, p.151). The mean firing rate for the place cell population was 0.92 (\pm 0.097) and 0.09 (\pm 0.008) for the silent cell population. The mean firing rate of interneurons was much higher: 10.52 (\pm 8.34) spikes per second.

8.2.3 Influence of goal location on place fields

To evaluate whether or not the population of place cells encoded the goal location via the localisation of their main place field, we computed the spatial distribution of the place field centroids in the arena (Fig. 99). Centroids from both reference conditions (ref (ext) or ref (high)) were indifferently taken into account so that each place cell was represented by one centroid (see methods, Sec. 8.1.1.5, p.151, for the definition of centroids). Starting from this global representation, we dissociated CA1 from CA3 centroids to assess whether or not they would distribute themselves differently in the environment. There was **no significant difference** from a uniform distribution of the centroids for any of the three populations of cells. There was no difference between CA1 and CA3 distribution of fields either (see Table 5).




Fig. 99: Distribution of place cell centroids.

Each blue dot represents the position of a place cell (in case of multiple place fields, only the larger one was considered). The two small black circles indicate goal positions. The separations in six sectors of equal size were used to count the number of centroids in each portion of the environment.

The number of centroids from two opposite portions was summed up to get three categories (a goal and two non-goals). We used a **chi-square** test to assess whether the distribution of centroids in these categories was different from a uniform distribution and an **exact Fisher** test to compare proportions between CA1 and CA3. The results are presented below (Table 5).

	All	CA1	CA3
2 goal sectors (2+5)	34	15	19
2 non goal sectors (1+4)	21	8	13
2 non goal sectors (3+6)	28	15	13
Chisquare test for uniform distribution (k= chisquare statistic)	k = 3.1, p = 0.22 n.s.	k = 2.6, p = 0.27 n.s.	k = 1.6, p = 0.45 n.s.
Fisher comparison between CA1 and CA3 proportions		p = 0.08 n.s.	

Table 5: Tests for homogeneity of the place cell representation.

See Fig. 99, above, for methods used to assess homogeneity of the distribution of centroids.

Thus, even though rats were fully trained in the task, we evidenced a **homogeneous representation of the environment**. Interestingly, we noted a tendency for the CA1 and CA3 representations to be different ($p = 0.08$). Visual inspection revealed that CA3 fields seemed to be more central and closer to goal locations, whereas CA1 fields appeared to be more peripheral.

8.2.4 Goal-related coding by the hippocampal population

We assessed the statistical significance of the goal firing for each cell whose field was not overlapping any of the two goal zones, and for each goal zone. Namely, we investigated whether the goal firing rate was significantly higher than the out-of-field firing rate (see methods, p. 153). For the silent cells,

this consisted in comparing the goal-related activity to the activity expressed anywhere else in the environment. The results are presented in Table 6 and Fig. 100. A large majority of pyramidal cells expressed a significant goal-related activity: 57/79, namely, **72.2 % of the cells fired significantly more at at least one of the two goals than anywhere else outside their place field**. Importantly, a majority of **CA3 place cells** also expressed goal-related activity (68.4 %). Moreover, **silent cells**, of CA1 and CA3 origins pooled together, also showed a significant goal-related activity in the same proportions (69.4 %). Comparing the proportions of cells coding for one, both, or none of the goal zones between the CA1 and CA3 pyramidal cells yielded a non significant difference (Fisher's exact test: $p = 0.29$). Comparing the goal coding proportions of place versus silent cells (pooling CA1 and CA3 cells for each of these populations) provided a similar result, i.e., there was no significant difference between the two populations (Fisher's exact test: $p = 0.7$).

	Left goal	Right goal	Both goals	Total
CA1 place cells	3	7	9	19/24
CA3 place cell	4	6	3	13/19
CA1 silent cells	9	3	6	18/26
CA3 silent cells	2	4	1	7/10
Total	18	20	19	

Table 6: Categories of hippocampal cells with significant goal firing.

The number of cells expressing a significant out of field goal-related firing at either only left, only right, or both goals are presented.

The total number of goal encoding cells of each category over the total number of analysed cells (i.e., cells whose field does not overlap in any of the goals) is summarised on the left. Data are graphically presented in Fig. 100.




Fig. 100: Goal coding proportions in the hippocampal pyramidal population.

This graph presents the proportions of cells significantly coding for one or the two goals (data from Table 6). Place cells and silent cells categories regroup indifferently cells from CA1 or CA3. Conversely, the graphs representing CA1 and CA3 proportions regroup indifferently place and silent cells.

These results suggest that the **whole hippocampal population of pyramidal cells** appeared to participate in goal coding. Interestingly, the majority of the cells participating in this coding selectively encoded one or the other of the two goals. Hence, we can conclude that the goal-related coding held a **spatial aspect**. None of the two goals seemed to be more represented than the other, either by the main fields of place cells, or by the significance of goal-related activity (18 of the cells selectively 'represented' the left goal while 20 of them represented the right goal). From now on, the

analyses will only concern cells with a significant goal-related activity (either for one of the two goals or for both). Notably, we tested whether a characteristic temporal profile of goal-related activity could be observed in this task and in the different populations of neurons recorded.

8.2.5 Temporal profile of goal-related activity and reward consumption-related activity

In order to characterise the time course of hippocampal goal-related activity, PETHs were computed **for the whole population of cells that expressed a significant goal-related activity** (see p. 152 for the corresponding methods). To assess whether a particular time-course of out-of-field firing could be related to reward consumption, we also computed the PETH time-locked to the reward consumption event (by only considering those events for which the rat was not in the place field of the cell during the 2s preceding the consumption of the pellet, see p. 152 for methods). The resulting PETHs are presented in Fig. 101. The detail of the population activity is also displayed below the corresponding histograms, to provide an overview of the individual behaviour of the whole population of cells. In these graphs, each line indicates the mean normalised firing rate for each neuron as a function of time using the same timescale as PETH plots.

For both left and right goal zones, the hippocampal goal-related activity had a specific temporal profile, with a marked increase of firing around 500 ms after entry into the goal zone, and a sharp decrease at the moment of activation of the pellet dispenser. The raster plots of the hippocampal population activity, along with the rather low variability of goal-related firing (as shown by the PETHs), suggest that this **characteristic temporal profile** was shared by the majority of the cells used for this analysis.

By contrast, this temporal profile did not apply for the out-of-field firing preceding reward consumption. Also, the PETHs combining indifferently in-field and out-field firing time-locked to reward consumption did not provide any evidence for a population coding of the reward consumption event (data not shown). Overall, this suggests that the out-of-field firing of the hippocampal pyramidal population did not encode reward consumption nor the anticipation of reward consumption.




Fig. 101: Temporal profile of goal-related and reward-related activity.

Top: Averaged, normalized peri-event time histograms time-locked either to left goal event, right goal event, or reward consumption event ($t=0$; 100 ms time bins). The blue overlay corresponds to the s.e.m.

Bottom: Corresponding population activity. The colour indicates normalized firing rate. Each horizontal line corresponds to the firing rate as a function of time for a given neuron (100 ms bins as above). CA1, CA3, place and silent cells are included in these plots.

To investigate whether this temporal profile was shared by all cells expressing goal-related firing, the individual PETHs for CA1 and CA3 place and silent cells were computed. They are presented in Fig. 102. The increase in firing at around 500 ms after entering the goal zone (either left or right) is visible on all plots. The temporal profile seems less pronounced on the CA3 plots but one must note the difference of scale between CA1 and CA3 – due to a single neuron in CA3 that fired strongly upon leaving the left goal zone (data not shown).

In conclusion, regardless which population of cells expressed goal-related activity (CA1, CA3, place or silent cells), the **specific temporal profile** of this firing – a slightly delayed increase following entry into the goal zone, and a sharp decrease at the moment of pellet dispenser activation – is preserved without any major difference.




Fig. 102: Separated PETH for subclasses of cells.

The peri-event time histograms are computed as previously (see methods, Sec. 8.1.3, p. 152) for subpopulations of cells that coded significantly at least for one of the two goal locations.

Finally, although this was not central to our study, we also computed the PETHs from cells recorded in **cued sessions**, presented in Fig. 103. Interestingly, they evidence that the temporal profile of goal-related activity is present, but not forward-shifted as in previous reports (see Fig. 67, p.101; Hok et al., 2007a, 2007b, 2013). This discrepancy might be due to the fact that, in our paradigm, the rats did not have much experience with the cues apart from the early learning phase and the few cued recording sessions. Thus, even when cues were present, the rats in our protocol could have still relied on a place strategy and not a cue-guided strategy. Conversely, the rats providing results in the cued condition from Hok and collaborators' series of studies had only experienced the cued version of the task. In any case, our results suggest that the shift of the temporal profile seen in previous studies is not due to the mere presence of cues at the goal locations but rather to the experience rats had with the cues and probably to whether or not they relied on cues to navigate.




Fig. 103: PETHs of goal-related activity in cued sessions.

Normalized firing rate averaged over neurons ($n = 11$), before, during and after the delay at the cued goal zone. Eleven pyramidal cells from CA1 and CA3 were taken into account, recorded in the **reference condition**, including ref (ext) and ref (high). Zero indicates the time of pellet dispenser activation. 100 ms bins were used. The blue overlay indicates the s.e.m.

8.2.6 Conclusion

The series of results obtained from the reference sessions in the two-goal navigation task provides answers about the nature of the hippocampal goal-related signal. First, we demonstrate a homogeneous distribution of place fields in the two-goal navigation task, similarly to results obtained in the one-goal version of the task (Hok et al, 2007a). There was **no over-representation** of the goal zones by place fields. Second, we show the existence of **goal-related activity in CA1 place cells** in this task, with a temporal profile similar to previous reports in the one-goal navigation task (Hok et al, 2007a, 2007b). We also bring new insights into the expression of goal-related activity by the hippocampal population. The two-goal design of the task allowed us to shed light on a possible **spatial selectivity** of goal-related activity: most ‘goal coding’ cells significantly fired for one of the two goals. Furthermore, we demonstrated that this goal-related activity was expressed by **CA3 cells** and **silent cells**, with similar properties regardless of the population; namely, all categories of cells demonstrated a **temporally organised** goal-related firing and a similar spatial distribution of goal-related activity. Altogether, these results bring strong support to the spatial hypothesis of the goal-related signal. Indeed, the fact that CA3 place and silent cells expressed goal-related activity suggests that this signal might not arise from a direct subcortical input, since CA3 receives fewer projections than CA1 from motivational structures such as the amygdala or the VTA. This issue is to be discussed later (Sec. 8.4.1.1, p. 174). Moreover, the existence of a preferred spatial goal in the goal-related activity of a large proportion of goal coding cells suggests that this signal incorporates a spatial component. Indeed, if it were purely motivational, there would be no reason why neurons would not fire for both goals, provided that rats’ behaviour is the same in the two locations (to be discussed in Sec. 8.4.2, p. 178). To go deeper into the understanding of the goal-related activity of pyramidal cells, further analyses were performed, addressing, in particular, to what extent the goal-related firing was modulated by goal value.

8.3 Results: is hippocampal goal-related activity modulated by goal value?

Here, we focus on the analyses concerning value-changing conditions, but also on data from reference conditions that can help interpreting the results. The main question is whether goal-related activity of hippocampal neurons is modulated by changes in the estimated value that rats associate to a given goal. To address this question we took into account only pyramidal cells with significant goal-related firing at least one of the two goal zones, and sessions where rats visited at least 10 times each goal zone (see Sec. 8.1.6.2, p. 155).

Table 7 summarises the number of cells obtained in the different experimental conditions. Note that the data from cue sequences were not taken into account. The number of cells specifically used for each analysis is indicated in the corresponding graphs.

	Place Cells with main field not overlapping on goals		Silent Cells		TOTAL
	CA1	CA3	CA1	CA3	
extinction sequences	24	16	21	6	67
high-value sequences	17	16	18	11	62
cue sequences	8	4	2	3	17

Table 7: Distribution of recorded pyramidal cells according to the protocol sequence.

8.3.1 Temporal profile of goal-related activity under value-changing conditions

To assess whether changes in goal value modulated the temporal profile of goal-related activity, we computed goal-related PETHs for the value-changing conditions (see methods, Sec. 8.1.3, p. 152). The results are presented in Fig. 104. The **overall temporal pattern of discharge was preserved** for the goals with changing values as well as for the goals that still provided one food pellet. The sharp decrease of firing at the end of the delay was less pronounced for the extinction condition, a phenomenon that might be linked to the absence of auditory feedback or to the fact that rats perhaps stayed longer in the goal zone because no reward was delivered. However, the previous study by Hok et al. (2007) in the one-goal navigation task evidenced that a short period of extinction (4 min) did not prevent rats from leaving the goal zone after the 2s delay. In our task, the extinction lasted for 16 minutes, which might explain the discrepancy in the results. To test for this, we computed the PETHs in the extinction condition by splitting up each session in four parts of equal duration (i.e., data from each of the 4 min parts were used to build PETHs). The results are presented in Fig. 105. Interestingly, the PETHs for the low value goal only expressed a ‘normal’ profile during the first 0 to 8 minutes. Afterwards, the decrease at the end of the delay seemed to fade. In any case, these results would require further research to be fully interpreted, notably a thorough analysis of speed profiles.

8.3 – Results: is hippocampal goal-related activity modulated by goal value?




Fig. 104: Temporal profile for value-changing conditions.

PETHs were computed for each goal in each value condition, averaged over all cells. The delay period is indicated by dotted lines. Time (abscissa) is in seconds with respect to pellet dispenser activation. Ordinates indicate normalized firing rate (see methods, p.152). The same scale is used for all plots. The coloured overlay indicates the s.e.m. Note that for extinction sessions, an event flag was recorded 2s after entry into the goal zone but the pellet dispenser was not activated.

The **top row** concerns PETHs from extinction sessions. The **bottom row** concerns high-value sessions. The four PETHs on the **left** concern the goal with the lowest value in each condition (no pellet for the extinction condition, 1 pellet for high-value condition). Conversely, the four **rightmost** PETHs concern the goals with the highest value (either 1 or 3).

Sample size (number of neurons) from left to right is as follows: 41, 40, 40, 41 (top row); 36, 41, 41, 36 (bottom row).




Fig. 105: Time course of the temporal profile for 4-min bins during the extinction condition.

Each inset corresponds to the PETH computed during 4 minutes.

Top row: PETHs for the unchanged value (goal activation triggers one pellet; sample size, from left to right: 50, 50, 50, 50)

Bottom row: PETHs for the lowered value (goal in extinction; sample size, from left to right: 49, 47, 43, 36)

The same scale is used for all plots. Note that the scarcity of goal-related activity in the last parts of the session is probably due to the low number of visits made to the extinguished goal.

8.3.2 Goal-firing preference as a function of goal value

To assess more quantitatively, on the firing rate level, whether goal value would modulate goal-related activity, we computed the **value preference index**, which indicates the preference of firing within a given session for each cell. The **side preference index** for reference sessions is also indicated to assess the spatial bias of goal-related activity. Fig. 106 presents the results for the whole pyramidal population in the form of box plots (see methods, Sec. 8.1.6.3, p. 155). We tested whether the firing preference rates were not statistically different from a zero mean population (which indicates no preference – see statistics, p.131). Statistical tests resulted in an absence of difference, i.e., in none of the conditions did the cell population fire more for one or the other goal. For the reference conditions, this meant that there was no bias in firing neither for the left nor the right goal position. For the value-changing conditions, this revealed an absence of bias for a specific goal value.

Interestingly, the variability of firing preference seems to be different between the reference conditions and the value-changing conditions. This difference in firing preference variability may indicate that the spatial component has a stronger impact on goal-related activity than the value aspect. Indeed, we saw in the previous section (see Fig. 100) that the majority of goal coding cells expressed a significant goal-related activity for either one or the other of the goals. The presence of extreme values of side preference in the reference condition probably reflects cells with a strong spatial bias of goal-related activity. To test this, we computed the preference index for the subset of cells which expressed significant goal-related activity at the two goal zones. The results are presented in Fig. 107. With this subset of cells, the variability of firing preference in the reference session is much more reduced. This is in line with the fact that the variability of the whole population is mainly due to spatial aspects of goal-related coding.

Furthermore, the results from value-changing sessions of the 2-goal coding population still do not evidence any specific modulation of the firing preference by goal value. Thus, even for cells with a significant goal-related activity at the two goal locations, the firing rate at the goal is independent from goal value.

We also tested whether or not subpopulations of cells differently reacted in the face of value modifications: CA1 and CA3 pyramidal cells, and place or silent cells from either field, were tested separately. A subset of the results is shown in Fig. 108. The entire results can be found in Appendix V, p. 205. None of these populations expressed a value-related firing. The CA1 population was significantly ($p < 0.05$) different from a no-preference condition but **only for the side preference index** (see Fig. 108). Interestingly, although the CA3 population was not detected as expressing a particular preference, the data from the reference sessions seemed bimodal (not tested; see Fig. 108). This apparent bimodality completely disappeared when the value preference index was concerned. Note that the quite small number of cells in this condition ($n=7$) precludes any hasty conclusion.

Overall, this first series of result is in line with an absence of influence of value on the goal-related activity of the whole pyramidal cell population. Goal-related activity rather seems to be modulated by spatial aspects, i.e., the position of the goals.




Fig. 106: Firing preference index for the goal coding hippocampal population.

The **side preference index** is presented for the two reference conditions and the **value preference index** for the value-changing conditions. Box plots are used to visualise the variability of the population (see methods, Sec. 8.1.6.3, p. 155). The red line indicates the median.

T-tests did not evidence any difference from a 0-mean population (from left to right: $t = 1.95$, $p = 0.058$; $t = -0.54$, $p = 0.58$; $t = -0.27$, $p = 0.78$; $t = 0.23$, $p = 0.82$).




Fig. 107: Firing preference index for the two-goals coding hippocampal population.

Results from the subset of cells with significant goal-related activity at the two goal locations are presented with the same parameters as Fig. 106.

No difference from a population of 0 values was evidenced (Wilcoxon tests, from left to right: $Z = 26$, $p = 0.17$; $Z = 49$, $p = 0.35$; $Z = 71$, $p = 0.79$; $Z = 68$, $p = 0.68$).




Fig. 108: Firing preference of CA1 and CA3 pyramidal cells.

The **side preference index** for the **ref (ext)** session, and the **value preference index** for the **ext** session, are presented for the CA1 and the CA3 hippocampal population (place and silent cells included). Only the data from CA1 cells in the reference session was significantly different from a no-preference population. The light grey dots represent individual (neuron) data.

CA1: ref (ext): t-test ($t = 2.09$, $p = 0.045$); ext: Wilcoxon ($Z = 201$, $p = 0.35$)

CA3: ref (ext): Wilcoxon ($Z = 11.4$, $p = 0.52$); ext: Wilcoxon ($Z = 11$, $p = 0.61$)

8.3.3 Activity of hippocampal interneurons at goal locations

Interneurons could play a role in goal-related activity. Indeed, a decrease of their inhibitory firing has been shown to be concurrent with the appearance of place fields in a new environment (Wilson and McNaughton, 1993; see section 5.2.2, p. 89). There might even be a tight coupling between interneuron firing and the spatially-selective firing of place cells (Hangya et al., 2010). Previous work in the one-goal navigation task did not evidence any consistent drop in the activity of interneurons at the goal location (Hok et al., 2007a). In the present work, visual analysis of the firing maps of interneurons often – but not always – evidenced a decrease of firing at the goal zone (see for example the interneuron presented in Fig. 94, p. 150). To gain more insight into this issue, we computed the PETHs of interneurons time-locked to the pellet dispenser activation and to reward consumption (similar to pyramidal cells PETHs, see Sec. 8.1.3, p. 152 for methods). The PETHs for one of the two goal zones and for reward consumption are presented in Fig. 109. Interneurons do not evidence any increase of firing during the 2 s delay at the goal. Actually, the firing during the delay period is decreased compared to the baseline and progressively comes back to baseline after goal activation. Thus, interneuron firing is modulated during the goal delay period. Concerning activity time-locked to reward consumption, no increase of interneuron firing is visible but there is a brief drop at the moment of the event flag, i.e., when the rat retrieved and ate the pellet.




Fig. 109: PETHs of interneurons at the goal and before reward consumption.

Left: PETH for the interneuron population ($n = 13$) time-locked to pellet dispenser activation in the reference condition combining ref (ext) and ref (high). The results are for the right goal; the ones from left goal were highly similar (data not shown). blue overlay indicating s.e.m.

Right: PETH time-locked to reward consumption in the same conditions.

All recorded interneurons were indifferently taken into account for this analysis.

Overall, these results probably mirror the decrease in speed related to the entry into the goal zone and to the reward consumption. Indeed, interneurons were shown to be correlated with rat's speed (McNaughton et al., 1983). As an approximation of rats' behaviour at the goal zone, one can look at the speed profile in the goal zone from Hok et al., 2007a (see Fig. 112, p. 181). In this study, the decrease in speed at the goal zone mirrors relatively well the decrease in interneuron firing seen in Fig. 109, albeit the interneurons appear to precede rats in this decrease. An analysis of speed in the present study would be required to confirm the speed hypothesis. Nevertheless, we observe that the decrease in firing is rather homogeneous while in the goal zone, contrary to the temporal profile of goal-related activity (see Fig. 101, p. 162). Thus, a release of interneuron inhibition could eventually participate in the goal-related firing of pyramidal cells (perhaps upon entry in the goal zone) but

other mechanisms are most probably involved. Moreover, the reward consumption PETH evidence a decrease in interneuron firing, whereas no effect of reward consumption was seen on the activity of pyramidal cells (see Fig. 101, p. 162).

To go a bit further into this issue, we computed the firing preference index of interneuron firing, for all sessions. The results are presented in Fig. 110. The side preference rate of the interneuron population was never different from a zero-mean population, indicating no preference of firing. Moreover, we note the much reduced variability compared to the place-related variability of the goal-firing of pyramidal cells. This means that even at the single-cell level, interneurons do not discriminate between the two goal zones, whereas pyramidal cells do. Finally, similarly to place and silent cells, interneurons do not appear to code for goal value.




Fig. 110: Firing preference index of interneurons.

Data are presented in a similar fashion than Fig. 106, p.169.

None of the conditions presents a significant difference from a 0-mean population (results of the Wilcoxon signed-rank test, from left to right: $Z = 43$, $p = 0.86$; $Z = 29$, $p = 0.24$; $Z = 36$, $p = 0.81$; $Z = 36$, $p = 0.81$.)

Note the much reduced variability compared to pyramidal cell firing preference, regardless of the smaller size of the sample.

To conclude, the firing of interneurons seems to be modulated during the delay period. However, the time course of this firing does not appear to be correlated to that of the goal-related activity. Moreover, the variability of the side preference index is quite low, indicating that most interneurons do not express a firing preference for one or the other goal. This is in contrast with the goal-related firing of place and silent cells, which was highly spatially-biased. Overall, this suggests that interneurons alone cannot be the source of the spatially selective goal-related firing of hippocampal pyramidal cells. Moreover, these results corroborate that information on goal value does not appear to be expressed in the dorsal hippocampus during the delay period.

8.3.4 Within-session evolution of firing preference

Our behavioural data showed that the preference for the highest goal was larger during the second part of sessions, which indicated that, by that time, the rats had learned the new goal value (see Fig. 91, p. 142). The fact that intra-session dynamics are important for the encoding of goal value was also suggested by the evolution of PETHs in the extinction session (see Fig. 105, p. 167). Indeed, a differential coding of value could happen at the beginning of sessions (where it could be used to encode the new value) or at the end (where it could indicate the expected value).

8.3 – Results: is hippocampal goal-related activity modulated by goal value?

Thus, we computed the value preference index for 8-minutes subparts of sessions. This was computed first for all cells available for the analysis, i.e., the cells that expressed a significant goal-related firing on at least one of the two goal zones, in sessions in which the rat visited at least 5 times each goal. Then, the analysis was performed under the same conditions for those cells with significant goal firing at the two goal zones. The results are presented in Fig. 111. They demonstrate that there was **no significant within-session shift of value preference of the goal-related activity**, regardless of the condition (extinction or high-value) and the group of cell used. Note that there was no difference between the two parts of the reference sessions either (data not shown). Interestingly, the variability of the results was quite large for the whole population and reduced when only the “two-goal” cells were taken into account. This could be due to the fact that, because the “one-goal” cells probably fire very few spikes in their non-preferred goal, our index tends to take extreme values, above all when the trials sample is reduced as is the case in the 8-min subdivisions of sessions. Note for example several of the cells that all shift to a value preference index of exactly one from one part to the next. Yet, this does not bias the results towards a difference between the two subdivisions.




Fig. 111: Within-session evolution of firing preference for goal value.

Box plots are used to present the data as in Fig. 106. Gray dots indicate individual (cell) values. The values from the same cell in the two parts of session are connected by a gray line. **No difference** was ever detected between samples from the two parts of the sessions.

8.3.5 Conclusion

The analysis of value-changing sessions allowed us to assess whether or not hippocampal goal-related activity encodes aspects of goal value, which would support the hypothesis of a motivational nature of the observed goal-related signal. Our findings suggest that, at the population level, this activity is independent from the goal value. This result holds even in the late phase of recording sessions, when the behaviour indicates that rats efficiently learned the new goal value. Furthermore, the analysis of the reference sessions shows that the spatial property of the goal is likely to be more important than its associated value in modulating goal-related activity. The recorded population of interneurons did not seem to be involved in goal-related firing, and they did not encode any goal value. The PSTHs did not evidence any difference in the time course of the goal-related activity either, apart perhaps for the extinguished goal, where the sharp decrease usually occurring at the time of pellet dispenser activation was less obvious. Further investigation is required to rule out behavioural factors (e.g., time of exit from the goal zone) before concluding on the matter. In any case, **no modulation due to goal value was detected as far as the firing rate was concerned**, as demonstrated by the results obtained from the value preference rate.

These results suggest that the main determinant of goal-related activity, in addition to the presence of the goal itself, is the location of the goal rather than its value. By comparing these results with those obtained through our behavioural analysis (exposed in Chapter 7), we can note a dissociation between behavioural preferences, unambiguously directed towards the goal with the highest value, and the goal-firing preference of the hippocampal population, which was not influenced by goal value. Conversely, rats' behaviour did not show any preferred spatial goal (apart from one subject), whereas hippocampal goal firing often expressed a spatial bias. This discrepancy supports the hypothesis that pyramidal cell's activity is not directly linked to sensory or behavioural parameters, but it can rather operate, at least partially, independently from overt behaviour.

8.4 Discussion

Our electrophysiological work mainly focused on the hippocampal goal-related activity (demonstrated by CA1 place cells in Hok and collaborators, 2007a). We evidenced a significant goal-related firing from CA1 place cells in the **two-goal navigation task**. We showed that the goal-related signal was expressed by a large majority of the population of hippocampal pyramidal cells, **place and silent cells** altogether, and by cells from the **CA3** field. Interneurons did not seem to participate in this activity, although ruling out this hypothesis would require further investigation. The goal-related activity had already been shown to demonstrate a first level of spatial selectivity in that it is only present at goal locations. However, it is different from a place field since it is not expressed during foraging episodes (Hok et al., 2007a). We evidenced a **second level of spatial selectivity** of the goal-related firing, since a majority of cells contributing to the goal signal (67%) fired only for one of the two goal zones¹⁹. Still, a subset of goal-coding cells significantly fired for the two goal zones. Finally, this population signal is clearly **independent from modulations of goal value**. As such, it probably cannot be considered as an anticipatory reward signal.

¹⁹ one-goal coding / all goal-coding cells : 38/57

The goal-related activity is probably a signal of major importance within the dorsal hippocampus. If we consider a rat navigating towards a goal in the reference condition, a subset of his hippocampal place cells continuously fire, depending on its current location. At this moment, none of the silent cells are active. Whenever the rat enters the goal zone, a large population of the previously silent cells suddenly fire, along with a similarly sudden discharge of the majority of place cells. These two concurrent, massively distributed signals then stop at the moment of the feedback (i.e., the pellet dispenser activation).

In this discussion, we will propose hypotheses concerning this hippocampal goal-related activity. First, we will address the question of which mechanisms could give rise to this massive hippocampal signal. Then, we will ask what this activity could ‘code for’. Finally, we will discuss how this activity could be of use to the brain in general and for goal-directed behaviour in particular.

8.4.1 Possible sources of hippocampal goal-related signal

8.4.1.1 *Extra-hippocampal origins*

Since the hippocampus is anatomically and functionally connected to the spatial processing circuit of the brain, and since it also communicates with structures from the decision-making circuit (see Chapter 4, p. 55), several origins of the observed goal-related activity are possible. It could originate from the prefrontal cortex (via the nucleus reuniens), which was shown to express a spatial signal at the goal zone in the one-goal navigation task (Hok et al., 2005). It could result from afferent signals from reward-related areas such as the amygdala, the substantia nigra or the VTA (see Sec. 4.3.3, p.72). Finally, this signal could be computed from information received from the spatial circuit, such as the entorhinal cortex.

Previous studies showed that the goal-related activity was probably not of **prefrontal** origin, since an inactivation of the hippocampus did remove the spatial signal at the goal in the medial prefrontal cortex, but an inactivation of the medial prefrontal cortex did not have any detectable effect on the goal-related activity of hippocampal place cells (Burton et al., 2009; Hok et al., 2013; see Sec 4.3.4.1, p. 75). Rather, the medial prefrontal cortex seems to be an output structure of the hippocampus, at least, as far as goal-related activity is concerned.

