

38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65
66
67
68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87

SUMMARY

There are rich structures in off-task neural activity which are hypothesised to reflect fundamental computations across a broad spectrum of cognitive functions. Here, we develop an analysis toolkit – Temporal Delayed Linear Modelling (TDLM) for analysing such activity. TDLM is a domain-general method for finding neural sequences that respect a pre-specified transition graph. It combines nonlinear classification and linear temporal modelling to test for statistical regularities in sequences of task-related reactivations. TDLM is developed on the non-invasive neuroimaging data and is designed to take care of confounds and maximize sequence detection ability. Notably, as a linear framework, TDLM can be easily extended, without loss of generality, to capture rodent replay in electrophysiology, including in continuous spaces, as well as addressing second-order inference questions, e.g., its temporal and spatial varying pattern. We hope TDLM will advance a deeper understanding of neural computation and promote a richer convergence between animal and human neuroscience.

INTRODUCTION

Human neuroscience has made remarkable progress in detailing the relationship between the representations of different stimuli during task performance¹⁻³. At the same time, it is increasingly clear that resting, off-task, brain activities are structurally rich^{4,5}. An ability to study spontaneous activity with respect to task-related representation is important for understanding cognitive process beyond current sensation⁶. However, unlike the case for task-based activity, little attention has been given to techniques that can measure representational content of resting brain activity in humans.

Unlike human neuroscience, representational content of resting activity is studied extensively in animal neuroscience. One seminal example is “hippocampal replay”⁷⁻¹⁰: During sleep, and quiet wakefulness, place cells in the hippocampus (that signal self-location during periods of activity) spontaneously recapitulate old, and explore new, trajectories through an environment. These internally generated sequences are hypothesized to reflect a fundamental feature of neural computation across tasks¹¹⁻¹⁵. Numerous methods have been proposed to analyse hippocampal replay¹⁶⁻¹⁸. However, they are not domain general in that they are designed to be most suited for specific needs, such as particular task design, data modality, or research question^{19,20}. Most commonly, these methods apply to invasive electrophysiology signals, aiming to detect sequences in a linear track during spatial navigation task²⁰. As a result, they cannot be directly adapted for analysing human resting activity collected using non-invasive neuroimaging techniques. Furthermore, in rodent neuroscience, it is non-trivial to adapt these algorithms to even small changes in tasks (such as 2D foraging). This may be a limiting factor in taking replay analyses to more interesting and complex tasks, such as complex mazes²¹.

Here, we introduce TDLM (temporal delayed linear modelling), a domain general analysis toolkit, for characterizing temporal structure of internally generated neural representations in rodent electrophysiology as well as human neuroimaging data. TDLM is inspired by existing replay detection methods^{8,16,17}, especially those analysis of population of relay events¹⁷. It is developed based on the General Linear Modelling (GLM) framework, and can therefore easily accommodate testing of "second-order" statistical questions¹⁹, such as whether there is

88 more forward than reverse replay, or is replay strength changing over time, or differs between
89 behavioural conditions. This type of question is ubiquitous in cognitive studies, but is
90 typically addressed ad-hoc in other replay detection methods¹⁹. In TDLM, such questions are
91 treated naturally as linear contrasts of effects in a GLM.

92

93 Here we show TDLM is suited to measure the **average** amount of replay across many events
94 (i.e., replay strength) in linear modelling. This makes it applicable to both rodent
95 electrophysiology and human neuroimaging. Applying TDLM on non-invasive neuroimaging
96 data in humans, we, and others, have shown it is possible to measure the average
97 sequenceness (propensity for replay) in spontaneous neural representations²²⁻²⁵. The results
98 resemble key characteristics found in rodent hippocampal replay and inform key
99 computational principles of human cognition²⁴.

100

101 In the following sections, we first introduce the logic and mechanics of TDLM in detail,
102 followed by a careful treatment of its statistical inference procedure. We test TDLM in both
103 simulation (see the Methods section “Simulating MEG data”) and real human MEG/EEG
104 data (see the Methods section “Human replay dataset”). We then turn to rodent
105 electrophysiology and compare TDLM to existing rodent replay methods, extending TDLM
106 to work on a continuous state space. Lastly, using our approach we re-analyse rodent
107 electrophysiology data from Ólafsdóttir, et al.²⁶ (see the Methods section “Rodent replay
108 dataset”), and show what TDLM can offer uniquely compared to existing methods in rodent
109 replay analysis.

110

111 To summarise, TDLM is a general, and flexible, tool for measuring neural sequences. It
112 facilitates cross-species investigations by linking large-scale measurements in humans to
113 single neuron measurements in non-human species. It provides a powerful tool for revealing
114 abstract cognitive processes that extend beyond sensory representation, potentially open
115 doors for new avenues of research in cognitive science.

116

117

118

119 **RESULTS**

120

121 **TDLM**

122

123 **Overview of TDLM**

124

125 Our primary goal is to test for temporal structure of neural representations in humans.
126 However, to facilitate cross-species investigation²⁷, we also want to extend this method to
127 enable measurement of sequences in other species (e.g., rodents). Consequently, this
128 sequence detection method has to be domain general. We choose to measure sequences in a
129 decoded state space (e.g., posterior estimated locations in rodents¹⁷ or time course of task-
130 related reactivations in humans²⁸) as this makes results from different data types comparable.

131

132 Ideally, a general sequence detection method should (1) uncover structural regularities in the
133 reactivation of neural activity, (2) control for confounds that are not of interest, and (3) test
134 whether this regularity conforms to a hypothesized structure. To achieve these goals, we
135 developed the method under a GLM framework, and henceforth refer to it as Temporal
136 Delayed Linear Modelling, i.e., TDLM. Although TDLM works on a decoded state space, it

137 still needs to take account of confounds inherent in the data where the state space is decoded
138 from. This is a main focus of TDLM.

139

140 The starting point of TDLM is a set of n time series, each corresponding to a decoded neural
141 representation of a task variable of interest. This is what we call the state space, X , with
142 dimension of time by states. These time series could themselves be obtained in several ways,
143 described in detail in a later section (“Getting the states”). The aim of TDLM is to identify
144 task-related regularities in sequences of these representations.

145

146 Consider, for example, a task in which participants have been trained such that $n=4$ distinct
147 sensory objects (A, B, C, and D) appear in a consistent order: $A \rightarrow B \rightarrow C \rightarrow D$ (Figure 1a, b).
148 If we are interested in replay of this sequence during subsequent resting periods (Figure 1c,
149 d), we might want to ask statistical questions of the following form: “Does the existence of a
150 neural representation of A, at time T , predict the occurrence of a representation of B at time
151 $T+\Delta t$ ”, and similarly for $B \rightarrow C$ and $C \rightarrow D$.

152

153 In TDLM we ask such questions using a two-step process. First, for each of the n^2 possible
154 pairs of variables X_i and X_j , we find the linear relation between the X_i time series and the Δt -
155 shifted X_j time series. These n^2 relations comprise an empirical transition matrix, describing
156 how likely each