The VTA and the amygdala were shown to express reward and value-related signals (see Sec. 2.4, p. 27; Sec. 4.3.2, p. 71 and Sec. 4.3.3, p. 72). They could provide the hippocampus with an anticipatory reward signal during the delay phase of the continuous navigation task. Concerning the **VTA**, we do not believe this hypothesis is supported by the literature. Indeed, the VTA has been shown to interact with the hippocampus rather after the occurrence of a feedback. In the model by Lisman and Grace (2005), the VTA sends signals to the hippocampus once the latter has detected novelty in the ‘world state’. Experiments showed that dopaminergic innervations had to be blocked specifically following learning to prevent hippocampal LTP (Ghanbarian and Motamedi, 2013) while a blockade prior to, or 20 min after learning did not alter LTP. Bethus and collaborators (2010) showed that inactivating the hippocampal dopaminergic receptors did not impair the encoding or retrieval of

information; rather, it impaired its long-term retention. Overall, these studies indicate that the VTA-hippocampal interactions are involved in the maintenance of memory. Why would a signal from the VTA aiming at consolidating memory arrive when no feedback has been received on performance yet?

Concerning the **amygdala**, it was recently shown to interact with the hippocampus via increased synchronisation when rats were expecting a highly probable reward (Terada et al., 2013). Also, stimulations of the amygdala impaired the stability of place cells (Kim et al., 2012). The amygdala mainly projects to CA1 but also sends light-to-moderate connections to the CA3 subfield (Pitkänen et al., 2000). These data indicate that inputs from the amygdala could contribute to the goal-related signal. Because the amygdala inputs are much stronger in CA1, one would expect the goal-related signal to be more widespread in CA1. However, our results showing that goal-related activity is expressed by CA3 neurons in similar proportions as compared to CA1 neurons are not in line with this hypothesis. Moreover, the hippocampal goal-related activity is observed in the dorsal hippocampus, while the amygdala innervates the intermediary-to-ventral hippocampus (Pitkänen et al., 2000; Strange et al., 2014). Overall, it is then unlikely that the goal-related signal directly results from amygdala inputs, although this question requires a deeper research. For example, the projections from amygdala to CA3 only originate from the basal nucleus of the amygdala, contrary to CA1 which receives inputs from different amygdaloid nuclei (Pitkänen et al., 2000). The specific role of these different inputs, to our knowledge, has not been addressed yet.

The **substantia nigra**, in contrast to the amygdala, was shown to project to the dorsal ('posterior') hippocampus (Scatton et al., 1980). Neurons expressing reward expectation signals were found in this dopaminergic region (Schultz et al., 1997). A recent study showed that a lesion of the substantia nigra disrupted the firing of dorsal CA3 place cells (Retailleau et al., 2013). More precisely, this lesion impaired the correlation of CA3 firing rate with performance in a Y-maze task. Thus, the substantia nigra could send reward signals to the hippocampus and contribute to the goal-related activity. However, the direct dopaminergic afferents to the hippocampus are quite sparse and poorly described (to our knowledge, only the study from Scatton and collaborators, 1980, reported those afferents). Therefore, we cannot conclude on the hypothesis that dopaminergic afferents from the substantia nigra could play a role in driving hippocampal goal-related activity.

All the above-mentioned 'motivational' structures express value-related signals. However, our results support the hypothesis that goal-related activity in the hippocampus is independent from motivational inputs. First, goal-directed activity is, in majority, expressed at only one of two goal zones. Thus, it holds a spatial component (see Fig. 100, p.160). This finding indicates that a putative motivational aspect does not prevail over the spatial control of goal-related firing. Second, goal-related activity is independent from goal value (see Sec. 8.3.4, p. 171), which is a major argument against the reward expectation coding hypothesis.

The expression of goal-related activity by silent cells can actually help to better understand the characteristics and the possible origin of this signal. To formulate a much speculative hypothesis, we would like to bring together the results from two seemingly unrelated studies (Epsztein et al., 2011;

Thome et al., 2014). In the first study, the authors recorded place and silent cells using intracellular recording techniques (see Sec. 5.2.1, p. 87). They showed that silent cells had a higher spiking threshold than the place cells active in an environment. Actually, the membrane potential of place cells was shown to depolarise before entry in the place field (Harvey et al., 2009; Epsztein et al., 2011). Of much interest is the result that extra-field spikes from place cells occur without the depolarisation specific to in-field spikes (Epsztein et al., 2011). Linking these results with goal-related activity supports the idea that at least two types of coding would coexist in goal-directed tasks in the hippocampus: the place field firing, related to a depolarisation at the cellular level, and the goal-related firing, not linked to a depolarisation and visible also in silent cells (the membrane potential of which is not depolarised). How, then, could these cells fire if it is more difficult to trigger a spike because of their higher spiking threshold? We suggest that the recent results from Thome and collaborators (2014) represent an interesting lead. These authors showed that around 50 % of CA1 cells (28 % in CA3) had an axon-carrying dendrite, i.e., their axon emanated from a basal dendrite and not from the soma itself, as is generally the case. They observed that input onto these privileged dendrites were more likely to trigger spiking of the cell than input from other dendrites. In a highly speculative view, we could suppose two modes of functioning of hippocampal pyramidal cells. One of them would be the depolarised mode of place cells when in the vicinity of their place field. In this case, the spatial code would be preponderant. The other mode would be the non-depolarised mode, in which only privileged inputs could trigger spiking, and which would be the signature of non-local events such as replay, forward probes, and possibly goal-related activity. To date, the origin of these privileged dendrites is still unknown. The authors propose that excitatory synapses to these dendrites could come from the contralateral CA3, the ipsilateral CA2, but also from CA1 collaterals, the amygdala, or the entorhinal cortex. In any case, it would be of much interest to record with the above-mentioned intracellular techniques in the two-goal navigation task to assess what are the intracellular mechanisms at stake during the goal period and what mechanisms make silent cells fire in these conditions and not others.

Altogether, these arguments suggest that the goal-related signal may be either internally computed in the hippocampus, or possibly arises from entorhinal inputs. Interestingly, Zhang and collaborators (2013) showed that all types of cells from the entorhinal cortex projected towards the hippocampus: grid cells, head direction cells, but also non-spatial cells. A possible way to test the importance of entorhinal inputs for goal-related activity would be to inactivate the entorhinal cortex while recording hippocampal cell activity. In addition, one could try to search for goal-related correlates in the entorhinal cortex. To our knowledge, the question of goal-related activity in the entorhinal cortex has not been asked yet.

Note that in this short overview of structures projecting to the hippocampus, we did not mention other structures that could be at the interface between the hippocampus and the decision-making circuit. The perirhinal cortex, for example, projects to CA1 and receives projections from the orbitofrontal cortex and the amygdala (Deacon et al., 1983).

8.4.1.2 *Intra-hippocampal origin of goal-related activity*

The goal-related signal could arise from intra-hippocampal computations. The fact that CA3 cells show this goal-dependent activity supports this hypothesis, since CA3 is much more isolated from external influences than CA1 (see Sec. 4.1, p.55). In our results, the distributions of place fields in the recorded CA1 and CA3 population of place cells were not significantly different from each other. However, visual inspection of the distributions of place fields' centroids revealed a possible difference, which might not be detectable in the polar distribution of fields, but probably in their radial position. More precisely, CA3 cells seemed to be more concentrated in the centre and possibly closer to the goals, while CA1 cells would be more peripheral (see Fig. 99, p. 159). Finer analysis perhaps combined with a larger cell sample would allow a conclusion to be drawn about this issue.

Interestingly, we observed that the CA3 population seemed to behave in a bimodal way (exemplified in the side firing preference index plots, Fig. 108, p. 169 and Appendix V, p.205). Indeed, CA3 goal-related activity occurred only at one of the two goals. The proportions of cells coding for one or the two goals is in line with this bimodal hypothesis, as shown in Fig. 100 (p. 160): CA3 cells tended to represent only one of the two goals and more rarely both. Note, however, that these proportions were not statistically different from those of the CA1 population. Overall, this bimodal distribution of the side preference of CA3 cells is in line with the theory stating that this hippocampal subfield would be organised as attractor networks (see Sec. 1.2.5.2 p. 19 and Sec. 5.2.4 p. 91). Kubie and collaborators (2007) indeed suggested that the goal-related activity of CA1 could reflect the activation of a goal-specific attractor in CA3. We believe this bimodal aspect of the goal-related coding requires further research. In particular, it could be interesting to see to what extent CA3 and CA1 cells differ when more than two goals are simultaneously accessible.

Another possible origin of the goal-related firing in the dorsal hippocampus could be the intermediary or the **ventral hippocampus**. It was recently shown that a temporary inactivation of the intermediate-to-ventral hippocampus impairs the ability of rats to update the value of a goal (De Saint Blanquat et al., 2013). Moreover, this region of the hippocampus is much more strongly connected to motivational inputs, as reviewed in Chapter 4 (p.55). A hypothesis would be that the ventral hippocampus receives goal-related motivational inputs during learning, and that intra-hippocampal processing would remove the motivational aspect of this firing in order to store only the spatial aspect in the dorsal hippocampus. In this view, if ventral hippocampal cells were shown to express goal-related activity (which is, to date, unknown), this activity should be modulated by goal value.

Overall, anatomical and functional arguments are rather in line with an intra-hippocampal origin of the goal-related activity, which could come from the ventral hippocampus and/or be transmitted from CA3 to CA1. In our study, the population PETHs of goal-related firing did not highlight any obvious difference in timing between CA1 and CA3. However, simultaneous recordings in CA1 and CA3 could be performed to assess the precise time-locking of CA1 and CA3 goal-related activity, and to possibly determine in which of the two subregions of the hippocampus this activity occurs first.

8.4.2 Information conveyed by hippocampal goal-related activity

There are several possibilities concerning the nature of the information conveyed by goal-related activity to the rest of the brain. Since it is only observed when rats are waiting at the goal zone but not when they run through it while foraging (Hok et al., 2007a), it cannot be considered as a spatial signal similar to that of a place field. As previously stated, it is probably neither entirely motivational nor linked to increased attention, since both phenomena would generate a homogeneous signal regardless of the goal location, and probably without the specific temporal profile observed on PETHs (Hok et al., 2007a, 2007b). The hypothesis put forward by Hok and collaborators (2007a) is that the goal-related activity could be a **signal indicating that the rat is at the proper goal location**. This hypothesis is supported by their results in the cued version of the task where the goal-related activity started to fire earlier in the cued goal zone compared to the uncued goal zone (see Fig. 67, p.101). We did not observe such time shift of firing in the cued version in our experiment (Fig. 103, p. 164). However, as stated previously, this could be due to differences in training; in our case, rats were not used to navigate with the cue and probably did not rely much on it to localise the goal. Other ways to test for this ‘spatial congruence hypothesis’ could be to modify some of the spatial aspects of the task such as the size of the goals or the size of the environment. Indeed, within a 76-cm arena, the size of the goals compared to the navigable space is quite large and goals are probably not too difficult to locate. A spatial congruence signal might be modulated by the complexity of localisation while other types of signals would not.

Another hypothesis is that, similarly to the place representation of place cells, the goal-related activity would **represent the goal**. Thus, this activity would be ‘recalled’ at the goal zone because of the co-occurrence of the internal and external inputs that first triggered its spiking. However, it could also be active in other situations (as the non-local events of place cells), for example when choosing a goal, or when planning a trajectory towards this goal.

The fact that silent cells selectively fire at the goal location supports the goal representation hypothesis. In our task, the majority of the recorded silent cells expressed a goal-related activity significantly higher than the activity occurring anywhere else in the environment (see Fig. 100, p.160).

Thus, in the two-goal navigation task, the main determinant of firing of the majority of silent cells is the goal. Paralleling the idea that place cells represent a place because space is the main determinant of their firing, silent cells (in these conditions) would represent the goal.

Other arguments do not support the hypothesis that goal-related activity reflects a goal representation. Indeed, the concept of goal probably encompasses many features, among which aspects of the goal such as its value. Our results demonstrate that the goal-related firing is not modulated by goal value. From a more general point of view, the hippocampus is known for its role in spatial processing, especially when a place strategy is involved, i.e., when the task requires to locate an unmarked goal with respect to distal cues. The results from the two-goal navigation task suggest that the spatial aspect of the goal influences the goal-related firing more strongly than other aspects of the goal for a majority of cells expressing goal-related activity. This allows the ‘goal

representation’ hypothesis to be refined into a ‘**spatial goal representation**’ hypothesis: the hippocampal goal signal would mainly represent the spatial characteristics of the goal. This hypothesis is supported by the results from Kobayashi and collaborators (2003), who also found a kind of goal-related activity in a task where rats had to shuttle between two uncued goal locations. In this study, a forward shift of the place fields of a subset of place cells was observed during learning, meaning that cells tended to fire closer to the goal locations. This was usually the case for either one or the other of the goal zones, but not necessarily both, indicating that there was a spatial component underlying this forward shift. It is unclear whether we can really relate these results to ours, since their analysis concerned the main field of the cells. However, it is possible that the observed main field shift resulted from the generation of (out-of-field) goal-related activity. In a more recent study, McKenzie and collaborators (2013) taught rats to memorize several goal locations on an annular apparatus. They appeared to detect goal-related activity, since they noticed cells that would fire on arrival to a subset of goal areas. Importantly, they specify that these cells were never active at all goal locations, indicating that the spatial component of the goal is important, coherently with our results. Furthermore, the estimation of position of the animal obtained from decoding the hippocampal population was the most accurate at positions where the goal-related firing was most likely to occur. Thus, goal-related activity allows for a better estimation of position.

Other arguments might indicate that the goal-related activity is encoding a ‘spatial goal representation’. Place cells were shown to be recruited in larger numbers when the size of the environment to be represented increased (Fenton et al., 2008; Park et al., 2011). Perhaps when the ‘goal-related’ demands of a task increase, the number of cells expressing a goal-related activity also increases? Although this was not directly tested, three elements suggest that it could be so. First, comparing the number of place cells expressing a goal-related activity in the place or in the cued navigation task of Hok and colleagues (2007a) indicates that there would be more ‘goal-coding’ place cells in the place task: 84% versus 40% in the cued version of the task²⁰. Second, comparing the proportions of silent cells detected in the two-goal version with those in the one-goal version indicates that more silent cells are recruited in the two-goal (28.5%) than the one-goal navigation task (12%). This is still the case when we only consider the goal-coding silent cells from our task (17.4% of the total recorded pyramidal cells)²¹. However, we did not perform statistical comparisons of these proportions. By contrast, the proportions of place cells significantly coding for any of the two goals in our conditions were actually lower than in Hok and collaborators’ (2007) study. We found that 74 % of the place cells with a non-overlapping field expressed a significant goal coding (compared to 84 % in the above-mentioned study)²². A third argument comes from a study recording

²⁰ Uncued version: $68/81 = 84\%$; cued version: $34/87 = 40\%$. However, the number used as cell total in the uncued version ($n = 81$) is the number of cells whose field did not overlap on the goal zone while it is not indicated whether the number used for the cued percentage ($n = 87$) represents all place cells regardless of the position of their main field or only those with a non-overlapping place field.

²¹ In Hok et al., 2007: 22 silent cells over 174 recorded pyramidal cells = 12 %.
In our study: 41 silent cells over 144 recorded pyramidal cells = 28.5%;
only considering the 25 goal-coding silent cells : 17.4 %.

²² In our study: 32 out of 43 place cells whose field did not overlap on the goal zone = 74%.

place cells in an annular water maze task where the position of the goal was changed. While the rat was actively searching for the platform because the previous position was not active any more, a subset of cells that were previously silent fired upon encountering the new goal location (Fyhn et al., 2002). However, this firing came back to baseline as the rat learned the new location. Thus, it might not be comparable to the one observed in our study.

In conclusion, we believe that further research is required to evaluate whether an increase in the ‘goal processing’ component of the task would significantly increase the recruitment of the silent cell population.

Another question concerning the kind of information conveyed by the goal-related signal is the type of **neural code** it relies on (see Sec. 1.2.2, p.11). In the place field firing, two types of codes coexist. First, the rate code indicates the current location via modulations of firing rates. Second, the phase precession phenomenon allows the activity of a place cell to provide spatial information through a temporal code. As stated previously, taking the spiking phase with respect to theta into account helps improving the decoding of the position of the rat (Jensen and Lisman, 2000). Concerning the goal-related activity, we characterised it relying on its firing rate. However, it could also be characterised through a temporal code, for example, the synchrony of spikes with the theta or between pairs of cells. As far as the latter proposition is concerned, Hok and collaborators (2007a) did evidence more synchronous firing at the goal, although this would require further research because of the quite small sample on which this analysis was performed. Also, the specific temporal profile visible in peri-event time histograms of goal-related activity may underlie the importance of a temporal code.

8.4.3 Use of the hippocampal goal-related signal

We believe that the ‘spatial congruence hypothesis’ and the ‘goal representation hypothesis’ are compatible and even complementary, and that they could participate in driving navigation. Indeed, the congruence between a goal representation (that could be stored in silent cells or in CA3) and current sensory inputs (processed by CA1), upon arrival at the goal, could be the signal indicating that the rat arrived at the proper place. In this view, the synchronisation of firing from both subfields would precede the stop of the rat at the goal location. This would be testable through simultaneous recordings in CA3 and CA1 complemented by a high-resolution analysis of speed profile. Insights might already be given by the speed analysis performed in a previous study in the one-goal task (Hok et al., 2007b). Fig. 112 represents the mean speed profile of rats entering the goal zone, in the cued and the uncued version of the task. Note that the two speed profiles did not differ from each other. The right panel presents the PETH for goal-related activity of CA1 place cells recorded in the place task. One can see that the beginning of firing occurs approximately 400 ms after the entry in the goal zone. Concerning the speed, it decreases between 0 and 500 ms after entry in the goal zone. Unfortunately, the precision of this analysis is not sufficient to estimate which occurred first. Thus, these data do not reject the possibility that goal-related firing can precede the decrease in speed, upon entrance in the goal zone.




Fig. 112: Speed profile and corresponding PETH of goal-related activity.

Left : Speed profiles in the place (uncued) and cue task during extinction periods. Both profiles are similar. 0 indicates the entry into the goal zone.

Right : PETH of goal-related activity in the place task (200 ms bins) during extinction trials.

Note that these speed profiles concern the extinction condition of this experiment because speed profiles were not available for the reference condition.

From Hok et al., 2007b.

The hippocampal goal representation could also participate in trajectory planning, by being transiently activated at the beginning of a trial in a spatial task. This could trigger prospective mechanisms such as forward probes or place cell sequences (Johnson and Redish, 2007; Pfeiffer and Foster, 2013) allowing several paths towards this goal to be virtually tested. The final decision could be the one that generates the highest coherence between a probed sequence and the goal representation.

Finally, the goal representation provided by the hippocampus appears to primarily encode the spatial aspect of the goal. Anchored on this spatial aspect, other features of the goal state could be signalled by other brain structures. In this view, the hippocampus would provide its contribution to goal representation to other decision-making structures either in the course of planning, or while actually navigating towards a previously defined goal. Congruency between the goal representation provided by the hippocampus and that provided from either other structures or external inputs at the moment in which a subject reaches its goal, could then be accompanied by the ‘feeling’ of reaching the proper decision.

‘The feeling of having reached the right decision might, thus, be associated with brain states that are distinguished by particularly high coherence. As is hypothesized for the integration of multimodal sensory signals into coherent percepts, such states are likely to emerge when the distributed processes occurring simultaneously in different cortical areas can be brought to match.’

Singer, 2005

8.4.4 Limits and further investigations

8.4.4.1 *How to differentiate goal-related activity from infield firing?*

In the current study, we only used cells whose main field did not overlap any of the two goal zones, which reduced the exploitable amount of cells. Moreover, we focused on their goal-related activity and did not analyse their place field firing. However, cells with their fields overlapping on goals probably also participate in navigation. It would be interesting to see whether or not one can differentiate spikes emitted in the main field (that can be characterised by their discharge in bursts, the existence of phase precession, and overdispersion) and goal-related spikes (that do not seem bursty, are possibly synchronous between cells, and shared by a larger population of cells). Actually, the existence of goal-related activity in the main field is questionable, given the elements exposed in Sec. 8.4.1 (p. 174) of the discussion, notably, the possibly different cellular mechanisms associated with extra-field spikes and silent cells. Thus, there might be a separation between cells whose field is located at a goal zone and cannot express goal-related activity, and cells whose field is far from the goal and can express a goal-related activity. Whether or not the phase precession phenomenon takes place in the goal-related activity could be another way to dissociate in-field from goal-related activity.

Note that we used the firing rate to evaluate the significant character of goal-related activity. However, if information in this activity relied on a temporal and not a rate code, the firing rate would not be much relevant to goal-related firing. Thus, the issue of synchrony of spikes at the goal requires further research and it would benefit from simultaneous, large-scale recordings of neuronal populations.

Another discrepancy between the goal-related firing and the main field activity of place cells appears to be their modulation by reward value. Indeed, Lee, Ghim and collaborators (2012) showed that the main fields of a subset of hippocampal place cells were modulated by value-related parameters such as action value and outcome value. Looking into single-cell data in our experiment indicates that it might also be the case for a subset of cells that seem to increase or decrease their firing preference for value between the first and the second part of sessions (see Fig. 111, p. 172). However, such modulations do not appear to be consistent at the population level. In that matter, our results are closer to the ones from Tabuchi et al. (2003), who did not evidence any reward magnitude – associated activity in dorsal CA1 pyramidal cells.

8.4.4.2 *The novelty hypothesis*

According to Lisman and Grace (2005), the hippocampus would be involved in a loop with the VTA. The hippocampus would send a **novelty signal** to the VTA, which in turn would trigger learning in the hippocampus (see Sec. 4.3.3 p. 72). Our experiment might be appropriate to address this hypothesis by focusing on the activity at the moment of feedback on goal activation (i.e. just after the pellet dispenser activation, or non-activation in the case of extinction sessions). We could compare the firing rate in the beginning of a value-changing session with the firing rate for the same goal at the end of the same session. If a novelty signal exists, it should be seen, and only for the goal whose

value was modified. Because CA3 was shown to project to the VTA via the lateral septum, this signal might well be more visible in CA3.

8.4.4.3 *Vicarious trial and error*

Because the two-goal navigation task involves repeated choices between the goal locations, we suggest that vicarious trial-and-error (VTE, see Sec. 2.6, p. 29) events might be detected during the choice phase. Although the precise moment of decision is hard to be established, we propose that it could be estimated using variations of the orientation of rat's head. After consumption of the reward, rats could engage in VTE; once their choice is made, they would probably perform a direct trajectory towards the chosen goal. During this choice period, place cell population events could be detected, such as forward probes or sequences (Johnson and Redish, 2007; Pfeiffer and Foster, 2013). We suggest that events involving the goal representation (i.e., a distributed activity of place and silent cells) could also occur during choice. Similarly to spatial decoding algorithms, the hippocampal population could be decoded to assess whether the representation of a specific goal is reactivated. Correlations could be made between the amount of reactivation of a given goal representation and the final choice; we might expect that the most reactivated goal would be the one eventually chosen by the rat (as in Pfeiffer and Foster, 2013). As mentioned previously, it is also possible that another type of coding is involved, and that it is not the mere reactivation of the goal representation that could be used to take a decision, but rather the amount of synchronous firing generated by a goal representation. In this scheme, the chosen goal could be the one whose reactivation would be correlated with a more synchronous firing between hippocampal neurons. Yet more probable, the chosen goal could correspond to the one generating the highest level of synchrony between population of cells in the hippocampus and structures from the decision-making circuit, such as the medial prefrontal cortex. Indeed, the medial prefrontal cortex probably participates in this 'goal representation' since goal-related firing has also been evidenced in this area (Hok et al., 2005).

Furthermore, we propose that VTE events would be more numerous **following a change in goal value**. In cases of a large difference between two goal values, the cost of choosing the least rewarded goal is higher. Then, rats might spend more time 'deliberating' in order to make sure to take the proper decision.

8.4.4.4 *Overdispersion*

Overdispersion refers to the intrinsic variability of firing of a cell for similar passes through its place field (see Sec. 5.1.5, p. 86). Interestingly, a change in reward contingency was demonstrated to decrease the overdispersion of place cells (Wikenheiser and Redish, 2011). At the moment of change, rats probably get more focused on the environment, searching for probable causes of the change, or encoding possibly new information. The overdispersion of place cells could be assessed in our task to possibly reproduce these results (which were collected in an annular task). One could also link the overdispersion of place cells to the behavioural profile of rats, notably their individual exploration / exploitation ratio (see discussion of Chapter 7, p. 145). Overdispersion could be related to the 'exploratory' tendency of place cells and their flexibility, that is to say, the fact that they could easily

encode new information – when relevant for the task. More overdispersion at the neural level could be correlated with a more exploratory profile at the behavioural level. Indeed, overdispersion was shown to be correlated with performance in a spatial working memory task (Hok et al., 2012).

Chapter 9 – Conclusions

It is only recently that the hippocampus, a structure well known for its crucial role in spatial memory, has attracted the interest of neuroscientists concerning its possible contributions to decision-making. According to Doll and collaborators (2014),

‘The emerging correspondence between memory and decision systems has important consequences for our understanding of the latter in particular, since the hippocampal relational memory system provides a concrete and relatively well characterized foundation for understanding the otherwise rather mysterious mechanisms for model-based decision-making.’

Doll et al., 2014

Our work, which combines behavioural and electrophysiological approaches, is at the interface between the spatial cognition and decision-making domains. We designed a task in which rats have to choose between two uncued spatial goals and adapt their behavioural strategy to modifications of goal values. This allowed us to demonstrate that rats are able to memorise the location of two goals in allocentric space, and they can flexibly decide which goal to aim at according to their expected value. Electrophysiological recordings of dorsal hippocampal neurons in this task provided new insights into the hippocampal goal representation. The main electrophysiological results of our work can be summarised in three points.

First, at the level of the goal-related activity of CA1 place cells, we showed that:

- i. CA1 place cells express a goal-related activity when a rat must locate two hidden locations and freely choose between them.
- ii. This goal-related activity holds a spatial aspect, which prevails over possible behavioural or motivational aspects of the goal.
- iii. This goal-related activity is independent from goal value.

Second, we found that the population of ‘silent’ pyramidal cells (i.e., neurons with no place field in the recording environment) also massively express a goal-related activity, with the same characteristics as CA1 place cells. Thus, silent cells in the two-goal navigation task appear to be an example of dense distributed coding of the goals, which parallels the sparse distributed place cell coding. The fact that silent cells might be recruited to represent the goal because of the spatial complexity of the task requires further research.

Last, we also provided evidence for goal-related activity in a majority of CA3 place and silent cells. Overall, the characteristics of the CA3 goal representation are similar to those of CA1. However, subtle differences indicate that the two representations could subserve different roles and calls for further investigation.

Throughout the thesis, we saw that the activity of the hippocampus can be read out at multiple levels. The activity of a place cell can participate in the representation of position. It is believed to be able to alternate between different frames of reference which give rise, at least in part, to the overdispersion phenomenon. Furthermore, place cells are not ‘stuck’ in the present: they can be used to ‘navigate’ through memory and towards the future, by replaying experienced events, or planning possible sequences to a goal. At the population level, the combination of silent and active pyramidal cells defines the memory of an environment. Our results on the goal-related activity of hippocampal pyramidal cells suggest an additional level of coding, which transcends the classical dissociation between place and silent cells: as far as goal-related activity is concerned, both populations behave similarly. We showed that this goal-related coding is mainly spatial and is not a reward expectation signal as had previously been proposed. Finally, this allows us to draw a general conclusion about the role of the hippocampus in spatial decision-making.

Although it was central to our work, the hippocampus cannot be said to be essential to the survival of an individual, as demonstrated by lesion studies in rats and cases of amnesic patients like H.M. However, the hippocampus gains gradual importance with the spatial requirements of a situation. While, through latent learning, it is able to continuously record and update a map-like representation of the external world, the hippocampus reveals its major role when flexible manipulation of memorised information is involved. In that case, the neural computations it performs combine sensory input and stored information to provide other brain structures with selected and relevant information about the ‘world state’. Through goal-related activity, the hippocampus can be said to contribute to decision-making, by signalling the spatial aspect of the goal to other structures of the decision-making network, such as the prefrontal cortex and possibly the striatum. During goal-directed navigation, the hippocampal goal signal might be of crucial importance to the decision of stopping at a particular location.

The work presented in this thesis can contribute to a better understanding of the neural bases of flexible spatial navigation, by clarifying the nature of the information processed by the hippocampus and its possible contributions to decision-making. The rat hippocampus seems to be able to dissociate spatial features from other parameters that might interfere with proper manipulation of information, such as the expectation of a large reward. In humans, one can sometimes regret a decision that was made under the influence of emotions. More extreme, pathological situations evidence the importance of a properly functioning decision-making system. The hippocampus and its interactions with structures from the decision-making circuit were recently shown to be relevant in schizophrenia (Dere et al., 2010; Lisman et al., 2010; Suh et al., 2013; Ledoux et al., 2014), a pathology most often associated with impairments in behavioural flexibility (Floresco et al., 2009). Thus, a better understanding of how the brain manages to separate, at times, and combine, at others, ‘cold’ reasoning with emotions might be of help in finding new techniques or treatments for decision-making related pathologies.

Résumé (long)

Lors de ce travail de doctorat, nous nous sommes penchés sur le rôle de l'hippocampe dans la navigation orientée vers un but et, plus précisément, sur la contribution apportée par les neurones hippocampiques à la représentation d'un but spatial. L'hippocampe est une structure cérébrale connue en particulier pour son rôle dans la mémoire spatiale, pour lequel il est très étudié chez les mammifères (notamment le Rat). Ces études ont permis de mettre en évidence l'existence de « **cellules de lieu** » hippocampiques, des neurones dont la décharge est fortement corrélée à la position de l'animal dans un environnement (O'Keefe & Dostrovsky, 1971). Cependant, les interactions entre la représentation neurale de l'espace sous-tendue par l'activité des neurones de l'hippocampe et le comportement de navigation orientée vers un but restent à élucider. Des études récentes ont permis de mieux comprendre ce lien, en montrant notamment que les cellules de lieu pouvaient s'activer hors de leur position « préférée ». Ces activations auraient plusieurs fonctionnalités : consolidation mnésique avec le phénomène de *replay* (Louie & Wilson, 2001 ; Foster & Wilson, 2006) ; recherche prospective du meilleur trajet pour atteindre un but avec les *forward probes* (Johnson & Redish, 2007) et l'activation en séquences des cellules de lieu (Pfeiffer & Foster, 2013) ; représentation du but avec l'**activité liée au but** (Hok et al., 2007a). Ce dernier phénomène correspond à une activation de la majeure partie des cellules de lieu hippocampiques lorsqu'un rat attend la libération d'une récompense dans une zone-but, i.e., un lieu non-indiqué que l'animal a appris à localiser par rapport à sa position dans l'espace. L'origine et le rôle de cette activité liée au but sont encore méconnus. Par une approche mêlant comportement et enregistrements électrophysiologiques de l'activité neuronale, nous avons cherché à mieux comprendre cette activité, qui pourrait avoir un rôle dans la prise de décision au cours de la navigation, mais qui pourrait aussi refléter l'anticipation de la récompense. Une meilleure compréhension de ce mécanisme pourrait aider à clarifier la manière dont les structures cérébrales traitant l'information spatiale et celles responsables de la prise de décision communiquent pour diriger le comportement.

Partie I : Introduction

Pour situer le projet dans son contexte, il est nécessaire de définir certaines notions essentielles. Il existe en particulier un concept sur lequel nous nous appuyons, et qui est couramment utilisé en sciences cognitives : celui de représentation neurale. Selon Roitblat (1982), une **représentation** est la trace d'une expérience passée qui permet à cette expérience d'avoir un effet sur le comportement futur. Plus récemment, deCharms et Zador (2000) ont défini une représentation comme un message qui utilise les règles du code neural pour transmettre de l'information. Ces mêmes auteurs définissent le **code neural** comme un système de règles et de mécanismes par lesquels un signal transporte de l'information. Une représentation est caractérisée par deux aspects : son contenu et sa fonction (deCharms & Zador, 2000). Le contenu d'une représentation concerne l'information transmise tandis que sa fonction correspond à l'impact qu'elle peut avoir sur d'autres traitements neuronaux et, éventuellement, sur le comportement.

Certaines régions cérébrales, notamment les aires sensorielles et les aires motrices, contiennent des neurones qui déchargent en réponse à la présentation d'un stimulus particulier

(Hubel & Wiesel, 1959) ou lors de la réalisation d'un mouvement donné (Georgopoulos, 1986). Dans la définition donnée précédemment, l'activité de ces neurones représente soit le stimulus (par exemple, la présence d'une barre ayant une certaine orientation à un certain endroit du champ visuel), soit l'action à réaliser (un mouvement du bras vers la droite). Ces représentations se situent dans un référentiel **égo-centré**, c'est-à-dire dans le système de coordonnées du récepteur sensoriel ou de l'actionneur concerné (Hartley et al., 2014). L'information peut aussi être transmise dans un référentiel **allocentré**, où elle est alors représentée par rapport à un référentiel externe, indépendamment de l'orientation du sujet ou de celle d'un récepteur ou actionneur donné. C'est le cas pour l'activité spatio-sélective des **cellules de lieu** de l'hippocampe, qui déchargent lorsque l'animal se trouve à un endroit bien particulier de l'environnement, quels que soient l'orientation de l'animal ou les mouvements qu'il réalise à cet endroit (pour revue, voir Arleo & Rondi-Reig, 2007 ; Burgess, 2008 ; Hartley et al., 2014). La position pour laquelle la décharge de ces cellules est maximale dans l'environnement est appelée « **champ de lieu** » (cf. Fig. 39, p. 61). L'activité combinée de la population de cellules de lieu permet de décoder précisément la position de l'animal (à environ 5 cm près ; Wilson & McNaughton, 1993). De plus, cette activité est stable d'une exposition à l'autre en l'absence de changements dans l'environnement (Thompson & Best, 1990). D'abord découvertes chez le Rat (O'Keefe & Dostrovsky, 1971), des cellules à l'activité similaire ont été mises en évidence chez plusieurs espèces de mammifères, y compris l'Homme (Ekstrom et al., 2003). Depuis la découverte des cellules de lieu de l'hippocampe, de nombreuses études se sont penchées sur les paramètres permettant à ces cellules de construire leur décharge spatio-sélective ainsi que sur les caractéristiques de cette décharge. Il en résulte que plusieurs modalités sensorielles contribuent à la décharge des cellules de lieu. Les informations issues de ces différentes modalités peuvent être classées en deux catégories : les informations **idiothétiques**, qui renseignent sur les mouvements propres de l'animal, et les informations **allothétiques**, qui concernent l'environnement (Arleo & Rondi-Reig, 2007). Les informations idiothétiques proviennent du système vestibulaire, de la proprioception, de la copie efférente de la commande motrice et des flux sensoriels (par exemple, le flux optique). La modalité visuelle est l'une des sources principales d'information allothétiques, qui comprend aussi les modalités olfactive, auditive et somatosensorielle. Si les repères visuels qui permettent de polariser un environnement (par exemple, une carte accrochée au mur d'une arène circulaire) subissent une rotation en l'absence de l'animal, les champs de lieu des cellules enregistrées avant le déplacement vont subir la même rotation lorsque l'animal sera replacé dans l'environnement (O'Keefe & Conway, 1978 ; Muller & Kubie, 1987). Dans le noir, l'activité spatio-sélective de ces cellules perdure (Quirk et al., 1990), mais est perturbée si la quantité d'informations allothétiques disponible diminue, notamment en supprimant les indices olfactifs (Save et al., 2000 ; Poucet et al., 2000). En résumé, dans ce type de tâche 'simple' où l'animal doit trouver et consommer des granulés de nourriture distribués aléatoirement dans l'environnement, l'activité des cellules de lieu correspond à une **représentation multisensorielle et allocentrée de la position** de l'animal dans l'environnement.

Dans une tâche de recherche de nourriture aléatoire, l'animal n'a pas nécessairement besoin de se localiser ni de localiser son objectif, ni même de mémoriser l'organisation spatiale de l'environnement. Les études lésionnelles montrent d'ailleurs que l'hippocampe est nécessaire pour la réalisation d'une tâche seulement lorsque celle-ci implique un certain niveau de traitement allocentrique de l'espace, c'est-à-dire, lorsqu'un but doit être localisé en utilisant sa position par

rapport à des indices de l'environnement (Morris et al., 1981). Dans des tâches de ce type, l'animal utiliserait une **stratégie de carte**, contrairement à d'autres paradigmes où il peut utiliser une stratégie de réponse pour naviguer (c'est-à-dire, mémoriser les mouvements à réaliser pour atteindre le but au lieu de construire une trajectoire prospective, basée sur la connaissance de la position du but). Pour étudier plus précisément le lien entre l'activité des cellules de lieu et les décisions prises lors de la navigation orientée vers un but, il est nécessaire d'enregistrer l'activité des cellules de lieu dans des tâches plus complexes que la recherche aléatoire de nourriture, si possible reposant sur une stratégie de carte. L'un des paradigmes couramment utilisés pour étudier la mémoire spatiale et la navigation orientée vers un but est la **tâche de navigation continue** (Rossier et al., 2000). Celle-ci consiste à apprendre au rat à localiser une « zone-but » non-marquée dans un environnement. Le rat doit attendre un certain temps (généralement, 2 secondes) à l'endroit du but pour qu'un distributeur situé au-dessus de l'arène libère un granulé de nourriture. Ce granulé tombe à un endroit aléatoire dans l'environnement, ce qui fait que le rat alterne les trajets but – nourriture – but de manière continue avec des trajectoires virtuellement nouvelles à chaque fois. Plusieurs études ont tenté de manipuler la position préférée des cellules de lieu, soit dans cette tâche (Lenck-Santini et al., 2002, Kubie et al., 2007), soit dans d'autres (labyrinthe en croix : O'Keefe & Speakman, 1987 ; labyrinthe en Y : Lenck-Santini et al., 2001a) pour étudier l'effet d'un changement de position des champs sur le comportement de navigation. Dans tous les cas précédemment cités, les déplacements des champs de lieu étaient généralement en cohérence avec les décisions de navigation des rats, c'est-à-dire que la position où ils cherchaient le but était toujours au même endroit par rapport aux champs d'activité des cellules de lieu, que ce soit avant ou après modification de position. D'autres études semblent indiquer que l'hippocampe et les cellules de lieu pourraient participer à l'évaluation de trajectoires possibles vers un but. Comme dit précédemment, il est possible d'estimer précisément la position de l'animal en « décodant » l'activité de population de ses cellules de lieu. En utilisant de tels algorithmes de décodage, Johnson et Redish (2007) ont montré qu'à une intersection où l'animal devait faire un choix entre deux voies, l'activité de population représentait successivement les deux chemins possibles. Ces représentations prospectives n'avaient pas de lien avec le chemin finalement choisi. Plus récemment, Pfeiffer et Foster (2013) ont montré qu'avant l'initiation d'une nouvelle trajectoire vers un but non indicé se trouvant dans un environnement ouvert, l'activité de population des cellules de lieu représentait de manière transitoire plusieurs trajectoires possibles à partir de la position actuelle du rat. La trajectoire finalement réalisée par le rat correspondait précisément à l'une de ces trajectoires « virtuelles ». Ces résultats montrent que même l'activité « hors-champ » des cellules de lieu peut avoir du sens et que la mémoire de l'espace exprimée par l'activité des cellules de lieu pourrait aussi être utilisée dans le cadre de la navigation.

Finalement, pour pouvoir planifier une trajectoire vers un but, il est nécessaire de savoir localiser le but en question. Plusieurs études se sont penchées sur l'effet de la présence d'un but sur l'activité des cellules de lieu. En particulier, Hok et collaborateurs (2007a) ont utilisé la tâche de navigation continue pour montrer qu'une grande majorité de la population de cellules de lieu du champ CA1 de l'hippocampe²³ émettait une activité une fois l'animal arrivé dans la zone de but (cf.

²³ L'hippocampe peut être divisé en plusieurs « champs » : CA1, CA2 et CA3. L'organisation cellulaire et la connectivité de ces champs sont différentes. Le flux principal d'information arrive par le cortex entorhinal vers

Fig. 66, p. 101). Cette **activité liée au but**, hors du champ de lieu des cellules, présente un profil temporel caractéristique (cf. Fig. 67, p. 101), avec une augmentation de la décharge environ 500 ms après l'entrée dans la zone, puis un arrêt brusque une fois que le distributeur de nourriture s'est déclenché. Plusieurs hypothèses ont été émises quand à la nature et à l'origine de ce signal et certaines d'entre elles ont déjà été écartées (par exemple, celle d'un signal d'estimation temporelle, Hok et al., 2007b). Nous retenons en particulier deux hypothèses sur la nature et le rôle de ce signal :

- i. un signal spatial indiquant la position du but et confirmant qu'elle a été atteinte par l'animal ;
- ii. un signal motivationnel d'anticipation de la récompense, permettant d'associer lieu et récompense.

Les travaux réalisés lors de la thèse avaient pour objectif de mieux comprendre l'implication des neurones de l'hippocampe dans la navigation orientée vers un but, et en particulier de clarifier le rôle et l'origine de l'activité liée au but des cellules de lieu. Dans un premier temps, nous avons développé différents protocoles pouvant permettre de tester la nature spatiale ou motivationnelle du signal lié au but des cellules de lieu. Dans un deuxième temps, nous avons sélectionné l'un de ces protocoles et enregistré l'activité unitaire des neurones de l'hippocampe de rats entraînés dans cette tâche, ce qui nous a permis de tirer plusieurs conclusions sur la nature du signal lié au but et de préciser les hypothèses sur son rôle.

Partie II : contribution expérimentale

Expériences de dévaluation

La première partie de la thèse a été consacrée à la mise au point d'une tâche dans laquelle les rats devaient mémoriser la position de différents buts et leur associer une valeur. Nous voulions pouvoir estimer la valeur que les rats associeraient à un but à partir de leur comportement, de manière à pouvoir évaluer l'influence de la valeur du but ou de l'anticipation de la récompense sur l'activité liée au but. Dans une première série d'expériences, nous avons utilisé le phénomène de **dévaluation par satiété** pour diminuer la valeur d'une récompense et, indirectement, la valeur d'un but associé à cette récompense (les paramètres de ces expériences sont présentés en annexe I, pp. 200 et 201). La dévaluation par satiété consiste à laisser un rat consommer une récompense en grande quantité pendant un temps limité. La consommation de cette même récompense, lorsqu'elle est testée juste après la dévaluation, est diminuée par rapport au niveau de consommation pré-dévaluation (Balleine & Dickinson, 1991). Ce paradigme est généralement utilisé dans des tâches d'apprentissage instrumental, c'est-à-dire lorsqu'on attend du rat qu'il appuie sur un levier pour obtenir une récompense (Balleine & Dickinson, 1998). La dévaluation par satiété a très rarement été appliquée au domaine de la navigation spatiale, à l'exception notable d'une série d'études sur la mémoire pseudo-épisodique chez le rat initiée par Babb et Crystal (2006a). Dans cette expérience, les rats devaient mémoriser la localisation et le type de récompense associés à différents bras d'un labyrinthe à 8 branches (cf. Fig. 4, p. 7). La dévaluation sélective d'un type de récompense avait pour effet de diminuer les visites du bras associé à cette récompense, sans diminuer significativement les

le gyrus denté, puis passe par CA3, continue dans CA1 et enfin sort de la formation hippocampique par le subiculum.

visites des autres bras. Les auteurs ont ainsi montré que les rats étaient capables de mémoriser le type de récompense associé à un lieu particulier. Par ailleurs, l'enregistrement de cellules de lieu dans une version plus classique du labyrinthe à huit branches a mis en évidence un codage différentiel des bras appâtés et non-appâtés (Hölscher et al., 2003). Ce résultat pourrait être interprété (entre autres) par un codage de la magnitude de la récompense par le champ des cellules de lieu.

Nous nous sommes inspirés de ces deux protocoles pour concevoir une **tâche de labyrinthe radial** adaptée à l'enregistrement électrophysiologique et où la valeur de la récompense serait modifiée par un processus de dévaluation par satiété. Quatre paradigmes différents ont été testés, dans lesquels deux types de nourriture étaient disponibles dans certains bras du labyrinthe et les rats pouvaient apprendre le lien lieu-récompense au cours de plusieurs sessions d'entraînement. Ensuite, l'une des récompenses était dévaluée. Nous nous attendions à voir un changement dans l'ordre des visites des rats entre un test dans le labyrinthe pré-dévaluation et post-dévaluation, en faveur des bras contenant la nourriture non-dévaluée. Seulement, cela n'a été le cas dans aucun des protocoles testés.

Une autre approche a alors été employée, s'appuyant sur le paradigme de la **tâche d'alternance continue** (cf. Fig. 32, p. 52), dans lequel un rat doit successivement tourner à droite ou à gauche à l'intersection d'un labyrinthe en T continu pour obtenir une récompense. Trois protocoles inspirés de cette tâche ont été testés, dans lesquels chacun des deux « buts » (à l'extrémité des branches du T) était associé à un type de récompense particulier (du lait concentré sucré, dilué dans l'eau, parfumé à la framboise ou à la vanille). Les rats étaient entraînés dans cette tâche, puis une procédure de dévaluation de l'une des deux récompenses était appliquée. Aucune des variantes de cette tâche n'a permis d'observer une diminution sélective du nombre de visites au but associé à la récompense dévaluée. Dans la dernière variante, nous avons cependant observé une diminution significative de la vitesse dans le bras central du labyrinthe lors des essais où le rat allait visiter le but dévalué, mais seulement pour deux rats sur six.

En conclusion, ces différents protocoles n'ont pas permis d'obtenir de résultats clairs sur l'ensemble de la population de rats testés concernant l'effet de la dévaluation d'une récompense sur le comportement spatial. Notons qu'en parallèle de ces données dans les tâches spatiales, des tests de préférence de consommation des récompenses dans la cage des animaux étaient systématiquement réalisés pour vérifier l'efficacité de la dévaluation. Ces tests n'ont pas donné de résultats statistiquement significatifs démontrant une diminution sélective de la préférence pour la récompense dévaluée, mais une tendance existait toujours dans ce sens. Pour une tâche d'enregistrement électrophysiologique, il est essentiel que le comportement soit bien maîtrisé et il semble que l'effet de la dévaluation par satiété soit sujet à trop de variabilité interindividuelle et dépende de trop nombreux paramètres difficilement maîtrisables, tels que les préférences des rats pour un type de récompense donné. De plus, comme les tâches ne nécessitaient pas spécialement d'apprentissage du lien lieu-type de récompense, il est possible que cet apprentissage n'ait pas été réalisé par les rats. Par ailleurs, les paramètres que nous observions (ordre des visites dans le labyrinthe radial, nombre de visites ou vitesse dans la tâche d'alternance) n'étaient peut-être pas adaptés à l'évaluation de la préférence des rats pour un but ou un autre. Les tests réalisés après la dévaluation l'étaient toujours en extinction, c'est-à-dire, en l'absence de récompense, de manière similaire aux protocoles de dévaluation en apprentissage instrumental. L'absence de récompense a

pu perturber le comportement de navigation, les animaux entrant alors dans un comportement d'exploration. Enfin, peut-être que ce type de tâche de navigation requiert un comportement plus flexible que les tâches classiques d'apprentissage par renforcement (où le rat doit appuyer sur un levier) et donc plus sensible à une baisse générale de motivation générée par la satiété. Pour toutes ces raisons, nous avons changé d'approche ainsi que de paradigme expérimental, toujours dans l'objectif d'étudier la représentation du but des cellules de lieu.

Tâche de navigation continue à deux buts : protocole

Une nouvelle tâche a alors été élaborée, sur la base de la tâche de navigation continue dans laquelle l'activité au but extra-champ des cellules de lieu avait été mise en évidence (Hok et al., 2007a). Cette nouvelle tâche a été intitulée « tâche de navigation à deux buts » (cf. Fig. 73, p. 120). Par rapport au protocole originel, deux modifications majeures ont été implémentées, ajoutant une composante explicite de prise de décision à cette tâche spatiale : tout d'abord, deux zones-buts étaient disponibles en même temps et le rat pouvait choisir librement d'activer l'un ou l'autre de ces buts; par ailleurs, la quantité de récompense délivrée lors de l'activation de chaque zone-but pouvait être modifiée, modulant ainsi la valeur du but correspondant. Les deux zones-buts étaient non-indicées et positionnées à égale distance d'une carte blanche peinte sur la paroi d'une arène circulaire (cf. Fig. 74 p. 121 et Fig. 75 p. 122). Trois types de conditions étaient possibles :

- i. dans la condition de **référence**, les deux buts étaient associés à la libération d'un granulé ;
- ii. dans la condition **d'extinction**, l'un des deux buts ne fournissait plus de granulé ; la valeur de ce but était alors diminuée ;
- iii. dans la condition « **valeur élevée** », l'un des deux buts fournissait 3 granulés simultanément ; la valeur de ce but était augmentée.

Au cours de la tâche, des sessions de 16 minutes en condition de référence alternaient avec des sessions de même durée dans lesquelles la valeur d'un but était modifiée. Deux types de **séquences de session** étaient possibles (cf. Fig. 79 p. 126 et Fig. 80 p. 126):

- i. la **séquence d'extinction**, composée de deux sessions d'extinction intercalées entre deux sessions de référence.
- ii. la **séquence « valeur élevée »**, où deux sessions de référence alternaient avec deux sessions à valeur élevée.

Dans les deux cas, la position du but dont la valeur était modifiée était contrebalancée au sein d'une séquence. La deuxième session de référence servait de session-tampon pour réinitialiser la préférence suite à une session à valeur modifiée.

L'expérience dans sa globalité peut être divisée en plusieurs étapes :

- i. entraînement des rats dans la condition de référence de la tâche (entraînement progressif d'environ 6 semaines où les deux buts étaient d'abord appris séparément, puis ensemble, en utilisant dans un premier temps un indice au sol indiquant leur position) ;
- ii. implantation des rats avec un plot de quatre tétrodes d'enregistrement (fabriqué dans le laboratoire) ;
- iii. réentraînement après une période de récupération post-chirurgie ;
- iv. passage des rats dans les deux séquences précédemment citées lorsque des signaux neuronaux d'intérêt étaient captés sur les électrodes d'enregistrement.

Dix rats ont été entraînés dans la tâche et 7 d'entre eux ont été implantés dans l'hippocampe dorsal droit (5 au-dessus du champ CA1, 2 au-dessus du champ CA3). Dans le cas où aucun signal d'intérêt n'était détecté, les rats étaient simplement entraînés dans la condition de référence pendant 16 minutes, puis les électrodes étaient légèrement descendues. Lorsque des signaux pouvant potentiellement être émis par des neurones hippocampiques étaient captés, le rat était soumis à une séquence de sessions (extinction ou valeur élevée) de 4 fois 16 minutes. Le lendemain, si les signaux étaient retrouvés, l'autre séquence de sessions était réalisée. Puis les électrodes étaient légèrement descendues pour capter de nouveaux signaux. Cette phase d'enregistrement a duré environ 7 mois, jusqu'à atteindre un avancement des électrodes indiquant que nous avons très probablement dépassé non seulement la couche CA1 mais aussi la couche CA3 de l'hippocampe. Les 7 rats ont alors été tués par injection d'une dose létale d'anesthésiant et leurs cerveaux ont été extraits et congelés. Des coupes coronales de cerveau proches de la zone d'implantation ont été prélevées et déposées sur lamelle, puis colorées au crésyl violet pour vérifier au microscope optique la zone d'implantation des électrodes. Cette analyse histologique, couplée à la localisation verticale de l'enregistrement des neurones (présentée en annexe IV, p. 204) a permis d'estimer la zone de provenance des neurones enregistrés (CA1 ou CA3). Les autres données récoltées au cours de la tâche étaient les données de position de l'animal au cours du temps, l'activité électrophysiologique digitalisée au cours du temps (temps de chaque potentiel d'action) et les marqueurs (*event-flags*) indiquant soit l'activation d'un but dans une condition de récompense donnée, soit la consommation de récompense (qui, rappelons-le, avait lieu à un endroit aléatoire dans l'environnement, et était donc dissociée du but). Une étape de prétraitement des données était nécessaire avant l'analyse proprement dite : il s'agissait du « *spike-sorting* », au cours de laquelle l'expérimentateur utilisait un logiciel commercial pour déterminer l'origine des signaux enregistrés (qui pouvait aussi être du bruit électronique, ou de l'activité non-dissociable provenant de plusieurs neurones, auxquels cas les signaux n'étaient pas pris en compte). Seules les sessions pour lesquelles nous avons enregistré au moins un neurone ont été prises en compte pour les analyses, qu'elles soient comportementales ou électrophysiologiques.

Tâche de navigation continue à deux buts : résultats comportementaux

Dans un premier temps, les **données comportementales** (position et marqueurs) ont été analysées pour évaluer la bonne compréhension de la tâche par les rats ainsi que l'adaptation de leur comportement aux modifications de valeur des buts. Cette première série d'analyses a permis de montrer que les rats étaient capables d'apprendre à localiser les deux buts non-indicés et d'adapter de manière flexible leur comportement en fonction de la quantité de récompense attendue. Notamment, le taux moyen **d'activations de but** par minute (c'est-à-dire, le nombre de fois où le rat était resté deux secondes dans l'une ou l'autre des zones-buts) était d'environ 4.5 dans la condition de référence (cf. Fig. 84, p. 135). Nous avons calculé un **indice de préférence comportemental** permettant d'évaluer la préférence relative pour l'un ou l'autre des buts au sein d'une session donnée. Cet indice pouvait indifféremment évaluer la **préférence spatiale** (positif si le rat préférait visiter le but de gauche par rapport à la carte, négatif s'il préférait visiter celui de droite, proche de zéro en absence de préférence) ou la **préférence de valeur** (uniquement pour les sessions à valeur modifiée : positif si le rat préférait le but à la valeur la plus haute, négatif s'il préférait celui ayant la valeur la plus basse, proche de zéro en absence de préférence). Cet indice était calculé sur la base du

nombre d'activations de chaque but. L'analyse du taux de préférence spatial (cf. Fig. 85, p. 136) et du taux de préférence de valeur (cf. Fig. 86, p. 136) a permis de conclure que les rats répartissaient bien leurs visites dans les deux buts dans la condition de référence (malgré un léger biais pour le but de droite) et qu'ils adaptaient bien leur comportement aux modifications de valeur de buts, en préférant significativement le but fournissant un granulé dans la condition d'extinction et le but fournissant trois granulés dans la condition à valeur élevée. Nous avons aussi observé que l'évolution de la préférence semblait se faire graduellement au cours de la session, de manière beaucoup plus rapide pour la condition d'extinction que pour la condition de valeur élevée (cf. Fig. 89, p. 140 et Fig. 91, p. 142). Ce schéma d'évolution se retrouvait à l'échelle individuelle pour chaque rat, bien que l'un d'entre eux montrait une préférence très marquée pour le but de droite dans la condition de référence. En conclusion, ces résultats dans la tâche de navigation à deux buts démontrent à la fois que les rats sont capables de **mémoriser la localisation de deux zones-buts** non-indicées mais aussi **d'adapter leur comportement aux variations de valeur de ces buts**. La mesure employée (indice de préférence) a permis d'évaluer précisément la préférence comportementale des rats pour l'un ou l'autre des buts au sein d'une même session et nous avons estimé qu'elle était un marqueur approprié de la valeur que les rats associent à chacun des buts.

Tâche de navigation continue à deux buts : activité liée au but dans la condition de référence

Dans une deuxième série d'analyses, les résultats électrophysiologiques recueillis dans la tâche de navigation à deux buts, en condition de référence, ont été analysés. Nous avons enregistré au total 194 neurones hippocampiques, donc 144 ont été catégorisés par l'expérimentateur comme étant des cellules pyramidales et 15 comme interneurons. Parmi les cellules pyramidales, 28% (41) étaient des **cellules dites « silencieuses »**, c'est-à-dire des cellules n'ayant pas de champ de lieu visible dans l'environnement (cf. Fig. 98, p. 157), contrairement au reste des neurones pyramidaux qui ont été catégorisés comme étant des **cellules de lieu** (cf. Fig. 97, p. 157). Les cellules silencieuses étaient caractérisées par un taux de décharge moyen très bas (0.09 Hz) comparé à celui des cellules de lieu (0.92 Hz). Leur activité sporadique semblait se concentrer à l'endroit des buts. Dans un premier temps, nous nous sommes concentrés sur la position du champ de lieu des cellules. En effet, dans la littérature, certaines études montrent une concentration des champs d'activité des cellules de lieu à l'endroit du but (Hollup et al., 2001b ; Kobayashi et al., 2003) alors que ce n'est pas le cas dans d'autres études (Lenck-Santini et al., 2001, 2002 ; Hok et al., 2007a). Les résultats que nous avons obtenus ont montré que, dans la tâche de navigation à deux buts, **la répartition des champs de lieux ne différait pas d'une répartition homogène** (cf. Fig. 99, p. 159). Nous avons noté une tendance, pour les champs de CA3, à être plus concentrés autour des buts, mais ce point particulier requiert des analyses supplémentaires. L'une des questions majeures de cette étude était de savoir si l'activité au but allait être présente de manière similaire pour les deux zones-buts, ce qui serait en faveur de l'hypothèse motivationnelle du signal au but, ou s'il y aurait une différence de « codage » spatial des deux buts. Pour répondre à cette question, nous avons analysé l'activité liée au but de chaque cellule dont le champ principal n'empiétait pas sur les buts. Nous avons comparé le taux de décharge hors-champ au but (lors des 2 s de délai pendant lesquelles le rat attend la libération de récompense) au taux de décharge hors-champ et hors-but. Les résultats montrent, d'une part, que **l'activité hors-champ de la majorité des cellules pyramidales (72 %) était significativement plus**

importante au but qu'en-dehors du but. D'autre part, pour toutes les populations de cellules étudiées (CA1, CA3, cellules de lieu et cellules silencieuses), la majorité codait pour l'un des deux buts et une minorité était significativement active aux deux zones-buts (cf. Fig. 100, p. 160). L'activité au but contient donc un **aspect spatial pouvant permettre de discriminer les deux zones buts**. De manière intéressante, les proportions de cellules qui codaient pour un but, les deux buts, ou aucun des deux étaient similaires entre ces différentes populations de cellules. Dans des études précédentes, l'activité au but n'avait été mise en évidence que dans le champ CA1 (Hok et al., 2007a,b ; Hok et al., 2013). Nous avons donc montré que les **cellules du champ CA3 et les cellules silencieuses émettaient aussi une activité significative au but**.

Dans une autre série d'analyses de la condition de référence, nous nous sommes penchés sur le profil de décharge aux buts, en ne prenant en compte que les cellules qui codaient significativement pour au moins l'une des deux zones-buts. Nous avons calculé l'histogramme cumulé pour l'activité liée au but ainsi que pour l'activité hors-champ liée à la consommation de récompense (cf. Fig. 101, p. 162). **L'activité liée au but présentait un profil de décharge caractéristique**, avec une augmentation de la décharge ayant lieu environ 500 ms après l'entrée dans la zone-but et une diminution brusque au moment du déclenchement du distributeur. Ce profil, similaire pour les deux zones-buts, était complètement absent pour les deux secondes précédant la consommation de la récompense. Il était présent pour les différentes sous-catégories de cellules, bien que semblant moins marqué pour les cellules de lieu de CA3 (cf. Fig. 102, p. 163).

Tâche de navigation continue à deux buts : activité liée au but et modifications de valeur des buts

Nous avons ensuite analysé l'activité au but dans les conditions à récompense variable, afin d'évaluer si l'activité au but encodait la valeur du but, ce qui serait en faveur d'une origine motivationnelle et d'anticipation de la récompense de cette activité. Pour ce faire, nous avons utilisé l'indice de préférence de décharge, similaire à l'indice de préférence comportemental, qui permet d'évaluer la préférence relative de décharge de l'activité au but (cf. méthodes, Sec. 8.1.6.1, p. 154, et Fig. 106, p. 169). Comme précédemment, cet indice pouvait s'appliquer à une préférence spatiale ou à une préférence pour la valeur du but. La distribution des indices de préférence pour la valeur n'était pas différente d'une absence de préférence, indiquant une **absence de codage de la valeur** dans l'activité de population liée au but. De plus, la variabilité de cette distribution était beaucoup plus faible que celle de la distribution des indices de préférence de décharge spatiale, en cohérence avec le résultat précédant indiquant qu'une majorité de la population des cellules codaient significativement pour seulement l'un des deux buts. Notons que les cellules prises en compte pour les analyses de taux de préférence étaient celles enregistrées lors de sessions où le rat a visité au moins 10 fois chaque but (pour s'assurer d'avoir un échantillonnage suffisant des deux zones-buts). La décharge des **interneurones** pendant le délai d'attente au but ne codait pas non plus pour la valeur du but. L'étendue de la distribution des indices de préférences spatiaux autour de 0 étant très faible, il ne semblait pas que les interneurones différencient les deux buts. En revanche, une diminution de la décharge au moment de l'entrée dans la zone-but était visible, de manière comparable à l'augmentation de décharge des cellules pyramidales (cf. Fig. 109, p. 170). Cependant, au moment du déclenchement du distributeur, la décharge ne revenait pas brusquement au niveau de base, contrairement au phénomène observé pour les cellules pyramidales. Il semble aussi que la

décharge des interneurons diminue brusquement au moment de consommation de la récompense, contrairement à ce qu'il se passe pour les cellules de lieu. L'analyse de la décharge interneuronale, même si elle a été réalisée sur un petit nombre de neurones ($n = 13$), permet de dire que ces cellules inhibitrices participent peut-être à l'augmentation de décharge liée à l'entrée dans la zone-but mais que **d'autres phénomènes contribuent très probablement à cette décharge**.

Pour terminer, nous nous sommes penchés sur l'évolution de la préférence de décharge pour la valeur au sein de la session. En effet, les résultats comportementaux montrent une augmentation significative de la préférence pour la zone-but la plus récompensée entre les huit premières et huit dernières minutes des sessions. Il serait donc possible que la préférence de décharge pour la valeur ne soit visible qu'au début ou qu'à la fin des sessions, démontrant soit un signal de mise à jour de la valeur, soit un signal d'attente de la récompense. Les analyses n'ont pas mis en évidence d'évolution significative de la préférence de décharge pour la valeur au cours des sessions (cf. Fig. 111, p. 172). Nous pouvons donc en conclure que, contrairement au comportement des rats, **la décharge liée au but est indépendante de la valeur des buts**.

Activité liée au but : conclusion et discussion

En résumé, les résultats obtenus à la fois dans la condition de référence (où les deux buts avaient la même valeur) et dans les conditions à valeur modifiée tendent vers la même conclusion : **l'activité de population liée au but, émise par une large proportion de cellules pyramidales, détient un aspect spatial mais n'encode pas la valeur du but**. Nous notons cependant qu'une analyse au niveau de chaque neurone serait nécessaire pour préciser cette conclusion au niveau individuel. En effet, une étude récente a montré que l'activité du champ principal d'une sous-population de cellules de lieu encodait certains aspects liés à la valeur de l'action ou de la récompense dans une tâche d'alternance spontanée où la probabilité de récompense variait (Lee, Ghim et al., 2012). Il serait intéressant de voir si l'activité au but de certaines cellules encode la valeur du but malgré l'absence de codage de valeur au niveau populationnel. Des analyses préliminaires vont dans le sens d'une absence de codage de la valeur au niveau individuel (résultats non-présentés ici). Dans l'ensemble, nos résultats vont dans le sens d'une étude précédente réalisée dans un labyrinthe en croix (Tabuchi et al., 2003) où l'activité des cellules pyramidales n'encodait pas la magnitude de la récompense attendue.

Le fait que l'activité au but soit aussi exprimée par les neurones pyramidaux de CA3 (cellules de lieu et cellules silencieuses) indique que **ce signal n'est probablement pas dû aux afférences hippocampiques sous-corticales** (telles que l'aire tegmentale ventrale ou l'amygdale, qui projettent vers CA1 mais beaucoup moins vers CA3 ; Gasbarri et al., 1994 ; Pitkänen et al., 2000) mais peut-être plutôt aux entrées provenant du cortex entorhinal (qui innerve ces deux champs hippocampiques). Des processus internes à l'hippocampe sont sûrement aussi à l'œuvre et pourraient probablement impliquer l'hippocampe ventral, qui est plus fortement lié aux structures sous-corticales citées ci-dessus (Strange et al., 2014). Notons cependant qu'une origine dopaminergique liée à la substance noire ne peut pas être écartée. L'enregistrement des cellules de grilles (d'autres cellules spatio-sélectives) du cortex entorhinal et de cellules de lieu dans l'hippocampe ventral, dans une tâche similaire à celle que nous avons utilisée, pourrait permettre d'apporter des arguments majeurs concernant l'origine du signal lié au but.

Concernant le rôle de cette activité, puisqu'elle ne semble pas encoder l'attente de la récompense, nos résultats vont dans le sens d'un **rôle spatial de l'activité liée au but**. Notamment, ce signal pourrait signaler l'arrivée à un endroit visé. Lors de la navigation, l'activité exprimée par le réseau de neurones de CA3 pourrait être comparée à l'activité dans CA1, et une augmentation de la synchronisation entre ces deux champs hippocampiques à l'entrée dans la zone-but pourrait transmettre l'information au cortex préfrontal médian (qui exprime une activité liée au but dépendante de celle de l'hippocampe ; Hok et al., 2005 ; Burton et al., 2009). Cette activité synchronisée fournirait l'aspect spatial de la représentation du but et pourrait contribuer à la prise de décision d'arrêter la navigation. Des enregistrements simultanés dans ces différentes zones pourraient permettre de confirmer ou d'infirmer cette hypothèse, en portant une attention particulière aux constantes de temps avec lesquelles cette activité est émise ainsi qu'à la synchronisation entre potentiels d'actions de l'activité au but dans ces aires.

En conclusion, nos travaux de thèse ont permis de mettre au point une tâche combinant navigation spatiale et prise de décision, dans laquelle la valeur que les rats attribuent à deux buts spatiaux peut être estimée par l'intermédiaire de leurs préférences comportementales. L'enregistrement de neurones de l'hippocampe dorsal, dans cette tâche, a permis de clarifier le rôle de l'hippocampe dans la représentation du but. En effet, nous avons montré que, même dans le cas de deux buts spatiaux accessibles simultanément, la représentation de l'espace par le champ d'activité des cellules de lieu reste homogène. En revanche, nous avons mis en évidence une activité hors-champ, liée au but, émise par la grande majorité de la population de cellules pyramidales de l'hippocampe dorsal, que ce soit dans les champs CA1 ou CA3. De manière surprenante, cette activité est aussi émise par la population de cellules silencieuses de l'hippocampe, qui n'expriment pas de champ spatial dans l'environnement. Cette activité liée au but possède les caractéristiques d'une représentation du but, surtout visible dans la population de cellules silencieuses, pour lesquelles l'attente au but est le facteur déterminant de la décharge. Cette représentation du but comprend un aspect spatial dans le sens où elle est, en majorité, significative pour l'un mais pas l'autre des buts. En revanche, au niveau de la population, elle est indépendante de la valeur du but.

Dans leur ensemble, ces résultats permettent donc de confirmer l'hypothèse d'une nature spatiale du signal au but et d'infirmer l'hypothèse d'une nature motivationnelle. Cette conclusion est en accord avec le rôle généralement attribué à l'hippocampe, celui d'une structure de traitement de l'information spatiale. La combinaison de l'aspect spatial de la représentation du but, fournie par l'hippocampe, avec des aspects de valeur, présents dans d'autres structures tels que l'amygdale ou le cortex orbitofrontal, pourrait être réalisée en aval du flux informationnel, dans le cortex préfrontal médian ou le striatum ventral, pour permettre de choisir le but vers lequel naviguer. Une fois la navigation initiée, l'activité liée au but des neurones pyramidaux pourrait être utilisée pour confirmer l'arrivée à la position souhaitée. Ces résultats contribuent à préciser le rôle de l'hippocampe dans la représentation du but et pourront être utilisés notamment dans les modèles neuromimétiques de l'hippocampe. Ils participent aussi à une meilleure compréhension du rôle de l'hippocampe, à l'interface entre mémoire spatiale et prise de décision. Sur le long terme, cette approche peut permettre de mieux délimiter les cibles thérapeutiques dans le traitement de certains troubles neuropsychologiques, tels que l'addiction, la dépression ou la schizophrénie, dans lesquelles l'hippocampe joue un rôle encore méconnu à l'heure actuelle.

Appendices

I. Parameters of devaluation experiments.....	200
II. Behavioural training of the two-goal navigation task.....	202
III. Cumulated time maps for all cued sessions	203
IV. Recorded unit types as a function of electrode depth.....	204
V. Side and value preference index in subpopulations of hippocampal pyramidal cells.	205

I. Parameters of devaluation experiments

Table 8 presents the summary of material, methods and results obtained from four experiments in the eight-arm radial maze. In none of them did the devaluation consistently affect rats' behaviour in the task.





Tasks	Rewards	Subjects	Reward Allocations	Devaluation	Post-devaluation test (in the task)	Remarks	Effect of devaluation in the task	Effect of devaluation on consumption
Exp. 1 Radial Maze	- A : "purified" pellets, 20mg - B : "grain-based" pellets, 20mg	10 Long-Evans rats		10 g for 1 hour	Extinction		No effect on the order of visits.	Tendency to prefer the non-devalued food
Exp. 2 Radial Maze	- A : "purified" pellets, 3*20mg - B : "grain-based" pellets, 3*20 mg - C : "neutral" pellets, 3 * 20 mg	10 L-E rats (same as Exp. 1)		<i>Ad libitum</i> for 1 hour	Extinction or rewarded		No effect on the order of visits.	Tendency to prefer the non-devalued food
Exp. 3 Radial Maze with obstacles	- A : "purified" pellets, 3 * 20mg - B : "grain-based" pellets, 3 * 20 mg	10 L-E rats (same as Exp. 1)		<i>Ad libitum</i> for 1 hour	Extinction or rewarded	Higher criteria to proceed to devaluation	No effect on the order of visits.	Tendency to prefer the non-devalued food
Exp. 4 Radial Maze with obstacles	- D : sucrose pellets (peanut butter flavour), 125 mg - E : sucrose pellets (fruit flavour), 125 mg	10 L-E rats (same as Exp. 1)		<i>Ad libitum</i> for 1 hour	Extinction	The rat is removed after 4 visits for 2 out of 4 trials	No effect on the order of visits.	Tendency to prefer the non-devalued food

Table 8: Parameters and results of devaluation experiments in the radial maze.

Table 9 presents the summary of material, methods and results obtained from three experiments in the continuous alternation maze. In none of them did the devaluation consistently affect rats' behaviour in the task.





Tasks	Rewards	Subjects	Reward distribution	Devaluation	Post-devaluation test (in the task)	Remarks	Effect of devaluation in the task	Effect of devaluation on consumption
Exp. 5 Continuous alternation maze	- F : sweetened condensed milk diluted (10%) in water and raspberry-flavoured - I : sweetened condensed milk diluted (10%) in water and vanilla-flavoured 0.1 mL per passage.	6 L-E rats (same as exp. 1)		Ad libitum for 1 hour	Experiment stopped before the test (the criteria was never reached)	2 laps for a reward.	Not tested	Not tested
Exp. 6 Continuous alternation maze	Same as Exp. 5	4 L-E rats (same as exp. 1 and different from exp. 5)		Ad libitum for 1 hour	In extinction for 10 min, then rewarded for 10 min	The rat was forced to alternate	No effect on the number of visits.	No decrease of the devalued reward consumption
Exp. 7 Continuous alternation maze	Same as Exp. 5	6 Long-Evans rats (naive)		Ad libitum for 1 hour	Extinction for 8 alternations, then rewarded.	Alternation and habituation to the extinction.	No effect on the number of visits. Specific decrease of speed for two out of six rats.	No decrease of the devalued reward consumption except for 2 out of 6 rats.

Table 9: Parameters and results of devaluation experiments in the continuous alternation maze.

II. Behavioural training of the two-goal navigation task


The following table details the behavioural training steps for the 2-goal continuous navigation task (before surgery). Two to three weeks after surgery, the rats were trained again. The **retraining** protocol consisted of one week of reference condition (two 16-min session with the two uncued goals), then one week of balancing protocol with the use of the cues for half of the sessions.

	Familiarisation	One-goal training	Two-goal training	Two-goal training	Testing protocol	Balancing protocol
 Time (minutes in a day)	10 min	16 min 1 cued goal 0-2s delay	16 min 2 cued goals	16 min 2 uncued goals	16 min reference (1vs1)	8 min reference (1vs1)
	10 min random foraging	16 min 1 uncued goal 0-2s delay	16 min 2 uncued goals	16 min 2 uncued goals	16 min high (1vs3)	8 min reference or extinction (1vs0)
					16 min extinction (1vs0)	8 min reference
				8 min reference or extinction		
Duration in days or criterion for next step	1 day	2 visits /min	1 visit /goal/min	7 days	5 days	1 visit /goal/min & no significant side preference in the last 2 reference sessions
 Time (days)						
Remarks		The delay is increased by 0.5 s steps from 0 to 2 s whenever criterion is reached. The active goal is changed each day.	2 second delay from now on	2 uncued goals from now on	The value-changing goals are chosen so as to even out the preference	Setting to 'extinction' the preferred goal for rats showing a significant preference (binomial test)

Note that the rat was temporarily removed from the arena whenever the cues were removed.


III. Cumulated time maps for all cued sessions

The occupancy (time) maps are presented for all cued sessions, grouped by condition.




IV. Recorded unit types as a function of electrode depth

The approximate number of neurons, as a function of the estimated depth of the electrodes in the brain, is shown by red circles (the larger the circle, the more numerous the neurons). Four types of neurons were discriminated: PC (place cells), GA (silent cells), IN (interneurons), Others. This was used in conjunction with histological brain slices to evaluate the possible layer (CA1, CA3) of origin of each unit.



V. Side and value preference index in subpopulations of hippocampal pyramidal cells.

These figures indicate the side and value firing preference index for various populations of hippocampal pyramidal cells. The **side preference index** was computed for reference sessions (in blue) and the **value preference index** for value-changing sessions (red or green). The small grey circles present single cell values. The samples were compared to a 0-mean population that indicates no preference. Only the side preference index from the ref (ext) condition in CA1 significantly differed from a no-preference condition. The data from CA1 and CA3 is presented in the main body of the manuscript (Fig. 108, p. 169).



References

-A-

- Aben, B., Stapert, S., Blokland, A., **2012**. About the distinction between working memory and short-term memory. *Frontiers in Psychology*, 3, 301.
- Abraham, W.C., Dragunow, M., Tate, W.P., **1991**. The role of immediate early genes in the stabilization of long-term potentiation. *Molecular Neurobiology*, 5, 297–314.
- Abraham, W.C., Logan, B., Greenwood, J.M., Dragunow, M., **2002**. Induction and Experience-Dependent Consolidation of Stable Long-Term Potentiation Lasting Months in the Hippocampus. *The Journal of Neuroscience*, 22, 9626–9634.
- Adams, C.D., Dickinson, A., **1981**. Instrumental responding following reinforcer devaluation. *The Quarterly Journal of Experimental Psychology Section B: Comparative and Physiological Psychology*, 33, 109.
- Aikath, D., Weible, A.P., Rowland, D.C., Kentros, C.G., **2014**. Role of self-generated odor cues in contextual representation. *Hippocampus*, 24, 1039–1051.
- Ainge, J.A., van der Meer, M.A.A., Langston, R.F., Wood, E.R., **2007**. Exploring the role of context-dependent hippocampal activity in spatial alternation behavior. *Hippocampus*, 17, 988–1002.
- Alcaraz, F., Naneix, F., Desfosses, E., Marchand, A.R., Wolff, M., Coutureau, E., **2014**. Dissociable effects of anterior and mediodorsal thalamic lesions on spatial goal-directed behavior. *Brain Structure and Function*, 1–11.
- Alvarez, P., Squire, L.R., **1994**. Memory consolidation and the medial temporal lobe: a simple network model. *Proceedings of the National Academy of Sciences*, 91, 7041–7045.
- Alvernhe, A., **2010**. Codage de la topologie de l'espace dans l'hippocampe: approches comportementale et électrophysiologique. Aix-Marseille 1. Doctoral thesis.
- Alvernhe, A., Sargolini, F., Poucet, B., **2012**. Rats build and update topological representations through exploration. *Animal Cognition*, 15, 359–368.
- Alvernhe, A., Save, E., Poucet, B., **2011**. Local remapping of place cell firing in the Tolman detour task. *European Journal of Neuroscience*, 33, 1696–1705.
- Alvernhe, A., Van Cauter, T., Save, E., Poucet, B., **2008**. Different CA1 and CA3 Representations of Novel Routes in a Shortcut Situation. *The Journal of Neuroscience*, 28, 7324–7333.
- Alyan, S., McNaughton, B.L., **1999**. Hippocampectomized rats are capable of homing by path integration. *Behavioral Neuroscience*, 113, 19–31.
- Amaral, D.G., Witter, M.P., **1989**. The three-dimensional organization of the hippocampal formation: A review of anatomical data. *Neuroscience*, 31, 571–591.
- Anderson, M.I., Jeffery, K.J., **2003**. Heterogeneous Modulation of Place Cell Firing by Changes in Context. *The Journal of Neuroscience*, 23, 8827–8835.
- Anderson, M.I., Killing, S., Morris, C., O'Donoghue, A., Onyiahga, D., Stevenson, R., Verriotis, M., Jeffery, K.J., **2006**. Behavioral correlates of the distributed coding of spatial context. *Hippocampus*, 16, 730–742.

- Araque, A., Parpura, V., Sanzgiri, R.P., Haydon, P.G., **1999**. Tripartite synapses: glia, the unacknowledged partner. *Trends in Neurosciences*, 22, 208–215.
- Arleo, A., Déjean, C., Allegraud, P., Khamassi, M., Zugaro, M.B., Wiener, S.I., **2013**. Optic flow stimuli update anterodorsal thalamus head direction neuronal activity in rats. *The Journal of Neuroscience*, 33, 16790–16795.
- Arleo, A., Rondi-Reig, L., **2007**. Multimodal sensory integration and concurrent navigation strategies for spatial cognition in real and artificial organisms. *International Journal of Neuroscience*, 06, 327.
- Atkinson, R.C., Shiffrin, R.M., **1968**. Human Memory: A Proposed System and its Control Processes, in: Kenneth W. Spence and Janet Taylor Spence (Eds.), *Psychology of Learning and Motivation*. Academic Press, pp. 89–195.
- Avni, R., Zadicario, P., Eilam, D., **2006**. Exploration in a dark open field: a shift from directional to positional progression and a proposed model of acquiring spatial information. *Behavioural Brain Research*, 171, 313–323.

-B-

- Babayan, B.M., **2014**. Unraveling the neural circuitry of sequence-based navigation using a combined fos imaging and computational approach. Paris Descartes. Doctoral Thesis.
- Babb, S.J., Crystal, J.D., **2006a**. Episodic-like Memory in the Rat. *Current Biology*, 16, 1317–1321.
- Babb, S.J., Crystal, J.D., **2006b**. Discrimination of what, when, and where is not based on time of day. *Learning & Behavior*, 34, 124–130.
- Baddeley, A., **2000**. The episodic buffer: a new component of working memory? *Trends in Cognitive Sciences*, 4, 417–423.
- Baddeley, A.D., Hitch, G., **1974**. Working Memory, in: Gordon H. Bower (Ed.), *Psychology of Learning and Motivation*. Academic Press, pp. 47–89.
- Baeg, E.H., Kim, Y.B., Huh, K., Mook-Jung, I., Kim, H.T., Jung, M.W., **2003**. Dynamics of Population Code for Working Memory in the Prefrontal Cortex. *Neuron*, 40, 177–188.
- Bahar, A.S., Shirvalkar, P.R., Shapiro, M.L., **2011**. Memory-Guided Learning: CA1 and CA3 Neuronal Ensembles Differentially Encode the Commonalities and Differences between Situations. *The Journal of Neuroscience*, 31, 12270–12281.
- Baker, J.G., Franken, R.E., **1968**. Spontaneous alternation and the influence of reward. *Psychological Reports*, 22, 821–825.
- Balleine, B., Dickinson, A., **1991**. Instrumental performance following reinforcer devaluation depends upon incentive learning. *The Quarterly Journal of Experimental Psychology Section B: Comparative and Physiological Psychology*, 43, 279.
- Balleine, B.W., Dickinson, A., **1998**. Goal-directed instrumental action: contingency and incentive learning and their cortical substrates. *Neuropharmacology*, 37, 407–419.
- Balleine, B.W., Killcross, S., **2006**. Parallel incentive processing: an integrated view of amygdala function. *Trends in Neurosciences*, 29, 272–279.
- Balleine, B.W., O’Doherty, J.P., **2010**. Human and rodent homologies in action control: corticostriatal determinants of goal-directed and habitual action. *Neuropsychopharmacology*, 35, 48–69.

- Bannerman, D.M., Rawlins, J.N.P., McHugh, S.B., Deacon, R.M.J., Yee, B.K., Bast, T., Zhang, W.-N., Pothuizen, H.H.J., Feldon, J., **2004**. Regional dissociations within the hippocampus—memory and anxiety. *Neuroscience & Biobehavioral Reviews*, 28, 273–283.
- Barnes, C.A., McNaughton, B.L., Mizumori, S.J., Leonard, B.W., Lin, L.H., **1990**. Comparison of spatial and temporal characteristics of neuronal activity in sequential stages of hippocampal processing. *Progress in Brain Research*, 83, 287–300.
- Barnes, C.A., Nadel, L., Honig, W.K., **1980**. Spatial memory deficit in senescent rats. *Canadian Journal of Psychology*, 34, 29–39.
- Barry, C., Ginzberg, L.L., O’Keefe, J., Burgess, N., **2012**. Grid cell firing patterns signal environmental novelty by expansion. *Proceedings of the National Academy of Sciences*, 109, 17687–17692.
- Barry, C., Hayman, R., Burgess, N., Jeffery, K.J., **2007**. Experience-dependent rescaling of entorhinal grids. *Nature Neuroscience*, 10, 682–684.
- Barry, C., Lever, C., Hayman, R., Hartley, T., Burton, S., O’Keefe, J., Jeffery, K., Burgess, N., **2006**. The boundary vector cell model of place cell firing and spatial memory. *Reviews in the Neurosciences*, 17, 71–97.
- Bast, T., Wilson, I.A., Witter, M.P., Morris, R.G.M., **2009**. From Rapid Place Learning to Behavioral Performance: A Key Role for the Intermediate Hippocampus. *PLoS Biology*, 7, e1000089.
- Battaglia, F.P., Benchenane, K., Sirota, A., Pennartz, C.M.A., Wiener, S.I., **2011**. The hippocampus: hub of brain network communication for memory. *Trends in Cognitive Sciences (Regul. Ed.)* 15, 310–318.
- Bauer, E.P., LeDoux, J.E., Nader, K., **2001**. Fear conditioning and LTP in the lateral amygdala are sensitive to the same stimulus contingencies. *Nature Neuroscience*, 4, 687–688.
- Beeler, J.A., Cools, R., Luciana, M., Ostlund, S.B., Petzinger, G., **2014**. A kinder, gentler dopamine... highlighting dopamine’s role in behavioral flexibility. *Frontiers in Neuroscience*, 8, 4.
- Benchenane, K., Peyrache, A., Khamassi, M., Tierney, P.L., Gioanni, Y., Battaglia, F.P., Wiener, S.I., **2010**. Coherent theta oscillations and reorganization of spike timing in the hippocampal-prefrontal network upon learning. *Neuron*, 66, 921–936.
- Benhamou, S., Poucet, B., **1998**. Landmark use by navigating rats (*rattus norvegicus*): contrasting geometric and featural information. *Journal of Comparative Psychology*, 112, 317–322.
- Bennett, A.T., **1996**. Do animals have cognitive maps? *The Journal of Experimental Biology*, 199, 219–224.
- Berger-Tal, O., Nathan, J., Meron, E., Saltz, D., **2014**. The Exploration-Exploitation Dilemma: A Multidisciplinary Framework. *PLoS One*, 9, e95693.
- Berthoz, A., Viaud-Delmon, I., **1999**. Multisensory integration in spatial orientation. *Current Opinion in Neurobiology*, 9, 708–712.
- Bethus, I., Tse, D., Morris, R.G.M., **2010**. Dopamine and Memory: Modulation of the Persistence of Memory for Novel Hippocampal NMDA Receptor-Dependent Paired Associates. *The Journal of Neuroscience*, 30, 1610–1618.
- Biegler, R., **1996**. Landmark Stability: Further Studies Pointing to a Role in Spatial Learning. *The Quarterly Journal of Experimental Psychology Section B*, 49, 307–345.
- Biegler, R., Morris, R.G.M., **1993**. Landmark stability is a prerequisite for spatial but not discrimination learning. *Nature*, 361, 631–633.

- Bi, G., Poo, M., **1998**. Synaptic Modifications in Cultured Hippocampal Neurons: Dependence on Spike Timing, Synaptic Strength, and Postsynaptic Cell Type. *The Journal of Neuroscience*, *18*, 10464–10472.
- Bird, L.R., Roberts, W.A., Abroms, B., Kit, K.A., Crupi, C., **2003**. Spatial memory for food hidden by rats (*Rattus norvegicus*) on the radial maze: Studies of memory for where, what, and when. *Journal of Comparative Psychology*, *117*, 176–187.
- Bissonette, G.B., Burton, A.C., Gentry, R.N., Goldstein, B.L., Hearn, T.N., Barnett, B.R., Kashtelyan, V., Roesch, M.R., **2013**. Separate Populations of Neurons in Ventral Striatum Encode Value and Motivation. *PLoS ONE*, *8*, e64673.
- Bjerknes, T.L., Moser, E.I., Moser, M.-B., **2014**. Representation of Geometric Borders in the Developing Rat. *Neuron*, *82*, 71–78.
- Bliss, T.V.P., Lømo, T., **1973**. Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *The Journal of Physiology*, *232*, 331–356.
- Bluck, S., **2003**. Autobiographical memory: Exploring Its functions in everyday life. *Memory*, *11*, 113.
- Blundell, P., Hall, G., Killcross, S., **2001**. Lesions of the Basolateral Amygdala Disrupt Selective Aspects of Reinforcer Representation in Rats. *The Journal of Neuroscience*, *21*, 9018–9026.
- Boccaro, C.N., Sargolini, F., Thoresen, V.H., Solstad, T., Witter, M.P., Moser, E.I., Moser, M.-B., **2010**. Grid cells in pre- and parasubiculum. *Nature Neuroscience*, *13*, 987–994.
- Bostock, E., Muller, R.U., Kubie, J.L., **1991**. Experience-dependent modifications of hippocampal place cell firing. *Hippocampus*, *1*, 193–205.
- Botreau, F., Gisquet-Verrier, P., **2010**. Re-Thinking the Role of the Dorsal Striatum in Egocentric/Response Strategy. *Frontiers in Behavioral Neuroscience*, *4*.
- Bouret, S., Richmond, B.J., Ravel, S., **2012**. Complementary neural correlates of motivation in dopaminergic and noradrenergic neurons of monkeys. *Frontiers in Behavioral Neuroscience*, *6*, 40.
- Bowden, E.M., Jung-Beeman, M., Fleck, J., Kounios, J., **2005**. New approaches to demystifying insight. *Trends in Cognitive Sciences*, *9*, 322–328.
- Bowers, J.S., **2009**. On the biological plausibility of grandmother cells: implications for neural network theories in psychology and neuroscience. *Psychological Review*, *116*, 220–251.
- Brandon, M.P., Koenig, J., Leutgeb, S., **2014**. Parallel and convergent processing in grid cell, head-direction cell, boundary cell, and place cell networks. *Wiley Interdisciplinary Reviews: Cognitive Science*, *5*, 207–219.
- Breukelaar, J., Dalrymple-Alford, J.C., **1998**. Timing ability and numerical competence in rats. *Journal of Experimental Psychology: Animal Behavior Processes*, *24*, 84–97.
- Brun, V.H., Leutgeb, S., Wu, H.-Q., Schwarcz, R., Witter, M.P., Moser, E.I., Moser, M.-B., **2008**. Impaired Spatial Representation in CA1 after Lesion of Direct Input from Entorhinal Cortex. *Neuron*, *57*, 290–302.
- Buhot, M.-C., Chapuis, N., Scardigli, P., Herrmann, T., **1991**. Spatial problem-solving in a wheel-shaped maze: quantitative and qualitative analyses of the behavioural changes following damage to the hippocampus in the rat. *Behavioural Brain Research*, *44*, 67–79.
- Bunsey, M., Eichenbaum, H., **1996**. Conservation of hippocampal memory function in rats and humans. *Nature*, *379*, 255–257.
- Burgess, N., **2006**. Spatial memory: how egocentric and allocentric combine. *Trends in Cognitive Sciences*, *10*, 551–557.

- Burgess, N., **2008**. Spatial Cognition and the Brain. *Annals of the New York Academy of Sciences*, 1124, 77–97.
- Burgess, N., Donnett, J.G., Jeffery, K.J., O'Keefe, J., **1997**. Robotic and neuronal simulation of the hippocampus and rat navigation. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 352, 1535–1543.
- Burton, B.G., Hok, V., Save, E., Poucet, B., **2009**. Lesion of the ventral and intermediate hippocampus abolishes anticipatory activity in the medial prefrontal cortex of the rat. *Behavioural Brain Research*, 199, 222–234.
- Burwell, R.D., Hafeman, D.M., **2003**. Positional firing properties of postrhinal cortex neurons. *Neuroscience*, 119, 577–588.
- Burwell, R.D., Shapiro, M.L., O'Malley, M.T., Eichenbaum, H., **1998**. Positional firing properties of perirhinal cortex neurons. *Neuroreport*, 9, 3013–3018.
- Bush, D., Barry, C., Burgess, N., **2014**. What do grid cells contribute to place cell firing? *Trends in Neurosciences*, 37, 136–145.
- Buzsáki, G., **1986**. Hippocampal sharp waves: Their origin and significance. *Brain Research*, 398, 242–252.
- Buzsáki, G., **2005**. Theta rhythm of navigation: link between path integration and landmark navigation, episodic and semantic memory. *Hippocampus*, 15, 827–840.
- Buzsáki, G., Lai-Wo S., L., Vanderwolf, C.H., **1983**. Cellular bases of hippocampal EEG in the behaving rat. *Brain Research Reviews*, 6, 139–171.
- Buzsáki, G., Moser, E.I., **2013**. Memory, navigation and theta rhythm in the hippocampal-entorhinal system. *Nature Neuroscience*, 16, 130–138.

-C-

- Calton, J.L., Stackman, R.W., Goodridge, J.P., Archev, W.B., Dudchenko, P.A., Taube, J.S., **2003**. Hippocampal Place Cell Instability after Lesions of the Head Direction Cell Network. *The Journal of Neuroscience*, 23, 9719–9731.
- Canto, C.B., Wouterlood, F.G., Witter, M.P., **2008**. What Does the Anatomical Organization of the Entorhinal Cortex Tell Us? *Neural Plasticity*, 2008.
- Capaldi, E.J., **1957**. The effect of different amounts of alternating partial reinforcement on resistance to extinction. *The American Journal of Psychology*, 70, 451–452.
- Carr, M.F., Jadhav, S.P., Frank, L.M., **2011**. Hippocampal replay in the awake state: a potential substrate for memory consolidation and retrieval. *Nature Neuroscience*, 14, 147–153.
- Cassel, J.-C., Pereira de Vasconcelos, A., Loureiro, M., Cholvin, T., Dalrymple-Alford, J.C., Vertes, R.P., **2013**. The reuniens and rhomboid nuclei: Neuroanatomy, electrophysiological characteristics and behavioral implications. *Progress in Neurobiology*, 111, 34–52.
- Caston, J., Chianale, C., Delhaye-Bouchaud, N., Mariani, J., **1998**. Role of the cerebellum in exploration behavior. *Brain Research*, 808, 232–237.
- Chai, S.-C., White, N.M., **2004**. Effects of Fimbria-Fornix, Hippocampus, and Amygdala Lesions on Discrimination Between Proximal Locations. *Behavioral Neuroscience*, 118, 770–784.

- Chang, Q., Gold, P.E., **2003**. Switching Memory Systems during Learning: Changes in Patterns of Brain Acetylcholine Release in the Hippocampus and Striatum in Rats. *The Journal of Neuroscience*, 23, 3001–3005.
- Chapuis, N., Scardigli, P., **1993**. Shortcut ability in hamsters (*Mesocricetus auratus*): The role of environmental and kinesthetic information. *Animal Learning & Behavior*, 21, 255–265.
- Cheng, K., **1986**. A purely geometric module in the rat's spatial representation. *Cognition*, 23, 149–178.
- Cheng, K., **2008**. Whither geometry? Troubles of the geometric module. *Trends in Cognitive Sciences*, 12, 355–361.
- Chen, G., King, J.A., Burgess, N., O'Keefe, J., **2013**. How vision and movement combine in the hippocampal place code. *Proceedings of the National Academy of Sciences*, 110, 378–383.
- Chen, L.L., Lin, L.H., Green, E.J., Barnes, C.A., McNaughton, B.L., **1994**. Head-direction cells in the rat posterior cortex. I. Anatomical distribution and behavioral modulation. *Experimental Brain Research*, 101, 8–23.
- Chevaleyre, V., Siegelbaum, S.A., **2010**. Strong CA2 pyramidal neuron synapses define a powerful disinaptic cortico-hippocampal loop. *Neuron*, 66, 560–572.
- Chklovskii, D.B., Mel, B.W., Svoboda, K., **2004**. Cortical rewiring and information storage. *Nature*, 431, 782–788.
- Churchwell, J.C., Kesner, R.P., **2011**. Hippocampal-prefrontal dynamics in spatial working memory: Interactions and independent parallel processing. *Behavioural Brain Research*, 225, 389–395.
- Churchwell, J.C., Morris, A.M., Musso, N.D., Kesner, R.P., **2010**. Prefrontal and hippocampal contributions to encoding and retrieval of spatial memory. *Neurobiology of Learning and Memory*, 93, 415–421.
- Ciancia, F., **1991**. Tolman and Honzik (1930) revisited: or The mazes of psychology (1930–1980). *The Psychological Record*, 41, 461–472.
- Clark, R.E., Broadbent, N.J., Squire, L.R., **2005**. Hippocampus and remote spatial memory in rats. *Hippocampus*, 15, 260–272.
- Clayton, N.S., Dickinson, A., **1998**. Episodic-like memory during cache recovery by scrub jays. *Nature*, 395, 272–274.
- Cohen, N.J., Eichenbaum, H., **1991**. The theory that wouldn't die: a critical look at the spatial mapping theory of hippocampal function. *Hippocampus*, 1, 265–268.
- Cole, M.R., Chappell-Stephenson, R., **2003**. Exploring the limits of spatial memory in rats, using very large mazes. *Animal Learning & Behavior*, 31, 349–368.
- Colgin, L.L., Moser, E.I., Moser, M.-B., **2008**. Understanding memory through hippocampal remapping. *Trends in Neurosciences*, 31, 469–477.
- Corballis, M.C., **2007**. Recursion, Language, and Starlings. *Cognitive Science*, 31, 697–704.
- Corbit, L.H., Balleine, B.W., **2000**. The role of the hippocampus in instrumental conditioning. *The Journal of Neuroscience*, 20, 4233–4239.
- Corbit, L.H., Muir, J.L., Balleine, B.W., **2001**. The Role of the Nucleus Accumbens in Instrumental Conditioning: Evidence of a Functional Dissociation between Accumbens Core and Shell. *The Journal of Neuroscience*, 21, 3251–3260.

- Corbit, L.H., Ostlund, S.B., Balleine, B.W., **2002**. Sensitivity to Instrumental Contingency Degradation Is Mediated by the Entorhinal Cortex and Its Efferents via the Dorsal Hippocampus. *The Journal of Neuroscience*, 22, 10976–10984.
- Corkin, S., **1968**. Acquisition of motor skill after bilateral medial temporal-lobe excision. *Neuropsychologia*, 6, 255–265.
- Coutureau, E., Esclassan, F., Di Scala, G., Marchand, A.R., **2012**. The Role of the Rat Medial Prefrontal Cortex in Adapting to Changes in Instrumental Contingency. *PLoS ONE*, 7, e33302.
- Cowan, N., **1988**. Evolving conceptions of memory storage, selective attention, and their mutual constraints within the human information-processing system. *Psychological Bulletin*, 104, 163–191.
- Cowan, N., **2008**. What are the differences between long-term, short-term, and working memory? *Progress in Brain Research*, 169, 323–338.
- Cowan, N., Chen, Z., Rouder, J.N., **2004**. Constant Capacity in an Immediate Serial-Recall Task A Logical Sequel to Miller (1956). *Psychological Science*, 15, 634–640.
- Cressant, A., Muller, R.U., Poucet, B., **1997**. Failure of Centrally Placed Objects to Control the Firing Fields of Hippocampal Place Cells. *The Journal of Neuroscience*, 17, 2531–2542.
- Crystal, J.D., **2010**. Episodic-like memory in animals. *Behavioural Brain Research*, 215, 235–243.
- Crystal, J.D., Alford, W.T., Zhou, W., Hohmann, A.G., **2013**. Source Memory in the Rat. *Current Biology*, 23, 387–391.
- Curtis, C.E., Lee, D., **2010**. Beyond working memory: the role of persistent activity in decision making. *Trends in Cognitive Sciences*, 14, 216–222.

-D-

- Daw, N.D., O'Doherty, J.P., Dayan, P., Seymour, B., Dolan, R.J., **2006**. Cortical substrates for exploratory decisions in humans. *Nature*, 441, 876–879.
- Deacon, T.W., Eichenbaum, H., Rosenberg, P., Eckmann, K.W., **1983**. Afferent connections of the perirhinal cortex in the rat. *Journal of Comparative Neurology*, 220, 168–190.
- Debanne, D., Gähwiler, B.H., Thompson, S.M., **1998**. Long-term synaptic plasticity between pairs of individual CA3 pyramidal cells in rat hippocampal slice cultures. *The Journal of Physiology*, 507, 237–247.
- deCharms, R.C., Zador, A., **2000**. Neural Representation and the Cortical Code. *Annual Review of Neuroscience*, 23, 613–647.
- Dember, W.N., Fowler, H., **1958**. Spontaneous alternation behavior. *Psychological Bulletin*, 55, 412–428.
- Dere, E., Huston, J.P., De Souza Silva, M.A., **2005**. Integrated memory for objects, places, and temporal order: Evidence for episodic-like memory in mice. *Neurobiology of Learning and Memory*, 84, 214–221.
- Dere, E., Pause, B.M., Pietrowsky, R., **2010**. Emotion and episodic memory in neuropsychiatric disorders. *Behavioural Brain Research, Special issue on episodic memory*, 215, 162–171.

- De Saint Blanquat, P., Hok, V., Alvernhe, A., Save, E., Poucet, B., **2010**. Tagging items in spatial working memory: A unit-recording study in the rat medial prefrontal cortex. *Behavioural Brain Research*, 209, 267–273.
- De Saint Blanquat, P., Hok, V., Save, E., Poucet, B., Chaillan, F.A., **2013**. Differential role of the dorsal hippocampus, ventro-intermediate hippocampus, and medial prefrontal cortex in updating the value of a spatial goal. *Hippocampus*, 23, 342–351.
- Devan, B.D., White, N.M., **1999**. Parallel Information Processing in the Dorsal Striatum: Relation to Hippocampal Function. *The Journal of Neuroscience*, 19, 2789–2798.
- DeVito, L.M., Kanter, B.R., Eichenbaum, H., **2010a**. The hippocampus contributes to memory expression during transitive inference in mice. *Hippocampus*, 20, 208–217.
- DeVito, L.M., Lykken, C., Kanter, B.R., Eichenbaum, H., **2010b**. Prefrontal cortex: Role in acquisition of overlapping associations and transitive inference. *Learning & Memory*, 17, 161–167.
- Diana, R.A., Yonelinas, A.P., Ranganath, C., **2007**. Imaging recollection and familiarity in the medial temporal lobe: a three-component model. *Trends in Cognitive Sciences*, 11, 379–386.
- Diba, K., Buzsáki, G., **2007**. Forward and reverse hippocampal place-cell sequences during ripples. *Nature Neuroscience*, 10, 1241–1242.
- Dickinson, A., **1985**. Actions and Habits: The Development of Behavioural Autonomy. *Philosophical Transactions of the Royal Society of London. B. Biological Sciences*, 308, 67–78.
- Dickinson, A., Balleine, B., **1994**. Motivational control of goal-directed action. *Animal Learning & Behavior*, 22, 1–18.
- Doll, B.B., Shohamy, D., Daw, N.D., **2014**. Multiple memory systems as substrates for multiple decision systems. *Neurobiology of Learning and Memory*. In press.
- Dombeck, D.A., Harvey, C.D., Tian, L., Looger, L.L., Tank, D.W., **2010**. Functional imaging of hippocampal place cells at cellular resolution during virtual navigation. *Nature Neuroscience*, 13, 1433–1440.
- Doya, K., **2008**. Modulators of decision making. *Nature neuroscience*, 11, 410–416.
- Dudai, Y., **2004**. The Neurobiology of Consolidations, Or, How Stable is the Engram? *Annual Review of Psychology*, 55, 51–86.
- Dudek, S.M., Bear, M.F., **1992**. Homosynaptic long-term depression in area CA1 of hippocampus and effects of N-methyl-D-aspartate receptor blockade. *Proceedings of the National Academy of Sciences*, 89, 4363–4367.
- Dupret, D., O'Neill, J., Csicsvari, J., **2013**. Dynamic Reconfiguration of Hippocampal Interneuron Circuits during Spatial Learning. *Neuron*, 78, 166–180.
- Durstewitz, D., Vittoz, N.M., Floresco, S.B., Seamans, J.K., **2010**. Abrupt transitions between prefrontal neural ensemble states accompany behavioral transitions during rule learning. *Neuron*, 66, 438–448.
- Dusek, J.A., Eichenbaum, H., **1997**. The hippocampus and memory for orderly stimulus relations. *Proceedings of the National Academy of Sciences*, 94, 7109–7114.

-E-

- Eacott, M.J., Easton, A., **2010**. Episodic memory in animals: Remembering which occasion. *Neuropsychologia*, 48, 2273–2280.
- Eacott, M.J., Easton, A., Zinkivskay, A., **2005**. Recollection in an episodic-like memory task in the rat. *Learning & Memory*, 12, 221–223.
- Easton, A., Webster, L.A.D., Eacott, M.J., **2012**. The episodic nature of episodic-like memories. *Learning & Memory*, 19, 146–150.
- Ego-Stengel, V., Wilson, M.A., **2010**. Disruption of ripple-associated hippocampal activity during rest impairs spatial learning in the rat. *Hippocampus*, 20, 1–10.
- Eichenbaum, H., **2001**. The hippocampus and declarative memory: cognitive mechanisms and neural codes. *Behavioural Brain Research*, 127, 199–207.
- Eichenbaum, H., Cohen, N.J., **2014**. Can We Reconcile the Declarative Memory and Spatial Navigation Views on Hippocampal Function? *Neuron*, 83, 764–770.
- Eichenbaum, H., Dudchenko, P., Wood, E., Shapiro, M., Tanila, H., **1999**. The hippocampus, memory, and place cells: is it spatial memory or a memory space? *Neuron*, 23, 209–226.
- Eichenbaum, H., Fortin, N., **2003**. Episodic Memory and the Hippocampus : It's About Time. *Current Directions in Psychological Science*, 12, 53–57.
- Eichenbaum, H., Kuperstein, M., Fagan, A., Nagode, J., **1987**. Cue-sampling and goal-approach correlates of hippocampal unit activity in rats performing an odor-discrimination task. *The Journal of Neuroscience*, 7, 716–732.
- Eichenbaum, H., Yonelinas, A.P., Ranganath, C., **2007**. The Medial Temporal Lobe and Recognition Memory. *Annual Review of Neuroscience*, 30, 123–152.
- Eilam, D., Golani, I., **1989**. Home base behavior of rats (*Rattus norvegicus*) exploring a novel environment. *Behavioural Brain Research*, 34, 199–211.
- Ekstrom, A.D., Kahana, M.J., Caplan, J.B., Fields, T.A., Isham, E.A., Newman, E.L., Fried, I., **2003**. Cellular networks underlying human spatial navigation. *Nature*, 425, 184–188.
- Elduayen, C., Save, E., **2014**. The retrosplenial cortex is necessary for path integration in the dark. *Behavioural Brain Research*, 272, 303–307.
- Ellen, P., Parko, E.M., Wages, C., Doherty, D., Herrmann, T., **1982**. Spatial problem solving by rats: Exploration and cognitive maps. *Learning and Motivation*, 13, 81–94.
- Elliot, A.J., Church, M.A., **1997**. A hierarchical model of approach and avoidance achievement motivation. *Journal of Personality and Social Psychology*, 72, 218–232.
- Elliot, A. J., & Fryer, J. W. **2008**. The goal construct in psychology. In Shah & Gardner eds: *Handbook of Motivation Science*, 18, 235–250.
- Ennaceur, A., Delacour, J., **1988**. A new one-trial test for neurobiological studies of memory in rats. 1: Behavioral data. *Behavioural Brain Research*, 31, 47–59.
- Epszstein, J., Brecht, M., Lee, A.K., **2011**. Intracellular Determinants of Hippocampal CA1 Place and Silent Cell Activity in a Novel Environment. *Neuron*, 70, 109–120.
- Erdem, U.M., Hasselmo, M., **2012**. A goal-directed spatial navigation model using forward trajectory planning based on grid cells. *European Journal of Neuroscience*, 35, 916–931.
- Etienne, A.S., Jeffery, K.J., **2004**. Path integration in mammals. *Hippocampus*, 14, 180–192.

- Etienne, A.S., Teroni, E., Hurni, C., Portenier, V., **1990**. The effect of a single light cue on homing behaviour of the golden hamster. *Animal Behaviour*, 39, 17–41.
- Eustache, F., Desgranges, B., **2008**. MNESIS: Towards the Integration of Current Multisystem Models of Memory. *Neuropsychology Review*, 18, 53–69.
- Everitt, B.J., Robbins, T.W., **2005**. Neural systems of reinforcement for drug addiction: from actions to habits to compulsion. *Nature Neuroscience*, 8, 1481–1489.

-F-

- Feierstein, C.E., Quirk, M.C., Uchida, N., Sosulski, D.L., Mainen, Z.F., **2006**. Representation of Spatial Goals in Rat Orbitofrontal Cortex. *Neuron*, 51, 495–507.
- Fellini, L., Florian, C., Courtney, J., Rouillet, P., **2009**. Pharmacological intervention of hippocampal CA3 NMDA receptors impairs acquisition and long-term memory retrieval of spatial pattern completion task. *Learning & Memory*, 16, 387–394.
- Fellini, L., Morellini, F., **2011**. Geometric information is required for allothetic navigation in mice. *Behavioural Brain Research*, 222, 380–384.
- Fellini, L., Morellini, F., **2013**. Mice Create What–Where–When Hippocampus-Dependent Memories of Unique Experiences. *The Journal of Neuroscience*, 33, 1038–1043.
- Fenton, A.A., Arolfo, M.P., Bures, J., **1994**. Place navigation in the morris water maze under minimum and redundant extra-maze cue conditions. *Behavioral and Neural Biology*, 62, 178–189.
- Fenton, A.A., Csizmadia, G., Muller, R.U., **2000a**. Conjoint Control of Hippocampal Place Cell Firing by Two Visual Stimuli I. the Effects of Moving the Stimuli on Firing Field Positions. *The Journal of General Physiology*, 116, 191–210.
- Fenton, A.A., Csizmadia, G., Muller, R.U., **2000b**. Conjoint Control of Hippocampal Place Cell Firing by Two Visual Stimuli II. a Vector-Field Theory That Predicts Modifications of the Representation of the Environment. *The Journal of General Physiology*, 116, 211–222.
- Fenton, A.A., Kao, H.-Y., Neymotin, S.A., Olypher, A., Vayntrub, Y., Lytton, W.W., Ludvig, N., **2008**. Unmasking the CA1 Ensemble Place Code by Exposures to Small and Large Environments: More Place Cells and Multiple, Irregularly Arranged, and Expanded Place Fields in the Larger Space. *The Journal of Neuroscience*, 28, 11250–11262.
- Fenton, A.A., Lytton, W.W., Barry, J.M., Lenck-Santini, P.P., Zinyuk, L.E., Kubík, Š., Bureš, J., Poucet, B., Muller, R.U., Olypher, A.V., **2010**. Attention-Like Modulation of Hippocampus Place Cell Discharge. *The Journal of Neuroscience*, 30, 4613–4625.
- Fenton, A.A., Muller, R.U., **1998**. Place cell discharge is extremely variable during individual passes of the rat through the firing field. *Proceedings of the National Academy of Sciences*, 95, 3182–3187.
- Fenton, A.A., Wesierska, M., Kaminsky, Y., Bures, J., **1998**. Both here and there: Simultaneous expression of autonomous spatial memories in rats. *Proceedings of the National Academy of Sciences*, 95, 11493–11498.
- Ferbinteanu, J., Ray, C., McDonald, R.J., **2003**. Both dorsal and ventral hippocampus contribute to spatial learning in Long–Evans rats. *Neuroscience Letters*, 345, 131–135.
- Ferbinteanu, J., Shapiro, M.L., **2003**. Prospective and Retrospective Memory Coding in the Hippocampus. *Neuron*, 40, 1227–1239.

- Ferretti, V., Roulet, P., Sargolini, F., Rinaldi, A., Perri, V., Fabbro, M.D., Costantini, V.J.A., Annese, V., Scesa, G., Stefano, M.E.D., Oliverio, A., Mele, A., **2010**. Ventral striatal plasticity and spatial memory. *Proceedings of the National Academy of Sciences*, 107, 7945–7950.
- Feretti, V., Sargolini, F., Oliverio, A., Mele, A., Roulet, P., **2007**. Effects of intra-accumbens NMDA and AMPA receptor antagonists on short-term spatial learning in the Morris water maze task. *Behavioural Brain Research*, 179, 43–49.
- Floresco, S.B., Zhang, Y., Enomoto, T., **2009**. Neural circuits subserving behavioral flexibility and their relevance to schizophrenia. *Behavioural Brain Research*, 204, 396–409.
- Fonio, E., Benjamini, Y., Golani, I., **2009**. Freedom of movement and the stability of its unfolding in free exploration of mice. *Proceedings of the National Academy of Sciences*, 106, 21335–21340.
- Fortin, N.J., Agster, K.L., Eichenbaum, H.B., **2002**. Critical role of the hippocampus in memory for sequences of events. *Nature Neuroscience*, 5, 458–462.
- Foster, D.J., Wilson, M.A., **2006**. Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature*, 440, 680–683.
- Foster, D.J., Wilson, M.A., **2007**. Hippocampal theta sequences. *Hippocampus*, 17, 1093–1099.
- Foster, T.C., Castro, C.A., McNaughton, B.L., **1989**. Spatial selectivity of rat hippocampal neurons: dependence on preparedness for movement. *Science*, 244, 1580–1582.
- Fouquet, C., Tobin, C., Rondi-Reig, L., **2010**. A new approach for modeling episodic memory from rodents to humans: The temporal order memory. *Behavioural Brain Research*, Special issue on episodic memory, 215, 172–179.
- Frankland, P.W., Bontempi, B., **2005**. The organization of recent and remote memories. *Nature Reviews Neuroscience*, 6, 119–130.
- Frank, L.M., Brown, E.N., Wilson, M., **2000**. Trajectory Encoding in the Hippocampus and Entorhinal Cortex. *Neuron*, 27, 169–178.
- Frank, M.J., Doll, B.B., Oas-Terpstra, J., Moreno, F., **2009**. Prefrontal and striatal dopaminergic genes predict individual differences in exploration and exploitation. *Nature Neuroscience*, 12, 1062–1068.
- Fujisawa, S., Buzsáki, G., **2011**. A 4 Hz oscillation adaptively synchronizes prefrontal, VTA, and hippocampal activities. *Neuron*, 72, 153–165.
- Funahashi, S., Bruce, C.J., Goldman-Rakic, P.S., **1989**. Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *Journal of Neurophysiology*, 61, 331–349.
- Funahashi, S., Kubota, K., **1994**. Working memory and prefrontal cortex. *Neuroscience Research*, 21, 1–11.
- Fuster, J.M., **1997**. The prefrontal cortex: Anatomy, physiology, and neuropsychology of the frontal lobe. *Raven*, New York.
- Fuster, J.M., Alexander, G.E., **1971**. Neuron Activity Related to Short-Term Memory. *Science*, 173, 652–654.
- Fyhn, M., Hafting, T., Treves, A., Moser, M.-B., Moser, E.I., **2007**. Hippocampal remapping and grid realignment in entorhinal cortex. *Nature*, 446, 190–194.
- Fyhn, M., Molden, S., Hollup, S.A., Moser, M.-B., Moser, E.I., **2002**. Hippocampal Neurons Responding to First-Time Dislocation of a Target Object. *Neuron*, 35, 555–566.
- Fyhn, M., Molden, S., Witter, M.P., Moser, E.I., Moser, M.-B., **2004**. Spatial Representation in the Entorhinal Cortex. *Science*, 305, 1258–1264.

-G-

- Gallistel, C.R., **1990**. The organization of learning. *MIT press Cambridge, MA*.
- Gallistel, C.R., **2001**. Mental Representations, Psychology Of, in: Smelser, N.J., Baltes, B. (Eds.), *International Encyclopedia of the Social and Behavioral Sciences*. pp. 9691–9695.
- Gasbarri, A., Packard, M.G., Campana, E., Pacitti, C., **1994**. Anterograde and retrograde tracing of projections from the ventral tegmental area to the hippocampal formation in the rat. *Brain Research Bulletin*, 33, 445–452.
- Gaskin, S., Chai, S.-C., White, N.M., **2005**. Inactivation of the dorsal hippocampus does not affect learning during exploration of a novel environment. *Hippocampus*, 15, 1085–1093.
- Gelbard-Sagiv, H., Mukamel, R., Harel, M., Malach, R., Fried, I., **2008**. Internally Generated Reactivation of Single Neurons in Human Hippocampus During Free Recall. *Science*, 322, 96–101.
- Gentry, G., Lynn, W., Lee, H., **1948**. Spatial location in the learning of a multiple-T maze. *Journal of Comparative and Physiological Psychology*, 41, 312–318.
- Georgopoulos, A.P., Schwartz, A.B., Kettner, R.E., **1986**. Neuronal population coding of movement direction. *Science*, 233, 1416–1419.
- Gerstner, W., Kempter, R., van Hemmen, J.L., Wagner, H., **1996**. A neuronal learning rule for sub-millisecond temporal coding. *Nature*, 383, 76–78.
- Gerstner, W., Kistler, W.M., **2002**. Spiking Neuron Models: Single Neurons, Populations, Plasticity. *Cambridge University Press*.
- Ghanbarian, E., Motamedi, F., **2013**. Ventral tegmental area inactivation suppresses the expression of CA1 long term potentiation in anesthetized rat. *PLoS ONE*, 8, e58844.
- Gibson, B.M., **2001**. Cognitive maps not used by humans (*Homo sapiens*) during a dynamic navigational task. *Journal of Comparative Psychology*, 115, 397–402.
- Gilbert, P.E., Campbell, A., Kesner, R.P., **2003**. The role of the amygdala in conditioned flavor preference. *Neurobiology of Learning and Memory*, 79, 118–121.
- Girardeau, G., Benchenane, K., Wiener, S.I., Buzsáki, G., Zugaro, M.B., **2009**. Selective suppression of hippocampal ripples impairs spatial memory. *Nature Neuroscience*, 12, 1222–1223.
- Girardeau, G., Zugaro, M., **2011**. Hippocampal ripples and memory consolidation. *Current Opinion in Neurobiology*, Behavioural and Cognitive Neuroscience, 21, 452–459.
- Gloveli, T., **2010**. Hippocampal spatial navigation: interneurons take responsibility. *The Journal of Physiology*, 588, 4609–4610.
- Golding, N.L., Mickus, T.J., Katz, Y., Kath, W.L., Spruston, N., **2005**. Factors mediating powerful voltage attenuation along CA1 pyramidal neuron dendrites. *The Journal of Physiology*, 568, 69–82.
- Goldman-Rakic, P.S., **1995**. Cellular basis of working memory. *Neuron*, 14, 477–485.
- Golob, E.J., Taube, J.S., **1999**. Head Direction Cells in Rats with Hippocampal or Overlying Neocortical Lesions: Evidence for Impaired Angular Path Integration. *The Journal of Neuroscience*, 19, 7198–7211.
- Griffiths, D., Dickinson, A., Clayton, N., **1999**. Episodic memory: what can animals remember about their past? *Trends in Cognitive Sciences*, 3, 74–80.

- Guillette, L.M., Reddon, A.R., Hurd, P.L., Sturdy, C.B., **2009**. Exploration of a novel space is associated with individual differences in learning speed in black-capped chickadees, *Poecile atricapillus*. *Behavioural Processes*, 82, 265–270.
- Gupta, A.S., van der Meer, M.A.A., Touretzky, D.S., Redish, A.D., **2012**. Segmentation of spatial experience by hippocampal theta sequences. *Nature Neuroscience*, 15, 1032–1039.
- Guzowski, J.F., Knierim, J.J., Moser, E.I., **2004**. Ensemble Dynamics of Hippocampal Regions CA3 and CA1. *Neuron*, 44, 581–584.
- Guzowski, J.F., McNaughton, B.L., Barnes, C.A., Worley, P.F., **1999**. Environment-specific expression of the immediate-early gene *Arc* in hippocampal neuronal ensembles. *Nature Neuroscience*, 2, 1120–1124.
- H-**
- Hafting, T., Fyhn, M., Molden, S., Moser, M.-B., Moser, E.I., **2005**. Microstructure of a spatial map in the entorhinal cortex. *Nature*, 436, 801–806.
- Hamilton, D.A., Akers, K.G., Johnson, T.E., Rice, J.P., Candelaria, F.T., Redhead, E.S., **2009**. Evidence for a shift from place navigation to directional responding in one variant of the Morris water task. *Journal of Experimental Psychology: Animal Behavior Processes*, 35, 271–278.
- Hamilton, D.A., Akers, K.G., Weisend, M.P., Sutherland, R.J., **2007**. How do room and apparatus cues control navigation in the Morris water task? Evidence for distinct contributions to a movement vector. *Journal of Experimental Psychology: Animal Behavior Processes*, 33, 100–114.
- Hamilton, D.A., Driscoll, I., Sutherland, R.J., **2002**. Human place learning in a virtual Morris water task: some important constraints on the flexibility of place navigation. *Behavioural Brain Research*, 129, 159–170.
- Hampson, R.E., Simeral, J.D., Deadwyler, S.A., **1999**. Distribution of spatial and nonspatial information in dorsal hippocampus. *Nature*, 402, 610–614.
- Hangya, B., Li, Y., Muller, R.U., Czurkó, A., **2010**. Complementary spatial firing in place cell–interneuron pairs. *The Journal of Physiology*, 588, 4165–4175.
- Harand, C., Bertran, F., La Joie, R., Landeau, B., Mézenge, F., Desgranges, B., Peigneux, P., Eustache, F., Rauchs, G., **2012**. The Hippocampus Remains Activated over the Long Term for the Retrieval of Truly Episodic Memories. *PLoS ONE*, 7, e43495.
- Hargreaves, E.L., Rao, G., Lee, I., Knierim, J.J., **2005**. Major Dissociation Between Medial and Lateral Entorhinal Input to Dorsal Hippocampus. *Science*, 308, 1792–1794.
- Hartley, T., Lever, C., Burgess, N., O’Keefe, J., **2014**. Space in the brain: how the hippocampal formation supports spatial cognition. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369, 20120510.
- Harvey, C.D., Collman, F., Dombeck, D.A., Tank, D.W., **2009**. Intracellular dynamics of hippocampal place cells during virtual navigation. *Nature*, 461, 941–946.
- Hauser, M.D., Chomsky, N., Fitch, W.T., **2002**. The Faculty of Language: What Is It, Who Has It, and How Did It Evolve? *Science*, 298, 1569–1579.
- Hebb, D., **1949**. The Organization of Behaviour: A Neuropsychological Theory. *Psychology Press*.

- Hembrook, J.R., Onos, K.D., Mair, R.G., **2012**. Inactivation of ventral midline thalamus produces selective spatial delayed conditional discrimination impairment in the rat. *Hippocampus*, 22, 853–860.
- Hetherington, P.A., Shapiro, M.L., **1997**. Hippocampal place fields are altered by the removal of single visual cues in a distance-dependent manner. *Behavioral Neuroscience*, 111, 20–34.
- Hermer, L., Spelke, E., **1996**. Modularity and development: the case of spatial reorientation. *Cognition*, 61, 195–232.
- Higgins, E.T., **2011**. *Beyond Pleasure and Pain: How Motivation Works*. Oxford University Press.
- Hill, A.J., **1978**. First occurrence of hippocampal spatial firing in a new environment. *Experimental Neurology*, 62, 282–297.
- Hillman, K.L., Bilkey, D.K., **2010**. Neurons in the Rat Anterior Cingulate Cortex Dynamically Encode Cost-Benefit in a Spatial Decision-Making Task. *The Journal of Neuroscience*, 30, 7705–7713.
- Hitti, F.L., Siegelbaum, S.A., **2014**. The hippocampal CA2 region is essential for social memory. *Nature*, 508, 88–92.
- Hirel, J., Gaussier, P., Quoy, M., Banquet, J.P., Save, E., Poucet, B., **2013**. The hippocampo-cortical loop: Spatio-temporal learning and goal-oriented planning in navigation. *Neural Networks*, 43, 8–21.
- Hok, V., Chah, E., Reilly, R.B., O'Mara, S.M., **2012**. Hippocampal Dynamics Predict Inter-individual Cognitive Differences in Rats. *The Journal of Neuroscience*, 32, 3540–3551.
- Hok, V., Chah, E., Save, E., Poucet, B., **2013**. Prefrontal Cortex Focally Modulates Hippocampal Place Cell Firing Patterns. *The Journal of Neuroscience*, 33, 3443–3451.
- Hok, V., Lenck-Santini, P.P., Roux, S., Save, E., Muller, R.U., Poucet, B., **2007a**. Goal-Related Activity in Hippocampal Place Cells. *The Journal of Neuroscience*, 27, 472–482.
- Hok, V., Lenck-Santini, P.-P., Save, E., Gaussier, P., Banquet, J.-P., Poucet, B., **2007b**. A test of the time estimation hypothesis of place cell goal-related activity. *Journal of Integrative Neuroscience*, 6, 367–378.
- Hok, V., Save, E., Lenck-Santini, P.P., Poucet, B., **2005**. Coding for spatial goals in the prelimbic/infralimbic area of the rat frontal cortex. *Proceedings of the National Academy of Sciences*, 102, 4602–4607.
- Hollup, S.A., Kjelstrup, K.G., Hoff, J., Moser, M.B., Moser, E.I., **2001a**. Impaired recognition of the goal location during spatial navigation in rats with hippocampal lesions. *The Journal of Neuroscience*, 21, 4505–4513.
- Hollup, S.A., Molden, S., Donnett, J.G., Moser, M.-B., Moser, E.I., **2001b**. Accumulation of Hippocampal Place Fields at the Goal Location in an Annular Watermaze Task. *The Journal of Neuroscience*, 21, 1635–1644.
- Hölscher, C., Jacob, W., Mallot, H.A., **2003**. Reward modulates neuronal activity in the hippocampus of the rat. *Behavioural Brain Research*, 142, 181–191.
- Hoover, W.B., Vertes, R.P., **2012**. Collateral projections from nucleus reuniens of thalamus to hippocampus and medial prefrontal cortex in the rat: a single and double retrograde fluorescent labeling study. *Brain Structure and Function*, 217, 191–209.
- Hopfield, J., **2007**. Hopfield network. *Scholarpedia*, 2, 1977.
- Hopfield, J.J., **1982**. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, 79, 2554–2558.

- Horne, M.R., Iordanova, M.D., Pearce, J.M., **2010**. Spatial learning based on boundaries in rats is hippocampus-dependent and prone to overshadowing. *Behavioral Neuroscience*, 124, 623–632.
- Horne, M.R., Martin, G.M., Harley, C.W., Skinner, D.M., **2007**. Where am I? Distal cues use requires sensitivity to start location change in the rat. *Journal of Experimental Psychology: Animal Behavior Processes*, 33, 92–99.
- Howe, M.W., Tierney, P.L., Sandberg, S.G., Phillips, P.E.M., Graybiel, A.M., **2013**. Prolonged dopamine signalling in striatum signals proximity and value of distant rewards. *Nature*, 500, 575–579.
- Hubel, D.H., Wiesel, T.N., **1959**. Receptive fields of single neurones in the cat's striate cortex. *The Journal of Physiology*, 148, 574–591.
- Hu, D., Amsel, A., **1995**. A simple test of the vicarious trial-and-error hypothesis of hippocampal function. *Proceedings of the National Academy of Sciences*, 92, 5506–5509.
- Hu, D., Xu, X., Gonzalez-Lima, F., **2006**. Vicarious trial-and-error behavior and hippocampal cytochrome oxidase activity during y-maze discrimination learning in the rat. *International Journal of Neuroscience*, 116, 265–280.
- Hull, C., **1943**. Principles of behavior: An introduction of behavior theory.
- Hummel, J.E., **2000**. Localism as a first step toward symbolic representation. *Behavioral and Brain Sciences*, 23, 480–481.
- Humphries, M.D., Prescott, T.J., **2010**. The ventral basal ganglia, a selection mechanism at the crossroads of space, strategy, and reward. *Progress in Neurobiology*, 90, 385–417.
- Hunsaker, M.R., Kesner, R.P., **2013**. The operation of pattern separation and pattern completion processes associated with different attributes or domains of memory. *Neuroscience & Biobehavioral Reviews*, 37, 36–58.
- Huxter, J.R., Senior, T.J., Allen, K., Csicsvari, J., **2008**. Theta phase-specific codes for two-dimensional position, trajectory and heading in the hippocampus. *Nature Neuroscience*, 11, 587–594.

-I-

- Inglis, I.R., Langton, S., Forkman, B., Lazarus, J., **2001**. An information primacy model of exploratory and foraging behaviour. *Animal Behaviour*, 62, 543–557.
- Ishikawa, A., Nakamura, S., **2006**. Ventral Hippocampal Neurons Project Axons Simultaneously to the Medial Prefrontal Cortex and Amygdala in the Rat. *Journal of Neurophysiology*, 96, 2134–2138.
- Ito, M., **2001**. Cerebellar Long-Term Depression: Characterization, Signal Transduction, and Functional Roles. *Physiological Reviews*, 81, 1143–1195.
- Itskov, P.M., Vinnik, E., Honey, C., Schnupp, J., Diamond, M.E., **2012**. Sound sensitivity of neurons in rat hippocampus during performance of a sound-guided task. *Journal of Neurophysiology*, 107, 1822–1834.

-J-

- Jacob, P.-Y., **2014**. Rôle du cortex entorhinal médian dans le traitement des informations spatiales: études comportementales et électrophysiologiques. Aix-Marseille Université, Doctoral thesis.

- Jacob, P.-Y., Poucet, B., Liberge, M., Save, E., Sargolini, F., **2014**. Vestibular control of entorhinal cortex activity in spatial navigation. *Frontiers in Integrative Neuroscience*, 8, 38.
- Jacobs, L.F., Schenk, F., **2003**. Unpacking the cognitive map: The parallel map theory of hippocampal function. *Psychological Review*, 110, 285–315.
- Jacobson, T.K., Gruenbaum, B.F., Markus, E.J., **2012**. Extensive training and hippocampus or striatum lesions: Effect on place and response strategies. *Physiology & Behavior*, 105(3), 645-652.
- Jankowski, M.M., Islam, M.N., Wright, N.F., Vann, S.D., Erichsen, J.T., Aggleton, J.P., O'Mara, S.M., **2014**. Nucleus reuniens of the thalamus contains head direction cells. *Elife*, 3, e03075.
- Jankowski, M.M., Ronnqvist, K.C., Tsanov, M., Vann, S.D., Wright, N.F., Erichsen, J.T., Aggleton, J.P., O'Mara, S.M., **2013**. The anterior thalamus provides a subcortical circuit supporting memory and spatial navigation. *Frontiers in Systems Neuroscience*, 7.
- Jarrard, L.E., Okaichi, H., Steward, O., Goldschmidt, R.B., **1984**. On the role of hippocampal connections in the performance of place and cue tasks: Comparisons with damage to hippocampus. *Behavioral Neuroscience*, 98, 946–954.
- Jay, T.M., Witter, M.P., **1991**. Distribution of hippocampal CA1 and subicular efferents in the prefrontal cortex of the rat studied by means of anterograde transport of *Phaseolus vulgaris*-leucoagglutinin. *The Journal of Comparative Neurology*, 313, 574–586.
- Jeewajee, A., Barry, C., Douchamps, V., Manson, D., Lever, C., Burgess, N., **2014**. Theta phase precession of grid and place cell firing in open environments. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369, 20120532.
- Jeffery, K.J., **1998**. Learning of landmark stability and instability by hippocampal place cells. *Neuropharmacology*, 37, 677–687.
- Jeffery, K.J., Anderson, M.I., **2003**. Dissociation of the geometric and contextual influences on place cells. *Hippocampus*, 13, 868–872.
- Jeffery, K.J., Gilbert, A., Burton, S., Strudwick, A., **2003**. Preserved performance in a hippocampal-dependent spatial task despite complete place cell remapping. *Hippocampus*, 13, 175–189.
- Jeffery, K.J., O'Keefe, J.M., **1999**. Learned interaction of visual and idiothetic cues in the control of place field orientation. *Experimental Brain Research*, 127, 151–161.
- Jensen, O., Lisman, J.E., **2000**. Position Reconstruction From an Ensemble of Hippocampal Place Cells: Contribution of Theta Phase Coding. *Journal of Neurophysiology*, 83, 2602–2609.
- Jepma, M., Te Beek, E.T., Wagenmakers, E.-J., van Gerven, J.M.A., Nieuwenhuis, S., **2010**. The role of the noradrenergic system in the exploration-exploitation trade-off: a psychopharmacological study. *Frontiers in Human Neuroscience*, 4, 170.
- Johnson, A., Crowe, D.A., **2009**. Revisiting Tolman, his theories and cognitive maps. *Cognitive Critique*, 1, 43–72.
- Johnson, A., Redish, A.D., **2007**. Neural Ensembles in CA3 Transiently Encode Paths Forward of the Animal at a Decision Point. *The Journal of Neuroscience*, 27, 12176–12189.
- Johnson, A., van der Meer, M.A., Redish, A.D., **2007**. Integrating hippocampus and striatum in decision-making. *Current Opinion in Neurobiology*, 17, 692–697.
- Johnson, A., Varberg, Z., Benhardus, J., Maahs, A., Schrater, P., **2012**. The hippocampus and exploration: dynamically evolving behavior and neural representations. *Frontiers in Human Neuroscience*, 6, 216.

- Jones, M.W., McHugh, T.J., **2011**. Updating hippocampal representations: CA2 joins the circuit. *Trends in Neurosciences*, Special Issue: Hippocampus and Memory, 34, 526–535.
- Jung, M.W., Qin, Y., McNaughton, B.L., Barnes, C.A., **1998**. Firing characteristics of deep layer neurons in prefrontal cortex in rats performing spatial working memory tasks. *Cerebral Cortex*, 8, 437 – 450.
- Jung, M.W., Wiener, S., McNaughton, B., **1994**. Comparison of spatial firing characteristics of units in dorsal and ventral hippocampus of the rat. *The Journal of Neuroscience*, 14, 7347 –7356.

-K-

- Kandel, E.R., Dudai, Y., Mayford, M.R., **2014**. The Molecular and Systems Biology of Memory. *Cell*, 157, 163–186.
- Kennedy, P.J., Shapiro, M.L., **2009**. Motivational states activate distinct hippocampal representations to guide goal-directed behaviors. *Proceedings of the National Academy of Sciences*, 106, 10805–10810.
- Kentros, C.G., Agnihotri, N.T., Streater, S., Hawkins, R.D., Kandel, E.R., **2004**. Increased Attention to Spatial Context Increases Both Place Field Stability and Spatial Memory. *Neuron*, 42, 283–295.
- Kentros, C., Hargreaves, E., Hawkins, R.D., Kandel, E.R., Shapiro, M., Muller, R.V., **1998**. Abolition of Long-Term Stability of New Hippocampal Place Cell Maps by NMDA Receptor Blockade. *Science*, 280, 2121 –2126.
- Kesner, R.P., Farnsworth, G., DiMattia, B.V., **1989**. Double dissociation of egocentric and allocentric space following medial prefrontal and parietal cortex lesions in the rat. *Behavioral Neuroscience*, 103, 956–961.
- Kesner, R.P., Williams, J.M., **1995**. Memory for magnitude of reinforcement: dissociation between the amygdala and hippocampus. *Neurobiology of Learning and Memory*, 64, 237–244.
- Killcross, S., Coutureau, E., **2003**. Coordination of Actions and Habits in the Medial Prefrontal Cortex of Rats. *Cerebral Cortex*, 13, 400 –408.
- Kimble, D.P., Greene, E.G., **1968**. Absence of latent learning in rats with hippocampal lesions. *Psychonomic Science*, 11, 99–100.
- Kim, H., Sul, J.H., Huh, N., Lee, D., Jung, M.W., **2009**. Role of striatum in updating values of chosen actions. *The Journal of Neuroscience*, 29, 14701–14712.
- Kim, S., Sapiurka, M., Clark, R.E., Squire, L.R., **2013**. Contrasting effects on path integration after hippocampal damage in humans and rats. *Proceedings of the National Academy of Sciences*, 110, 4732–4737.
- Kishi, T., Tsumori, T., Yokota, S., Yasui, Y., **2006**. Topographical projection from the hippocampal formation to the amygdala: A combined anterograde and retrograde tracing study in the rat. *Journal of Comparative Neurology*, 496, 349–368.
- Kjelstrup, K.B., Solstad, T., Brun, V.H., Hafting, T., Leutgeb, S., Witter, M.P., Moser, E.I., Moser, M.-B., **2008**. Finite Scale of Spatial Representation in the Hippocampus. *Science*, 321, 140–143.
- Kjelstrup, K.G., Tuvnes, F.A., Steffenach, H.-A., Murison, R., Moser, E.I., Moser, M.-B., **2002**. Reduced fear expression after lesions of the ventral hippocampus. *Proceedings of the National Academy of Sciences*, 99, 10825 –10830.

- Klausberger, T., Somogyi, P., **2008**. Neuronal Diversity and Temporal Dynamics: The Unity of Hippocampal Circuit Operations. *Science*, 321, 53–57.
- Klein, S., **2013**. Making the case that episodic recollection is attributable to operations occurring at retrieval rather than to content stored in a dedicated subsystem of long-term memory. *Frontiers in Behavioral Neuroscience*, 7, 3.
- Klein, S.B., Nichols, S., **2012**. Memory and the Sense of Personal Identity. *Mind* fzs080.
- Kloosterman, F., van Haften, T., Lopes da Silva, F.H., **2004**. Two reentrant pathways in the hippocampal-entorhinal system. *Hippocampus*, 14, 1026–1039.
- Kloosterman, F., Witter, M.P., Van Haften, T., **2003**. Topographical and laminar organization of subicular projections to the parahippocampal region of the rat. *Journal of Comparative Neurology*, 455, 156–171.
- Knierim, J.J., Hamilton, D.A., **2011**. Framing Spatial Cognition: Neural Representations of Proximal and Distal Frames of Reference and Their Roles in Navigation. *Physiological Reviews*, 91, 1245 – 1279.
- Knierim, J.J., Kudrimoti, H.S., McNaughton, B.L., **1995**. Place cells, head direction cells, and the learning of landmark stability. *The Journal of Neuroscience*, 15, 1648–1659.
- Knight, R., Hayman, R., Ginzberg, L.L., Jeffery, K., **2011**. Geometric Cues Influence Head Direction Cells Only Weakly in Nondisoriented Rats. *The Journal of Neuroscience*, 31, 15681–15692.
- Knowlton, B.J., Squire, L.R., **1995**. Remembering and knowing: two different expressions of declarative memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21, 699–710.
- Kobayashi, T., Tran, A., Nishijo, H., Ono, T., Matsumoto, G., **2003**. Contribution of hippocampal place cell activity to learning and formation of goal-directed navigation in rats. *Neuroscience*, 117, 1025–1035.
- Kohara, K., Pignatelli, M., Rivest, A.J., Jung, H.-Y., Kitamura, T., Suh, J., Frank, D., Kajikawa, K., Mise, N., Obata, Y., Wickersham, I.R., Tonegawa, S., **2014**. Cell type-specific genetic and optogenetic tools reveal hippocampal CA2 circuits. *Nature Neuroscience*, 17, 269–279.
- Komorowski, R.W., Manns, J.R., Eichenbaum, H., **2009**. Robust conjunctive item-place coding by hippocampal neurons parallels learning what happens where. *The Journal of Neuroscience*, 29, 9918–9929.
- Kraus, B.J., Robinson II, R.J., White, J.A., Eichenbaum, H., Hasselmo, M.E., **2013**. Hippocampal “Time Cells”: Time versus Path Integration. *Neuron*, 78, 1090–1101.
- Kubie, J.L., Fenton, A., Novikov, N., Touretzky, D., Muller, R.U., **2007**. Changes in goal selection induced by cue conflicts are in register with predictions from changes in place cell field locations. *Behavioral Neuroscience*, 121, 751–763.
- Kumaran, D., McClelland, J.L., **2012**. Generalization through the recurrent interaction of episodic memories: A model of the hippocampal system. *Psychological Review* 119, 573–616.
- Kyd, R.J., Bilkey, D.K., **2003**. Prefrontal Cortex Lesions Modify the Spatial Properties of Hippocampal Place Cells. *Cerebral Cortex*, 13, 444–451.
- Kyd, R.J., Bilkey, D.K., **2005**. Hippocampal Place Cells Show Increased Sensitivity to Changes in the Local Environment Following Prefrontal Cortex Lesions. *Cerebral Cortex*, 15, 720–731.

-L-

- Lalonde, R., **2002**. The neurobiological basis of spontaneous alternation. *Neuroscience & Biobehavioral Reviews*, 26, 91–104.
- Lansink, C.S., Goltstein, P.M., Lankelma, J.V., Joosten, R.N.J.M.A., McNaughton, B.L., Pennartz, C.M.A., **2008**. Preferential Reactivation of Motivationally Relevant Information in the Ventral Striatum. *The Journal of Neuroscience*, 28, 6372–6382.
- Lansink, C.S., Goltstein, P.M., Lankelma, J.V., McNaughton, B.L., Pennartz, C.M.A., **2009**. Hippocampus Leads Ventral Striatum in Replay of Place-Reward Information. *PLoS Biology*, 7.
- Laroche, S., Davis, S., Jay, T.M., **2000**. Plasticity at hippocampal to prefrontal cortex synapses: dual roles in working memory and consolidation. *Hippocampus*, 10, 438–446.
- Larocque, J.J., Lewis-Peacock, J.A., Postle, B.R., **2014**. Multiple neural states of representation in short-term memory? It's a matter of attention. *Frontiers in Human Neuroscience*, 8, 5.
- Lavenex, P., Schenk, F., **1998**. Olfactory traces and spatial learning in rats. *Animal Behaviour*, 56, 1129–1136.
- Lavoie, A.M., Mizumori, S.J.Y., **1994**. Spatial, movement- and reward-sensitive discharge by medial ventral striatum neurons of rats. *Brain Research*, 638, 157–168.
- Leaton, R.N., **1965**. Exploratory behavior in rats with hippocampal lesions. *Journal of Comparative and Physiological Psychology*, 59, 325–330.
- Ledoux, A.-A., Boyer, P., Phillips, J.L., Smith, A., Labelle, A., Bohbot, V.D., **2014**. Structural hippocampal anomalies in a schizophrenia population correlate with navigation performance on a wayfinding task. *Frontiers in Behavioral Neuroscience*, 8, 88.
- Lee, A.K., Wilson, M.A., **2002**. Memory of Sequential Experience in the Hippocampus during Slow Wave Sleep. *Neuron*, 36, 1183–1194.
- Lee, D., Lin, B.-J., Lee, A.K., **2012**. Hippocampal Place Fields Emerge upon Single-Cell Manipulation of Excitability During Behavior. *Science*, 337, 849–853.
- Lee, H., Ghim, J.-W., Kim, H., Lee, D., Jung, M., **2012**. Hippocampal Neural Correlates for Values of Experienced Events. *The Journal of Neuroscience*, 32, 15053–15065.
- Lee, I., Rao, G., Knierim, J.J., **2004a**. A Double Dissociation between Hippocampal Subfields: Differential Time Course of CA3 and CA1 Place Cells for Processing Changed Environments. *Neuron*, 42, 803–815.
- Lee, I., Yoganarasimha, D., Rao, G., Knierim, J.J., **2004b**. Comparison of population coherence of place cells in hippocampal subfields CA1 and CA3. *Nature*, 430, 456–459.
- Lee, S.A., Spelke, E.S., **2010**. Two systems of spatial representation underlying navigation. *Experimental Brain Research*, 206, 179–188.
- Legault, M., Rompré, P.-P., Wise, R.A., **2000**. Chemical Stimulation of the Ventral Hippocampus Elevates Nucleus Accumbens Dopamine by Activating Dopaminergic Neurons of the Ventral Tegmental Area. *The Journal of Neuroscience*, 20, 1635–1642.
- Lenck-Santini, P.P., Muller, R.U., Save, E., Poucet, B., **2002**. Relationships between place cell firing fields and navigational decisions by rats. *The Journal of Neuroscience*, 22, 9035–9047.
- Lenck-Santini, P.P., Rivard, B., Muller, R.U., Poucet, B., **2005**. Study of CA1 place cell activity and exploratory behavior following spatial and nonspatial changes in the environment. *Hippocampus*, 15, 356–369.

- Lenck-Santini, P.P., Save, E., Poucet, B., **2001a**. Evidence for a relationship between place-cell spatial firing and spatial memory performance. *Hippocampus*, 11, 377–390.
- Lenck-Santini, P.P., Save, E., Poucet, B., **2001b**. Place-cell firing does not depend on the direction of turn in a Y-maze alternation task. *European Journal of Neuroscience*, 13, 1055–1058.
- Leutgeb, S., Leutgeb, J.K., Barnes, C.A., Moser, E.I., McNaughton, B.L., Moser, M.-B., **2005**. Independent Codes for Spatial and Episodic Memory in Hippocampal Neuronal Ensembles. *Science*, 309, 619–623.
- Leutgeb, S., Ragozzino, K.E., Mizumori, S.J.Y., **2000**. Convergence of head direction and place information in the CA1 region of hippocampus. *Neuroscience*, 100, 11–19.
- Lever, C., Burton, S., Jeewajee, A., O'Keefe, J., Burgess, N., **2009**. Boundary Vector Cells in the Subiculum of the Hippocampal Formation. *The Journal of Neuroscience*, 29, 9771–9777.
- Lever, C., Burgess, N., Cacucci, F., Hartley, T., O'Keefe, J., **2002b**. What can the hippocampal representation of environmental geometry tell us about Hebbian learning? *Biological Cybernetics*, 87, 356–372.
- Lever, C., Wills, T., Cacucci, F., Burgess, N., O'Keefe, J., **2002a**. Long-term plasticity in hippocampal place-cell representation of environmental geometry. *Nature*, 416, 90–94.
- Lindvall O., Bjirklund A., **1978**. Organization of catecholamine neurons in the rat central nervous system. In: Iversen L.L., Iversen S.D. and Snyder S.H. (Eds.), *Handbook of Psychopharmacology*. Plenum Press, New York Vol.9, pp.139–231.
- Lipska, B.K., Aultman, J.M., Verma, A., Weinberger, D.R., Moghaddam, B., **2002**. Neonatal Damage of the Ventral Hippocampus Impairs Working Memory in the Rat. *Neuropsychopharmacology*, 27, 47–54.
- Lisman, J.E., Grace, A.A., **2005**. The Hippocampal-VTA Loop: Controlling the Entry of Information into Long-Term Memory. *Neuron*, 46, 703–713.
- Lisman, J., Grace, A.A., Duzel, E., **2011**. A neoHebbian framework for episodic memory; role of dopamine-dependent late LTP. *Trends in Neurosciences*, Special Issue: Hippocampus and Memory, 34, 536–547.
- Lisman, J.E., Pi, H.J., Zhang, Y., Otmakhova, N.A., **2010**. A thalamo-hippocampal-ventral tegmental area loop may produce the positive feedback that underlies the psychotic break in schizophrenia. *Biological Psychiatry* 68, 17–24.
- Lomo T., **1966**. Frequency potentiation of excitatory synaptic activity in the dentate area of the hippocampal formation. *Acta Physiologica Scandinavica*, 68 Suppl 277: 128.
- Lorente de Nó, R., **1934**. Studies on the structure of the cerebral cortex. II. Continuation of the study of the ammonic system. *Journal für Psychologie und Neurologie*.
- Louie, K., Wilson, M.A., **2001**. Temporally Structured Replay of Awake Hippocampal Ensemble Activity during Rapid Eye Movement Sleep. *Neuron*, 29, 145–156.
- Loureiro, M., Cholvin, T., Lopez, J., Merienne, N., Latreche, A., Cosquer, B., Geiger, K., Kelche, C., Cassel, J.-C., Vasconcelos, A.P. de, **2012**. The Ventral Midline Thalamus (Reuniens and Rhomboid Nuclei) Contributes to the Persistence of Spatial Memory in Rats. *The Journal of Neuroscience*, 32, 9947–9959.
- Loureiro, M., Lecourtier, L., Engeln, M., Lopez, J., Cosquer, B., Geiger, K., Kelche, C., Cassel, J.-C., Pereira de Vasconcelos, A., **2011**. The ventral hippocampus is necessary for expressing a spatial memory. *Brain Structure and Function*, 217(1), 93–106.

- Luksys, G., Sandi, C., **2011**. Neural mechanisms and computations underlying stress effects on learning and memory. *Current Opinion in Neurobiology, Behavioural and cognitive neuroscience*, 21, 502–508.
- Luo, A.H., Tahsili-Fahadan, P., Wise, R.A., Lupica, C.R., Aston-Jones, G., **2011**. Linking Context with Reward: A Functional Circuit from Hippocampal CA3 to Ventral Tegmental Area. *Science*, 333, 353–357.
- Lynch, M.A., **2004**. Long-Term Potentiation and Memory. *Physiological Reviews*, 84, 87–136.

-M-

- Maaswinkel, H., Jarrard, L.E., Whishaw, I.Q., **1999**. Hippocampectomized rats are impaired in homing by path integration. *Hippocampus*, 9, 553–561.
- Maaswinkel, H., Whishaw, I.Q., **1999**. Homing with locale, taxon, and dead reckoning strategies by foraging rats: sensory hierarchy in spatial navigation. *Behavioural Brain Research*, 99, 143–152.
- MacDonald, C.J., Lepage, K.Q., Eden, U.T., Eichenbaum, H., **2011**. Hippocampal “Time Cells” Bridge the Gap in Memory for Discontiguous Events. *Neuron*, 71, 737–749.
- Malenka, R.C., Bear, M.F., **2004**. LTP and LTD: An Embarrassment of Riches. *Neuron*, 44, 5–21.
- Malhotra, S., Cross, R.W.A., van der Meer, M.A.A., **2012**. Theta phase precession beyond the hippocampus. *Reviews in the Neurosciences*, 23, 39–65.
- Mällo, T., Alftoa, A., Köiv, K., Tõnissaar, M., Eller, M., Harro, J., **2007**. Rats with persistently low or high exploratory activity: Behaviour in tests of anxiety and depression, and extracellular levels of dopamine. *Behavioural Brain Research*, 177, 269–281.
- Mannella, F., Gurney, K., Baldassarre, G., **2013**. The nucleus accumbens as a nexus between values and goals in goal-directed behavior: a review and a new hypothesis. *Front. Behav Neurosci*, 7, 135.
- Markram, H., Gerstner, W., Sjöström, P.J., **2012**. Spike-timing-dependent plasticity: a comprehensive overview. *Frontiers in Synaptic Neuroscience*, 4. Editorial article.
- Markram, H., Lübke, J., Frotscher, M., Sakmann, B., **1997**. Regulation of Synaptic Efficacy by Coincidence of Postsynaptic APs and EPSPs. *Science*, 275, 213–215.
- Markus, E.J., Barnes, C.A., McNaughton, B.L., Gladden, V.L., Skaggs, W.E., **1994**. Spatial information content and reliability of hippocampal CA1 neurons: Effects of visual input. *Hippocampus*, 4, 410–421.
- Markus, E., Qin, Y., Leonard, B., Skaggs, W., McNaughton, B., Barnes, C., **1995**. Interactions between location and task affect the spatial and directional firing of hippocampal neurons. *The Journal of Neuroscience*, 15, 7079–7094.
- Marr, D., **1971**. Simple memory: a theory for archicortex. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 262, 23–81.
- Martig, A.K., Mizumori, S.J.Y., **2011**. Ventral tegmental area disruption selectively affects CA1/CA2 but not CA3 place fields during a differential reward working memory task. *Hippocampus*, 21, 172–184.
- Martinet, L.-E., Sheynikhovich, D., Benchenane, K., Arleo, A., **2011**. Spatial learning and action planning in a prefrontal cortical network model. *PLoS Computational Biology*, 7, e1002045.

- Masquelier, T., Guyonneau, R., Thorpe, S.J., **2009**. Competitive STDP-based spike pattern learning. *Neural Computation*, 21, 1259–1276.
- Masquelier, T., Thorpe, S.J., **2007**. Unsupervised learning of visual features through spike timing dependent plasticity. *PLoS Computational Biology*, 3, e31.
- Matthews, D.B., Best, P.J., **1997**. Evidence for the flexible use of spatial knowledge in the rat. *Psychobiology*, 25, 294–302.
- Maurer, R., Derivaz, V., **2000**. Rats in a transparent morris water maze use elemental and configural geometry of landmarks as well as distance to the pool wall. *Spatial Cognition and Computation*, 2, 135–156.
- McClelland, D.C., Koestner, R., Weinberger, J., **1989**. How do self-attributed and implicit motives differ? *Psychological Review*, 96, 690–702.
- McClelland, J.L., McNaughton, B.L., O'Reilly, R.C., **1995**. Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, 102, 419–457.
- McDonald, R.J., White, N.M., **1993**. A triple dissociation of memory systems: hippocampus, amygdala, and dorsal striatum. *Behavioral Neuroscience*, 107, 3–22.
- McDonald, R.J., White, N.M., **2013**. A triple dissociation of memory systems: Hippocampus, amygdala, and dorsal striatum. *Behavioral Neuroscience*, 127, 835–853.
- McGaugh, J.L., **2004**. The Amygdala Modulates the Consolidation of Memories of Emotionally Arousing Experiences. *Annual Review of Neuroscience*, 27, 1–28.
- McGregor, A., Hayward, A.J., Pearce, J.M., Good, M.A., **2004**. Hippocampal lesions disrupt navigation based on the shape of the environment. *Behavioral Neuroscience*, 118, 1011–1021.
- McHugh, T.J., Blum, K.I., Tsien, J.Z., Tonegawa, S., Wilson, M.A., **1996**. Impaired Hippocampal Representation of Space in CA1-Specific NMDAR1 Knockout Mice. *Cell*, 87, 1339–1349.
- McKenzie, S., Eichenbaum, H., **2011**. Consolidation and Reconsolidation: Two Lives of Memories? *Neuron*, 71, 224–233.
- McNaughton, B.L., Battaglia, F.P., Jensen, O., Moser, E.I., Moser, M.-B., **2006**. Path integration and the neural basis of the “cognitive map.” *Nature Reviews Neuroscience*, 7, 663–678.
- McNaughton, D.B.L., Barnes, C.A., O'Keefe, J., **1983**. The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely-moving rats. *Experimental Brain Research*, 52, 41–49.
- Meck, W.H., MacDonald, C.J., **2007**. Amygdala inactivation reverses fear's ability to impair divided attention and make time stand still. *Behavioral Neuroscience*, 121, 707–720.
- Mehta, M.R., Barnes, C.A., McNaughton, B.L., **1997**. Experience-dependent, asymmetric expansion of hippocampal place fields. *Proceedings of the National Academy of Sciences*, 94, 8918–8921.
- Mehta, M.R., Quirk, M.C., Wilson, M.A., **2000**. Experience-Dependent Asymmetric Shape of Hippocampal Receptive Fields. *Neuron*, 25, 707–715.
- Mehta, M.R., Lee, A.K., Wilson, M.A., **2002**. Role of experience and oscillations in transforming a rate code into a temporal code. *Nature*, 417, 741–746.
- Mei, B., Li, F., Gu, Y., Cui, Z., Tsien, J.Z., **2011**. NMDA Receptors Are Not Required for Pattern Completion During Associative Memory Recall. *PLoS ONE*, 6, e19326.
- Meliza, C.D., Dan, Y., **2006**. Receptive-field modification in rat visual cortex induced by paired visual stimulation and single-cell spiking. *Neuron*, 49, 183–189.

- Miles, R., Wong, R.K., **1986**. Excitatory synaptic interactions between CA3 neurones in the guinea-pig hippocampus. *The Journal of Physiology*, 373, 397–418.
- Miller, G.A., **1956**. The magical number seven, plus or minus two: some limits on our capacity for processing information. *Psychological Review*, 63, 81–97.
- Mittelstaedt, H., Mittelstaedt, M.-L., **1982**. Homing by Path Integration, in: Papi, P.D.F., Wallraff, D.H.G. (Eds.), *Avian Navigation, Proceedings in Life Sciences*. Springer Berlin Heidelberg, pp. 290–297.
- Mizumori, S.J., Williams, J.D., **1993**. Directionally selective mnemonic properties of neurons in the lateral dorsal nucleus of the thalamus of rats. *The Journal of Neuroscience*, 13, 4015–4028.
- Mizuseki, K., Royer, S., Diba, K., Buzsaki, G., **2012**. Activity Dynamics and Behavioral Correlates of CA3 and CA1 Hippocampal Pyramidal Neurons. *Hippocampus*, 22, 1659–1680. doi:10.1002/hipo.22002
- Monaco, J.D., Rao, G., Roth, E.D., Knierim, J.J., **2014**. Attentive scanning behavior drives one-trial potentiation of hippocampal place fields. *Nature Neuroscience*, 17, 725–731.
- Mongillo, G., Barak, O., Tsodyks, M., **2008**. Synaptic Theory of Working Memory. *Science*, 319, 1543–1546.
- Morris JC, Storandt M, Miller J, et al, **2001**. Mild cognitive impairment represents early-stage alzheimer disease. *Archives of Neurology*, 58, 397–405.
- Morris, R.G.M., **1981**. Spatial localization does not require the presence of local cues. *Learning and Motivation*, 12, 239–260.
- Morris, R.G.M., Anderson, E., Lynch, G.S., Baudry, M., **1986**. Selective impairment of learning and blockade of long-term potentiation by an N-methyl-D-aspartate receptor antagonist, AP5. *Nature*, 319, 774–776.
- Morris, R.G.M., Garrud, P., Rawlins, J.N.P., O’Keefe, J., **1982**. Place navigation impaired in rats with hippocampal lesions. *Nature*, 297, 681–683.
- Morris, R.G.M., Schenk, F., Tweedie, F., Jarrard, L.E., **1990**. Ibotenate Lesions of Hippocampus and/or Subiculum: Dissociating Components of Allocentric Spatial Learning. *European Journal of Neuroscience*, 2, 1016–1028.
- Moscovitch, M., Nadel, L., Winocur, G., Gilboa, A., Rosenbaum, R.S., **2006**. The cognitive neuroscience of remote episodic, semantic and spatial memory. *Current Opinion in Neurobiology, Cognitive neuroscience*, 16, 179–190.
- Moser, E.I., Kropff, E., Moser, M.-B., **2008**. Place Cells, Grid Cells, and the Brain’s Spatial Representation System. *Annual Review of Neuroscience*, 31, 69–89.
- Moser, E., Moser, M.B., Andersen, P., **1993**. Spatial learning impairment parallels the magnitude of dorsal hippocampal lesions, but is hardly present following ventral lesions. *The Journal of Neuroscience*, 13, 3916–3925.
- Moser, M.-B., Moser, E.I., **1998**. Functional differentiation in the hippocampus. *Hippocampus*, 8, 608–619.
- Moser, M.B., Moser, E.I., Forrest, E., Andersen, P., Morris, R.G., **1995**. Spatial learning with a minislab in the dorsal hippocampus. *Proceedings of the National Academy of Sciences*, 92, 9697–9701.
- Muenzinger, K.F., **1938**. Vicarious Trial and Error at a Point of Choice: I. A General Survey of its Relation to Learning Efficiency. *The Pedagogical Seminary and Journal of Genetic Psychology*, 53, 75–86.
- Muir, G.M., Taube, J.S., **2004**. Head direction cell activity and behavior in a navigation task requiring a cognitive mapping strategy. *Behavioural Brain Research*, 153, 249–253.

- Müller, M., Wehner, R., **1988**. Path integration in desert ants, *Cataglyphis fortis*. *Proceedings of the National Academy of Sciences*, 85, 5287–5290.
- Muller, R., Kubie, J., **1987**. The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *The Journal of Neuroscience*, 7, 1951–1968.
- Muller, R.U., Bostock, E., Taube, J.S., Kubie, J.L., **1994**. On the directional firing properties of hippocampal place cells. *The Journal of Neuroscience*, 14, 7235–7251.
- Muller, R.U., Kubie, J.L., Ranck, J.B., **1987**. Spatial firing patterns of hippocampal complex-spike cells in a fixed environment. *The Journal of Neuroscience*, 7, 1935–1950.

-N-

- Naber, P.A., Lopes da Silva, F.H., Witter, M.P., **2001**. Reciprocal connections between the entorhinal cortex and hippocampal fields CA1 and the subiculum are in register with the projections from CA1 to the subiculum. *Hippocampus*, 11, 99–104.
- Nadel, L., Moscovitch, M., **1997**. Memory consolidation, retrograde amnesia and the hippocampal complex. *Current Opinion in Neurobiology*, 7, 217–227.
- Naqshbandi, M., Feeney, M.C., McKenzie, T.L.B., Roberts, W.A., **2007**. Testing for episodic-like memory in rats in the absence of time of day cues: replication of Babb and Crystal. *Behavioural Processes*, 74, 217–225.
- Nemati, F., Whishaw, I.Q., **2007**. The point of entry contributes to the organization of exploratory behavior of rats on an open field: An example of spontaneous episodic memory. *Behavioural Brain Research*, 182, 119–128.
- Nerad, L., Liu, P., Bilkey, D.K., **2009**. Bilateral NMDA lesions centered on the postrhinal cortex have minimal effects on hippocampal place cell firing. *Hippocampus*, 19, 221–227.
- Nestor, P.J., Scheltens, P., Hodges, J.R., **2004**. Advances in the early detection of Alzheimer’s disease. *Nature Medicine*. 10 Suppl, S34–41.
- Nyberg, L., McIntosh, A.R., Cabeza, R., Habib, R., Houle, S., Tulving, E., **1996**. General and specific brain regions involved in encoding and retrieval of events: what, where, and when. *Proceedings of the National Academy of Sciences*, 93, 11280–11285.

-O-

- Oberauer, K., **2002**. Access to information in working memory: Exploring the focus of attention. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 28, 411–421.
- O’Keefe, J., **1976**. Place units in the hippocampus of the freely moving rat. *Experimental Neurology*, 51, 78–109.
- O’Keefe, J., **1979**. A review of the hippocampal place cells. *Progress in Neurobiology*, 13, 419–439.
- O’Keefe, J., **1999**. Do hippocampal pyramidal cells signal non-spatial as well as spatial information? *Hippocampus*, 9, 352–364.
- O’Keefe, J., Burgess, N., **1996**. Geometric determinants of the place fields of hippocampal neurons. *Nature*, 381, 425–428.

- O'Keefe, J., Conway, D.H., **1978**. Hippocampal place units in the freely moving rat: why they fire where they fire. *Experimental Brain Research*, 31, 573–590.
- O'Keefe, J., Dostrovsky, J., **1971**. The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Research*, 34, 171–175.
- O'Keefe, J., Nadel, L., **1978**. *The hippocampus as a cognitive map*. Clarendon Press Oxford.
- O'Keefe, J., Recce, M.L., **1993**. Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus*, 3, 317–330.
- O'Keefe, J., Speakman, A., **1987**. Single unit activity in the rat hippocampus during a spatial memory task. *Experimental Brain Research*, 68, 1–27.
- Olypher, A.V., Lánský, P., Fenton, A.A., **2002**. Properties of the extra-positional signal in hippocampal place cell discharge derived from the overdispersion in location-specific firing. *Neuroscience*, 111, 553–566.
- Olsen, R.K., Moses, S.N., Riggs, L., Ryan, J.D., **2012**. The hippocampus supports multiple cognitive processes through relational binding and comparison. *Frontiers in Human Neuroscience*, 6, 146.
- Olton, D.S., Becker, J.T., Handelmann, G.E., **1979**. A re-examination of the role of hippocampus in working memory. *Behavioral and Brain Sciences*, 2, 352–365.
- Olton, D.S., Collison, C., Werz, M.A., **1977**. Spatial memory and radial arm maze performance of rats. *Learning and Motivation*, 8, 289–314.
- Olton, D.S., Samuelson, R.J., **1976**. Remembrance of places passed: Spatial memory in rats. *Journal of Experimental Psychology: Animal Behavior Processes*, 2, 97–116.
- O'Neill, J., Pleydell-Bouverie, B., Dupret, D., Csicsvari, J., **2010**. Play it again: reactivation of waking experience and memory. *Trends in Neurosciences*, 33, 220–229.

-P-

- Packard, M.G., McGaugh, J.L., **1992**. Double dissociation of fornix and caudate nucleus lesions on acquisition of two water maze tasks: Further evidence for multiple memory systems. *Behavioral Neuroscience*, 106, 439–446.
- Packard, M.G., McGaugh, J.L., **1996**. Inactivation of Hippocampus or Caudate Nucleus with Lidocaine Differentially Affects Expression of Place and Response Learning. *Neurobiology of Learning and Memory*, 65, 65–72.
- Padoa-Schioppa, C., **2007**. Orbitofrontal cortex and the computation of economic value. *Annals of the New York Academy of Sciences*, 1121, 232–253.
- Papez, J.W., **1937**. A proposed mechanism of emotion. *Archives of Neurology & Psychiatry*, 38, 725–743.
- Park, E., Dvorak, D., Fenton, A.A., **2011**. Ensemble Place Codes in Hippocampus: CA1, CA3, and Dentate Gyrus Place Cells Have Multiple Place Fields in Large Environments. *PLoS ONE*, 6, e22349.
- Parron, C., Poucet, B., Save, E., **2006**. Cooperation between the hippocampus and the entorhinal cortex in spatial memory: A disconnection study. *Behavioural Brain Research*, 170, 99–109.
- Parron, C., Save, E., **2004**. Evidence for entorhinal and parietal cortices involvement in path integration in the rat. *Experimental Brain Research*, 159, 349–359.

- Passino, E., Middei, S., Restivo, L., Bertaina-Anglade, V., Ammassari-Teule, M., **2002**. Genetic approach to variability of memory systems: Analysis of place vs. response learning and Fos-related expression in hippocampal and striatal areas of C57BL/6 and DBA/2 mice. *Hippocampus*, 12, 63–75.
- Passot, J.-B., Sheynikhovich, D., Duvelle, É., Arleo, A., **2012**. Contribution of Cerebellar Sensorimotor Adaptation to Hippocampal Spatial Memory. *PLoS ONE*, 7, e32560.
- Pause, B.M., Zlomuzica, A., Kinugawa, K., Mariani, J., Pietrowsky, R., Dere, E., **2013**. Perspectives on episodic-like and episodic memory. *Frontiers in Behavioral Neuroscience*, 7, 33.
- Pawlak, C.R., Ho, Y.-J., Schwarting, R.K.W., **2008**. Animal models of human psychopathology based on individual differences in novelty-seeking and anxiety. *Neuroscience & Biobehavioral Reviews*, Special section: *The European Workshop in Imagery and Cognition: Neurocognition and Visual Imagery*, 32, 1544–1568.
- Pearce, J.M., Roberts, A.D.L., Good, M., **1998**. Hippocampal lesions disrupt navigation based on cognitive maps but not heading vectors. *Nature*, 396, 75–77.
- Perea, G., Navarrete, M., Araque, A., **2009**. Tripartite synapses: astrocytes process and control synaptic information. *Trends in Neurosciences*, 32, 421–431.
- Peters, J., Büchel, C., **2010**. Neural representations of subjective reward value. *Behavioural Brain Research*, 213, 135–141.
- Pfeiffer, B.E., Foster, D.J., **2013**. Hippocampal place-cell sequences depict future paths to remembered goals. *Nature*, 497, 74–79.
- Pinker, S., Jackendoff, R., **2005**. The faculty of language: what's special about it? *Cognition*, 95, 201–236.
- Pitkänen, A., Pikkarainen, M., Nurminen, N., Ylinen, A., **2000**. Reciprocal Connections between the Amygdala and the Hippocampal Formation, Perirhinal Cortex, and Postrhinal Cortex in Rat: A Review. *Annals of the New York Academy of Sciences*, 911, 369–391.
- Poucet, B., **1993**. Spatial cognitive maps in animals: New hypotheses on their structure and neural mechanisms. *Psychological Review*, 100, 163–182.
- Poucet, B., **1997**. Searching for spatial unit firing in the prelimbic area of the rat medial prefrontal cortex. *Behavioural Brain Research*, 84, 151–159.
- Poucet, B., Benhamou, S., **1997**. The neuropsychology of spatial cognition in the rat. *Critical Reviews in Neurobiology*, 11, 101–120.
- Poucet, B., Durup, M., Thinus-Blanc, C., **1988**. Short-term and long-term habituation of exploration in rats, hamsters and gerbils. *Behavioural Processes*, 16, 203–211.
- Poucet, B., Herrmann, T., **2001**. Exploratory patterns of rats on a complex maze provide evidence for topological coding. *Behavioural Processes*, 53, 155–162.
- Poucet, B., Hok, V., Sargolini, F., Save, E., **2012**. Stability and variability of place cell activity during behavior: Functional implications for dynamic coding of spatial information. *Journal of Physiology-Paris*, Neuronal Ensemble Recordings in Integrative Neuroscience, 106, 62–71.
- Poucet, B., Lenck-Santini, P.P., Hok, V., Save, E., Banquet, J.P., Gaussier, P., Muller, R.U., **2004**. Spatial navigation and hippocampal place cell firing: the problem of goal encoding. *Reviews in the Neurosciences*, 15, 89–107.
- Poucet, B., Sargolini, F., Song, E.Y., Hangya, B., Fox, S., Muller, R.U., **2014**. Independence of landmark and self-motion-guided navigation: a different role for grid cells. *Philosophical Transactions of the Royal Society of London. B. Biological Sciences*, 369, 20130370.

- Poucet, B., Save, E., **2005**. Attractors in Memory. *Science*, 308, 799–800.
- Poucet, B., Save, E., Lenck-Santini, P.P., **2000**. Sensory and memory properties of hippocampal place cells. *Reviews in the Neurosciences*, 11, 95–111.
- Pratt, W.E., Mizumori, S.J., **1998**. Characteristics of basolateral amygdala neuronal firing on a spatial memory task involving differential reward. *Behavioral Neuroscience*, 112, 554–570.
- Preuss, T.M., **1995**. Do Rats Have Prefrontal Cortex? The Rose-Woolsey-Akert Program Reconsidered. *Journal of Cognitive Neuroscience* 7, 1–24.
- Puryear, C.B., Kim, M.J., Mizumori, S.J.Y., **2010**. Conjunctive encoding of movement and reward by ventral tegmental area neurons in the freely navigating rodent. *Behavioral Neuroscience*, 124, 234–247.
- Pynn, L.K., DeSouza, J.F.X., **2013**. The function of efference copy signals: Implications for symptoms of schizophrenia. *Vision Research*, 76, 124–133.

-Q-

- Quirk, G.J., Muller, R.U., Kubie, J.L., **1990**. The firing of hippocampal place cells in the dark depends on the rat's recent experience. *The Journal of Neuroscience*, 10, 2008–2017.
- Quiroga, R.Q., **2012**. Concept cells: the building blocks of declarative memory functions. *Nature Reviews Neuroscience*, 13, 587–597.
- Quiroga, R.Q., Kreiman, G., **2010**. Measuring sparseness in the brain: Comment on Bowers (2009). *Psychological Review*, 117, 291–297.
- Quiroga, R.Q., Kreiman, G., Koch, C., Fried, I., **2008**. Sparse but not “Grandmother-cell” coding in the medial temporal lobe. *Trends in Cognitive Sciences*, 12, 87–91.
- Quiroga, R.Q., Reddy, L., Kreiman, G., Koch, C., Fried, I., **2005**. Invariant visual representation by single neurons in the human brain. *Nature*, 435, 1102–1107.

-R-

- Ramirez-Amaya, V., **2007**. Molecular Mechanisms of Synaptic Plasticity Underlying Long-Term Memory Formation, in: Bermúdez-Rattoni, F. (Ed.), *Neural Plasticity and Memory: From Genes to Brain Imaging*, *Frontiers in Neuroscience*. CRC Press, Boca Raton (FL).
- Ramirez, S., Liu, X., Lin, P.-A., Suh, J., Pignatelli, M., Redondo, R.L., Ryan, T.J., Tonegawa, S., **2013**. Creating a False Memory in the Hippocampus. *Science*, 341, 387–391.
- Ramnerö, J., Törneke, N., **2014**. On Having a Goal: Goals as Representations or Behavior. *The Psychological Record*, 1–11.
- Ramón y Cajal, S., **1894**. The Croonian Lecture: La Fine Structure des Centres Nerveux. *Royal Society of London Proceedings Series I*, 55, 444–468.
- Ramos, J.M.J., **2010**. Preserved learning about allocentric cues but impaired flexible memory expression in rats with hippocampal lesions. *Neurobiology of Learning and Memory*, 93, 506–514.

- Ramus, S.J., Davis, J.B., Donahue, R.J., Disenza, C.B., Waite, A.A., **2007**. Interactions between the Orbitofrontal Cortex and the Hippocampal Memory System during the Storage of Long-Term Memory. *Annals of the New York Academy of Sciences*, 1121, 216–231.
- Rasmussen, M., Barnes, C.A., McNaughton, B.L., **1989**. A systematic test of cognitive mapping, working-memory, and temporal discontinuity theories of hippocampal function. *Psychobiology*, 17, 335–348.
- Ravassard, P., Kees, A., Willers, B., Ho, D., Aharoni, D.A., Cushman, J., Aghajan, Z.M., Mehta, M.R., **2013**. Multisensory Control of Hippocampal Spatiotemporal Selectivity. *Science*, 340(6138), 1342–1346.
- Rawlins, J.N.P., **1985**. Associations across time: The hippocampus as a temporary memory store. *Behavioral and Brain Sciences*, 8, 479–497.
- Redish, A.D., **1999**. Beyond the Cognitive Map: From Place Cells to Episodic Memory. *MIT Press*.
- Redish, A.D., **2013**. The mind within the Brain: How we make decisions and how those decisions go wrong. *Oxford University Press*.
- Redish, A.D., Mizumori, S.J.Y., **2014**. Memory and decision making. *Neurobiology of Learning and Memory*. In Press.
- Redondo, R.L., Kim, J., Arons, A.L., Ramirez, S., Liu, X., Tonegawa, S., **2014**. Bidirectional switch of the valence associated with a hippocampal contextual memory engram. *Nature*, 513, 426–430.
- Reep, R.L., Corwin, J.V., King, V., **1996**. Neuronal connections of orbital cortex in rats: topography of cortical and thalamic afferents. *Experimental Brain Research*, 111, 215–232.
- Reichelt, A.C., Lin, T.E., Harrison, J.J., Honey, R.C., Good, M.A., **2011**. Differential role of the hippocampus in response-outcome and context-outcome learning: Evidence from selective satiation procedures. *Neurobiology of Learning and Memory*, 96, 248–253.
- Renaudineau, S., **2008**. Flexibilité fonctionnelle des cellules de lieu et mémoire spatiale: étude des mécanismes d'adaptation et des aspects moléculaires. Université d'Aix-Marseille I. Doctoral Thesis.
- Renaudineau, S., Poucet, B., Save, E., **2007**. Flexible use of proximal objects and distal cues by hippocampal place cells. *Hippocampus*, 17, 381–395.
- Retailleau, A., Dejean, C., Fourneaux, B., Leinekugel, X., Boraud, T., **2013**. Why am I lost without dopamine? Effects of 6-OHDA lesion on the encoding of reward and decision process in CA3. *Neurobiology of Disease*, 59, 151–164.
- Río, C. del, Oliveras, I., Cañete, T., Blázquez, G., Tobeña, A., Fernández-Teruel, A., **2014**. Genetic Rat Models of Schizophrenia-Relevant Symptoms. *World Journal of Neuroscience*, 04, 261–278.
- Rinaldi, A., Oliverio, A., Mele, A., **2012**. Spatial memory, plasticity and nucleus accumbens. *Reviews in the Neurosciences*, 23, 527–541.
- Rivard, B., Li, Y., Lenck-Santini, P.-P., Poucet, B., Muller, R.U., **2004**. Representation of Objects in Space by Two Classes of Hippocampal Pyramidal Cells. *The Journal of General Physiology*, 124, 9–25.
- Roberts, W.A., **2012**. Evidence for future cognition in animals. *Learning and Motivation*, 43, 169–180.
- Roberts, W.W., Dember, W.N., Brodwick, M., **1962**. Alternation and exploration in rats with hippocampal lesions. *Journal of Comparative and Physiological Psychology*, 55, 695–700.
- Rocheffort, C., Arabo, A., André, M., Poucet, B., Save, E., Rondi-Reig, L., **2011**. Cerebellum shapes hippocampal spatial code. *Science* 334, 385–389.

- Roesch, M.R., Calu, D.J., Schoenbaum, G., **2007**. Dopamine neurons encode the better option in rats deciding between differently delayed or sized rewards. *Nature Neuroscience*, 10, 1615–1624.
- Roesch, M.R., Singh, T., Brown, P.L., Mullins, S.E., Schoenbaum, G., **2009**. Ventral striatal neurons encode the value of the chosen action in rats deciding between differently delayed or sized rewards. *The Journal of Neuroscience*, 29, 13365–13376.
- Rolls, E.T., **1999**. Spatial view cells and the representation of place in the primate hippocampus. *Hippocampus* 9, 467–480.
- Rondi-Reig, L., Petit, G.H., Tobin, C., Tonegawa, S., Mariani, J., Berthoz, A., **2006**. Impaired Sequential Egocentric and Allocentric Memories in Forebrain-Specific–NMDA Receptor Knock-Out Mice during a New Task Dissociating Strategies of Navigation. *The Journal of Neuroscience*, 26, 4071–4081.
- Rossier, J., Kaminsky, Y., Schenk, F., Bures, J., **2000**. The place preference task: a new tool for studying the relation between behavior and place cell activity in rats. *Behavioral Neuroscience*, 114, 273–284.
- Rossier, J., Haeberli, C., Schenk, F., **2000**. Auditory cues support place navigation in rats when associated with a visual cue. *Behavioural Brain Research*, 117, 209–214.
- Rossier, J., Schenk, F., **2003**. Olfactory and/or visual cues for spatial navigation through ontogeny: olfactory cues enable the use of visual cues. *Behavioral Neuroscience*, 117, 412–425.

-S-

- Sargolini, F., Fyhn, M., Hafting, T., McNaughton, B.L., Witter, M.P., Moser, M.-B., Moser, E.I., **2006**. Conjunctive Representation of Position, Direction, and Velocity in Entorhinal Cortex. *Science*, 312, 758–762.
- Save, E., Cressant, A., Thinus-Blanc, C., Poucet, B., **1998**. Spatial Firing of Hippocampal Place Cells in Blind Rats. *The Journal of Neuroscience*, 18, 1818–1826.
- Save, E., Guazzelli, A., Poucet, B., **2001**. Dissociation of the effects of bilateral lesions of the dorsal hippocampus and parietal cortex on path integration in the rat. *Behavioral Neuroscience*, 115, 1212–1223.
- Save, E., Nerad, L., Poucet, B., **2000**. Contribution of multiple sensory information to place field stability in hippocampal place cells. *Hippocampus*, 10, 64–76.
- Save, E., Poucet, B., Foreman, N., Buhot, M.-C., **1992**. Object exploration and reactions to spatial and nonspatial changes in hooded rats following damage to parietal cortex or hippocampal formation. *Behavioral Neuroscience*, 106, 447–456.
- Save, E., Poucet, B., **2008**. Do place cells guide spatial behaviors ? *SJY Mizumori (Ed.), Hippocampal place fields: Relevance to learning and memory* (pp. 138-149). Oxford University Press.
- Savelli, F., Yoganarasimha, D., Knierim, J.J., **2008**. Influence of boundary removal on the spatial representations of the medial entorhinal cortex. *Hippocampus*, 18, 1270–1282.
- Scatton, B., Simon, H., Le Moal, M., Bischoff, S., **1980**. Origin of dopaminergic innervation of the rat hippocampal formation. *Neuroscience Letters*, 18, 125–131.
- Schacter, D.L., Eich, J.E., Tulving, E., **1978**. Richard Semon’s theory of memory. *Journal of Verbal Learning and Verbal Behavior*, 17, 721–743.

- Schmidt, B., Papale, A., Redish, A.D., Markus, E.J., **2013**. Conflict between place and response navigation strategies: effects on vicarious trial and error (VTE) behaviors. *Learning & Memory*, 20, 130–138.
- Schultz, W., **2007**. Multiple Dopamine Functions at Different Time Courses. *Annual Review of Neuroscience*, 30, 259–288.
- Schultz, W., Dayan, P., Montague, P.R., **1997**. A Neural Substrate of Prediction and Reward. *Science*, 275, 1593–1599.
- Scoville, W.B., Milner, B., **1957**. Loss of recent memory after bilateral hippocampal lesions. *Journal of Neurology, Neurosurgery & Psychiatry*, 20, 11–21.
- Shah, A., Jhavar, S.S., Goel, A., **2012**. Analysis of the anatomy of the Papez circuit and adjoining limbic system by fiber dissection techniques. *Journal of Clinical Neuroscience*, 19, 289–298.
- Shapiro, M.L., Tanila, H., Eichenbaum, H., **1997**. Cues that hippocampal place cells encode: dynamic and hierarchical representation of local and distal stimuli. *Hippocampus*, 7, 624–642.
- Sharp, P.E., **1997**. Subicular cells generate similar spatial firing patterns in two geometrically and visually distinctive environments: Comparison with hippocampal place cells. *Behavioural Brain Research*, 85, 71–92.
- Sharp, P.E., **2006**. Subicular place cells generate the same “map” for different environments: Comparison with hippocampal cells. *Behavioural Brain Research*, 174, 206–214.
- Sharp, P.E., Blair, H.T., Etkin, D., Tzanetos, D.B., **1995**. Influences of vestibular and visual motion information on the spatial firing patterns of hippocampal place cells. *The Journal of Neuroscience*, 15, 173–189.
- Sharp, P.E., Green, C., **1994**. Spatial correlates of firing patterns of single cells in the subiculum of the freely moving rat. *The Journal of Neuroscience*, 14, 2339–2356.
- Sharp, P.E., Koester, K., **2008**. Lesions of the mammillary body region severely disrupt the cortical head direction, but not place cell signal. *Hippocampus*, 18, 766–784.
- Sheynikhovich, D., Chavarriaga, R., Strösslin, T., Arleo, A., Gerstner, W., **2009**. Is there a geometric module for spatial orientation? Insights from a rodent navigation model. *Psychological Review*, 116, 540–566.
- Shohamy, D., Adcock, R.A., **2010**. Dopamine and adaptive memory. *Trends in Cognitive Sciences*, 14, 464–472.
- Siegel, J.J., Neunuebel, J.P., Knierim, J.J., **2008**. Dominance of the Proximal Coordinate Frame in Determining the Locations of Hippocampal Place Cell Activity During Navigation. *Journal of Neurophysiology*, 99, 60–76.
- Singer, W., **2005**. How does the brain know when it is right?, in: Changeux, J.-P., Damasio, A.R., Singer, W., Christen, Y. (Eds.), *Neurobiology of Human Values*. Springer-Verlag, Berlin/Heidelberg, pp. 125–135.
- Skaggs, W.E., McNaughton, B.L., Wilson, M.A., Barnes, C.A., **1996**. Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus*, 6, 149–172.
- Skov-Rackette, S.I., Shettleworth, S.J., **2005**. What Do Rats Learn About the Geometry of Object Arrays? Tests With Exploratory Behavior. *Journal of Experimental Psychology: Animal Behavior Processes*, 31, 142–154.
- Solstad, T., Boccara, C.N., Kropff, E., Moser, M.-B., Moser, E.I., **2008**. Representation of Geometric Borders in the Entorhinal Cortex. *Science*, 322, 1865–1868.

- Song, S., Miller, K.D., Abbott, L.F., **2000**. Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. *Nature Neuroscience*, 3, 919–926.
- Speakman, A., O’Keefe, J., **1990**. Hippocampal Complex Spike Cells do not Change Their Place Fields if the Goal is Moved Within a Cue Controlled Environment. *European Journal of Neuroscience*, 2, 544–555.
- Spence, K.W., Lippitt, R., **1946**. An experimental test of the sign-gestalt theory of trial and error learning. *Journal of Experimental Psychology*, 36, 491–502.
- Spiers, H.J., Bendor, D., **2014**. Enhance, delete, incept: Manipulating hippocampus-dependent memories. *Brain Research Bulletin*, 105, 2–7.
- Squire, L.R., **2004**. Memory systems of the brain: a brief history and current perspective. *Neurobiology of Learning and Memory*, 82, 171–177.
- Squire, L.R., **2009**. Memory and Brain Systems: 1969–2009. *The Journal of Neuroscience*, 29, 12711–12716.
- Squire, L.R., Alvarez, P., **1995**. Retrograde amnesia and memory consolidation: a neurobiological perspective. *Current Opinion in Neurobiology*, 5, 169–177.
- Squire, L.R., Knowlton, B.J., **1995**. Memory, hippocampus, and brain systems, in: *The Cognitive Neurosciences*. The MIT Press, Cambridge, MA, USA, 825–837.
- Squire, L.R., Zola, S.M., **1996**. Structure and function of declarative and nondeclarative memory systems. *Proceedings of the National Academy of Sciences*, 93, 13515–13522.
- Stackman, R.W., Clark, A.S., Taube, J.S., **2002**. Hippocampal spatial representations require vestibular input. *Hippocampus*, 12, 291–303.
- Stackman, R.W., Taube, J.S., **1998**. Firing properties of rat lateral mammillary single units: head direction, head pitch, and angular head velocity. *The Journal of Neuroscience*, 18, 9020–9037.
- Stanton, P.K., Sejnowski, T.J., **1989**. Associative long-term depression in the hippocampus induced by hebbian covariance. *Nature*, 339, 215–218.
- Steffenach, H.-A., Witter, M., Moser, M.-B., Moser, E.I., **2005**. Spatial Memory in the Rat Requires the Dorsolateral Band of the Entorhinal Cortex. *Neuron*, 45, 301–313.
- Steiner, A.P., Redish, A.D., **2014**. Behavioral and neurophysiological correlates of regret in rat decision-making on a neuroeconomic task. *Nature Neuroscience*, 17, 995–1002.
- Stensola, H., Stensola, T., Solstad, T., Frøland, K., Moser, M.-B., Moser, E.I., **2012**. The entorhinal grid map is discretized. *Nature*, 492, 72–78.
- Stevenson, E.L., Caldwell, H.K., **2014**. Lesions to the CA2 region of the hippocampus impair social memory in mice. *European Journal of Neuroscience*. In Press.
- Storandt, M., **2008**. Cognitive Deficits in the Early Stages of Alzheimer’s Disease. *Current Directions in Psychological Science*, 17, 198–202.
- Strange, B.A., Witter, M.P., Lein, E.S., Moser, E.I., **2014**. Functional organization of the hippocampal longitudinal axis. *Nature Reviews Neuroscience*, 15, 655–669.
- Stuchlik, A., **2014**. Dynamic learning and memory, synaptic plasticity and neurogenesis: an update. *Frontiers in Behavioral Neuroscience*, 8, 106.
- Sul, J.H., Kim, H., Huh, N., Lee, D., Jung, M.W., **2010**. Distinct Roles of Rodent Orbitofrontal and Medial Prefrontal Cortex in Decision Making. *Neuron*, 66, 449–460.
- Sutherland, R.J., Chew, G.L., Baker, J.C., Linggard, R.C., **1987**. Some limitations on the use of distal cues in place navigation by rats. *Psychobiology*, 15, 48–57.

- Sutherland, R.J., Hamilton, D.A., **2004**. Rodent spatial navigation: at the crossroads of cognition and movement. *Neuroscience & Biobehavioral Reviews*, 28, 687–697.
- Sutherland, R.J., Rodriguez, A.J., **1989**. The role of the fornix/fimbria and some related subcortical structures in place learning and memory. *Behavioural Brain Research*, 32, 265–277.
- Sutton, J.E., **2009**. What is geometric information and how do animals use it? *Behavioural Processes*, 80, 339–343.
- Sutton, R.S., Barto, A.G., **1998**. Introduction to reinforcement learning. *MIT Press*.
- Swanson, L. W., Köhler, C., and Björklund, A., **1987**. The limbic region. I. The septohippocampal system. In: Björklund, A., Hökfelt, T., and Swanson, L. W., (Eds.) *Handbook of Chemical Neuroanatomy, Elsevier, Publisher, Amsterdam*, Vol. 5, Part I., pp. 125–227.
- T-**
- Tabuchi, E., Mulder, A.B., Wiener, S.I., **2003**. Reward value invariant place responses and reward site associated activity in hippocampal neurons of behaving rats. *Hippocampus*, 13, 117–132.
- Takeuchi, T., Duzskiewicz, A.J., Morris, R.G.M., **2014**. The synaptic plasticity and memory hypothesis: encoding, storage and persistence. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369, 20130288.
- Takeuchi, T., Morris, R.G.M., **2014**. Neuroscience: Shedding light on a change of mind. *Nature*, 513, 323–324.
- Tamosiunaite, M., Ainge, J., Kulvicius, T., Porr, B., Dudchenko, P., Wörgötter, F., **2008**. Path-finding in real and simulated rats: assessing the influence of path characteristics on navigation learning. *Journal of Computational Neuroscience*, 25, 562–582.
- Taube, J.S., **1995**. Head direction cells recorded in the anterior thalamic nuclei of freely moving rats. *The Journal of Neuroscience*, 15, 70–86.
- Taube, J.S., **1995**. Place cells recorded in the parasubiculum of freely moving rats. *Hippocampus*, 5, 569–583.
- Taube, J.S., **1998**. Head direction cells and the neurophysiological basis for a sense of direction. *Progress in Neurobiology*, 55, 225–256.
- Taube, J.S., **2007**. The Head Direction Signal: Origins and Sensory-Motor Integration. *Annual Review of Neuroscience*, 30, 181–207.
- Taube, J.S., Muller, R.U., Ranck, J.B., **1990**. Head-direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis. *The Journal of Neuroscience*, 10, 420–435.
- Templer, V.L., Hampton, R.R., **2013**. Episodic Memory in Nonhuman Animals. *Current Biology*, 23, R801–R806.
- Terada, S., Takahashi, S., Sakurai, Y., **2013**. Oscillatory interaction between amygdala and hippocampus coordinates behavioral modulation based on reward expectation. *Frontiers in Behavioral Neuroscience*, 7, 177.
- Thinus-Blanc, C., Bouzouba, L., Chaix, K., Chapuis, N., Durup, M., Poucet, B., **1987**. A study of spatial parameters encoded during exploration in hamsters. *Journal of Experimental Psychology: Animal Behavior Processes*, 13, 418–427.

- Thomas, G.J., Gash, D.M., **1988**. Differential effects of hippocampal ablations on dispositional and representational memory in the rat. *Behavioral Neuroscience*, 102, 635–642.
- Thome, C., Kelly, T., Yanez, A., Schultz, C., Engelhardt, M., Cambridge, S.B., Both, M., Draguhn, A., Beck, H., Egorov, A.V., **2014**. Axon-carrying dendrites convey privileged synaptic input in hippocampal neurons. *Neuron*, 83, 1418–1430.
- Thompson, L.T., Best, P.J., **1989**. Place cells and silent cells in the hippocampus of freely-behaving rats. *The Journal of Neuroscience*, 9, 2382–2390.
- Thompson, L.T., Best, P.J., **1990**. Long-term stability of the place-field activity of single units recorded from the dorsal hippocampus of freely behaving rats. *Brain Research*, 509, 299–308.
- Thorpe, S., Delorme, A., Van Rullen, R., **2001**. Spike-based strategies for rapid processing. *Neural Networks*, 14, 715–725.
- Thorpe, S., Fize, D., Marlot, C., **1996**. Speed of processing in the human visual system. *Nature*, 381, 520–522.
- Tobler, P.N., Fiorillo, C.D., Schultz, W., **2005**. Adaptive Coding of Reward Value by Dopamine Neurons. *Science*, 307, 1642–1645.
- Tolman, E.C., **1938**. The determiners of behavior at a choice point. *Psychological Review*, 45, 1–41.
- Tolman, E.C., **1939**. Prediction of vicarious trial and error by means of the schematic sowbug. *Psychological Review*, 46, 318–336.
- Tolman, E.C., **1948**. Cognitive maps in rats and men. *Psychological Review*, 55, 189–208.
- Tolman, E.C., Honzik, C.H., **1930**. “Insight” in rats. *University of California Publication in Psychology*, 4, 215–232.
- Tolman, E.C., Ritchie, B.F., Kalish, D., **1946**. Studies in spatial learning; place learning versus response learning. *Journal of Experimental Psychology*, 36, 221–229.
- Traissard, N., Herbeaux, K., Cosquer, B., Jeltsch, H., Ferry, B., Galani, R., Pernon, A., Majchrzak, M., Cassel, J.-C., **2007**. Combined damage to entorhinal cortex and cholinergic basal forebrain neurons, two early neurodegenerative features accompanying Alzheimer’s disease: effects on locomotor activity and memory functions in rats. *Neuropsychopharmacology*, 32, 851–871.
- Tulving, E., **1972**. Episodic and semantic memory, in: Organization of Memory. *Academic Press*, Oxford, England, pp. xiii, 423.
- Tulving, E., **1985**. Memory and consciousness. *Canadian Psychology/Psychologie canadienne*, 26, 1–12.
- Tulving, E., **1995**. Organization of memory: Quo vadis?, in: The Cognitive Neurosciences. *The MIT Press*, Cambridge, MA, US, pp. 839–853.
- Tulving, E., **2002**. Episodic Memory: From Mind to Brain. *Annual Review of Psychology*, 53, 1–25.
- Tulving, E., A, C., Macdonald, C.A., **1991**. Long-lasting perceptual priming and semantic learning in amnesia: A case experiment. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 17, 595–617.
- Tulving, E. (2005). Episodic memory and auto-noesis: Uniquely human? In H. S. Terrace & J. Metcalfe (Eds.), *The missing link in cognition: Origins of self-reflective consciousness* (pp. 3–56). New York: *Oxford University Press*.
- Tye, K.M., Cone, J.J., Schairer, W.W., Janak, P.H., **2010**. Amygdala Neural Encoding of the Absence of Reward during Extinction. *The Journal of Neuroscience*, 30, 116–125.

Tye, K.M., Stuber, G.D., de Ridder, B., Bonci, A., Janak, P.H., **2008**. Rapid strengthening of thalamo-amygdala synapses mediates cue–reward learning. *Nature*, 453, 1253–1257.

-U-

Uylings, H.B.M., Groenewegen, H.J., Kolb, B., **2003**. Do rats have a prefrontal cortex? *Behavioural Brain Research*, 146, 3–17.

Ulanovsky, N., Moss, C.F., **2007**. Hippocampal cellular and network activity in freely moving echolocating bats. *Nature Neuroscience*, 10, 224–233.

-V-

Vafaei, A.A., Rashidy-Pour, A., **2004**. Reversible lesion of the rat's orbitofrontal cortex interferes with hippocampus-dependent spatial memory. *Behavioural Brain Research*, 149, 61–68.

Van Cauter, T., Camon, J., Alvernhe, A., Elduayen, C., Sargolini, F., Save, E., **2013**. Distinct Roles of Medial and Lateral Entorhinal Cortex in Spatial Cognition. *Cerebral Cortex*, 23, 451–459.

Van Cauter, T., Poucet, B., Save, E., **2008**. Unstable CA1 place cell representation in rats with entorhinal cortex lesions. *European Journal of Neuroscience*, 27, 1933–1946.

Van der Jeugd, A., Goddyn, H., Laeremans, A., Arckens, L., D'Hooge, R., Verguts, T., **2009**. Hippocampal involvement in the acquisition of relational associations, but not in the expression of a transitive inference task in mice. *Behavioral Neuroscience*, 123, 109–114.

Van der Meer, M.A.A., Johnson, A., Schmitzer-Torbert, N.C., Redish, A.D., **2010**. Triple dissociation of information processing in dorsal striatum, ventral striatum, and hippocampus on a learned spatial decision task. *Neuron*, 67, 25–32.

Van der Meer, M.A.A., Redish, A.D., **2011**. Theta phase precession in rat ventral striatum links place and reward information. *The Journal of Neuroscience*, 31, 2843–2854.

Vanderwolf, C.H., Kramis, R., Gillespie, L.A., Bland, B.H., **1975**. Hippocampal Rhythmic Slow Activity and Neocortical Low-Voltage Fast Activity: Relations to Behavior, in: Isaacson, R.L., Pribram, K.H. (Eds.), *The Hippocampus*. Springer US, pp. 101–128.

Van Duuren, E., Escámez, F.A.N., Joosten, R.N.J.M.A., Visser, R., Mulder, A.B., Pennartz, C.M.A., **2007**. Neural coding of reward magnitude in the orbitofrontal cortex of the rat during a five-odor olfactory discrimination task. *Learning & Memory*, 14, 446–456.

Van Duuren, E., Lankelma, J., Pennartz, C.M.A., **2008**. Population coding of reward magnitude in the orbitofrontal cortex of the rat. *The Journal of Neuroscience*, 28, 8590–8603.

Van Duuren, E., van der Plasse, G., Lankelma, J., Joosten, R.N.J.M.A., Feenstra, M.G.P., Pennartz, C.M.A., **2009**. Single-cell and population coding of expected reward probability in the orbitofrontal cortex of the rat. *The Journal of Neuroscience*, 29, 8965–8976.

Vann, S.D., **2013**. Dismantling the Papez circuit for memory in rats. *Elife*, 2.

Van Strien, N.M., Cappaert, N.L.M., Witter, M.P., **2009**. The anatomy of memory: an interactive overview of the parahippocampal–hippocampal network. *Nature Reviews Neuroscience*, 10, 272–282.

- Vargha-Khadem, F., Gadian, D.G., Watkins, K.E., Connelly, A., Paesschen, W.V., Mishkin, M., **1997**. Differential Effects of Early Hippocampal Pathology on Episodic and Semantic Memory. *Science*, 277, 376–380.
- Verney, C., Baulac, M., Berger, B., Alvarez, C., Vigny, A., Helle, K.B., **1985**. Morphological evidence for a dopaminergic terminal field in the hippocampal formation of young and adult rat. *Neuroscience*, 14, 1039–1052.
- Verschure, P.F.M.J., Pennartz, C.M.A., Pezzulo, G., **2014**. The why, what, where, when and how of goal-directed choice: neuronal and computational principles. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369, 20130483.
- Vertes, R.P., **2004**. Differential projections of the infralimbic and prelimbic cortex in the rat. *Synapse*, 51, 32–58.
- Vertes, R.P., **2006**. Interactions among the medial prefrontal cortex, hippocampus and midline thalamus in emotional and cognitive processing in the rat. *Neuroscience*, 142, 1–20.
- Voorn, P., Vanderschuren, L.J.M.J., Groenewegen, H.J., Robbins, T.W., Pennartz, C.M.A., **2004**. Putting a spin on the dorsal–ventral divide of the striatum. *Trends in Neurosciences*, 27, 468–474.

-W-

- Wallace, D.G., Hamilton, D.A., Whishaw, I.Q., **2006**. Movement characteristics support a role for dead reckoning in organizing exploratory behavior. *Animal Cognition*, 9, 219–228.
- Wallace D.G, Whishaw I.Q., **2004**. Dead reckoning. In: Whishaw IQ, Kolb B, Eds. *The behavior of the laboratory rat: a handbook with tests*. New York: Oxford University Press.
- Wallace, D.G., Whishaw, I.Q., **2003**. NMDA lesions of Ammon’s horn and the dentate gyrus disrupt the direct and temporally paced homing displayed by rats exploring a novel environment: evidence for a role of the hippocampus in dead reckoning. *European Journal of Neuroscience*, 18, 513–523.
- Wallenstein, G.V., Hasselmo, M.E., Eichenbaum, H., **1998**. The hippocampus as an associator of discontinuous events. *Trends in Neurosciences*, 21, 317–323.
- Wallis, J.D., **2012**. Cross-species studies of orbitofrontal cortex and value-based decision-making. *Nature Neuroscience*, 15, 13–19.
- Wang, G.-W., Cai, J.-X., **2006**. Disconnection of the hippocampal–prefrontal cortical circuits impairs spatial working memory performance in rats. *Behavioural Brain Research*, 175, 329–336.
- Wang, X.-J., **2008**. Decision Making in Recurrent Neuronal Circuits. *Neuron*, 60, 215–234.
- Wehner, R., **2003**. Desert ant navigation: how miniature brains solve complex tasks. *Journal of Comparative Physiology A*, 189, 579–588.
- Weldon, D.A., Patterson, C.A., Colligan, E.A., Nemeth, C.L., Rizio, A.A., **2008**. Single unit activity in the rat superior colliculus during reward magnitude task performance. *Behavioral Neuroscience*, 122, 183–190.
- Whishaw, I.Q., Gharbawie, O.A., Clark, B.J., Lehmann, H., **2006**. The exploratory behavior of rats in an open environment optimizes security. *Behavioural Brain Research*, 171, 230–239.
- Whishaw, I.Q., Hines, D.J., Wallace, D.G., **2001**. Dead reckoning (path integration) requires the hippocampal formation: evidence from spontaneous exploration and spatial learning tasks in light (allothetic) and dark (idiothetic) tests. *Behavioural Brain Research*, 127, 49–69.

- White, N., **2007**. Multiple memory systems. *Scholarpedia*, 2, 2663.
- White, N.M., McDonald, R.J., **2002**. Multiple Parallel Memory Systems in the Brain of the Rat. *Neurobiology of Learning and Memory*, 77, 125–184.
- White, N.M., Packard, M.G., McDonald, R.J., **2013**. Dissociation of memory systems: The story unfolds. *Behavioral Neuroscience*, 127, 813–834.
- Whitlock, J.R., Heynen, A.J., Shuler, M.G., Bear, M.F., **2006**. Learning Induces Long-Term Potentiation in the Hippocampus. *Science* 313, 1093–1097.
- Wiener, S.I., **1993**. Spatial and behavioral correlates of striatal neurons in rats performing a self-initiated navigation task. *The Journal of Neuroscience*, 13, 3802–3817.
- Wiener, S.I., Paul, C.A., Eichenbaum, H., **1989**. Spatial and behavioral correlates of hippocampal neuronal activity. *The Journal of Neuroscience*, 9, 2737–2763.
- Wiener, S.I., Taube, J.S., **2005**. Head Direction Cells and the Neural Mechanisms of Spatial Orientation (Bradford Books). *The MIT Press*.
- Wikenheiser, A.M., Redish, A.D., **2011**. Changes in reward contingency modulate the trial-to-trial variability of hippocampal place cells. *Journal of Neurophysiology*, 106, 589–598.
- Wilent, W.B., Nitz, D.A., **2007**. Discrete Place Fields of Hippocampal Formation Interneurons. *Journal of Neurophysiology*, 97, 4152–4161.
- Wilkerson, A., Levin, E.D., **1999**. Ventral hippocampal dopamine D1 and D2 systems and spatial working memory in rats. *Neuroscience*, 89, 743–749.
- Wills, T.J., Lever, C., Cacucci, F., Burgess, N., O’Keefe, J., **2005**. Attractor Dynamics in the Hippocampal Representation of the Local Environment. *Science*, 308, 873–876.
- Wilson, A.G., Crystal, J.D., **2012**. Prospective memory in the rat. *Animal Cognition*, 15, 349–358.
- Wilson, M.A., McNaughton, B.L., **1993**. Dynamics of the hippocampal ensemble code for space. *Science*, 261, 1055–1058.
- Wilson, M.A., Tonegawa, S., **1997**. Synaptic plasticity, place cells and spatial memory: study with second generation knockouts. *Trends in Neurosciences*, 20, 102–106.
- Wilz, K.J., Bolton, R.L., **1971**. Exploratory behavior in response to the spatial rearrangement of familiar stimuli. *Psychonomic Science*, 24, 117–118.
- Winocur, G., Gilbert, M., **1984**. The hippocampus, context, and information processing. *Behavioral and Neural Biology*, 40, 27–43.
- Winocur, G., Moscovitch, M., Bontempi, B., **2010a**. Memory formation and long-term retention in humans and animals: Convergence towards a transformation account of hippocampal–neocortical interactions. *Neuropsychologia, Animal Models of Amnesia*, 48, 2339–2356.
- Winocur, G., Moscovitch, M., Rosenbaum, R.S., Sekeres, M., **2010b**. An investigation of the effects of hippocampal lesions in rats on pre- and postoperatively acquired spatial memory in a complex environment. *Hippocampus*, 20, 1350–1365.
- Wirsching, B.A., Beninger, R.J., Jhamandas, K., Boegman, R.J., El-Defrawy, S.R., **1984**. Differential effects of scopolamine on working and reference memory of rats in the radial maze. *Pharmacology Biochemistry and Behavior*, 20, 659–662.
- Wise, R.A., **1987**. Sensorimotor modulation and the variable action pattern (VAP): Toward a noncircular definition of drive and motivation. *Psychobiology*, 15, 7–20.
- Wise, S.P., **2008**. Forward frontal fields: phylogeny and fundamental function. *Trends in Neurosciences*, 31, 599–608.

- Witter M.P., Amaral D.G., **2004**. Hippocampal formation. In: Paxinos G. (Ed.), *The rat nervous system*, Elsevier, Amsterdam, pp 635–704.
- Witter, M.P., Canto, C.B., Couey, J.J., Koganezawa, N., O'Reilly, K.C., **2014**. Architecture of spatial circuits in the hippocampal region. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369(1635), 20120515.
- Witter, M.P., Room, P., Groenewegen, H.J., Lohman, A.H., **1988**. Reciprocal connections of the insular and piriform claustrum with limbic cortex: an anatomical study in the cat. *Neuroscience*, 24, 519–539.
- Wixted, J.T., Squire, L.R., Jang, Y., Papesh, M.H., Goldinger, S.D., Kuhn, J.R., Smith, K.A., Treiman, D.M., Steinmetz, P.N., **2014**. Sparse and distributed coding of episodic memory in neurons of the human hippocampus. *Proceedings of the National Academy of Sciences*, 111, 9621–9626.
- Wood, E.R., Dudchenko, P.A., Robitsek, R.J., Eichenbaum, H., **2000**. Hippocampal Neurons Encode Information about Different Types of Memory Episodes Occurring in the Same Location. *Neuron*, 27, 623–633.

-X-

- Xu, W., Südhof, T.C., **2013**. A Neural Circuit for Memory Specificity and Generalization. *Science*, 339, 1290–1295.

-Y-

- Yamada, H., Inokawa, H., Matsumoto, N., Ueda, Y., Enomoto, K., Kimura, M., **2013**. Coding of the long-term value of multiple future rewards in the primate striatum. *Journal of Neurophysiology*, 109, 1140–1151.

-Z-

- Zaehle, T., Rach, S., Herrmann, C.S., **2010**. Transcranial Alternating Current Stimulation Enhances Individual Alpha Activity in Human EEG. *PLoS ONE*, 5.
- Zeithamova, D., Schlichting, M.L., Preston, A.R., **2012**. The hippocampus and inferential reasoning: building memories to navigate future decisions. *Frontiers in Human Neuroscience*, 6, 70.
- Zhang, S.-J., Ye, J., Miao, C., Tsao, A., Cerniauskas, I., Ledergerber, D., Moser, M.-B., Moser, E.I., **2013**. Optogenetic Dissection of Entorhinal-Hippocampal Functional Connectivity. *Science*, 340.
- Zhang, S., Schoenfeld, F., Wiskott, L., Manahan-Vaughan, D., **2014**. Spatial representations of place cells in darkness are supported by path integration and border information. *Frontiers in Behavioral Neuroscience*, 8, 222.
- Zhou, W., Hohmann, A.G., Crystal, J.D., **2012**. Rats Answer an Unexpected Question after Incidental Encoding. *Current Biology*, 22, 1149–1153.
- Zinyuk, L., Kubik, S., Kaminsky, Y., Fenton, A.A., Bures, J., **2000**. Understanding hippocampal activity by using purposeful behavior: Place navigation induces place cell discharge in both task-relevant

- and task-irrelevant spatial reference frames. *Proceedings of the National Academy of Sciences*, 97, 3771–3776.
- Zola-Morgan, S., Squire, L.R., Amaral, D.G., **1986**. Human amnesia and the medial temporal region: enduring memory impairment following a bilateral lesion limited to field CA1 of the hippocampus. *The Journal of Neuroscience*, 6, 2950–2967.
- Zucker, R.S., Regehr, W.G., **2002**. Short-Term Synaptic Plasticity. *Annual Review of Physiology*, 64, 355–405.
- Zugaro, M.B., Tabuchi, E., Fouquier, C., Berthoz, A., Wiener, S.I., **2001**. Active Locomotion Increases Peak Firing Rates of Anterodorsal Thalamic Head Direction Cells. *Journal of Neurophysiology*, 86, 692–702.
- Zugaro, M.B., Arleo, A., Berthoz, A., Wiener, S.I., **2003**. Rapid Spatial Reorientation and Head Direction Cells. *The Journal of Neuroscience*, 23, 3478–3482.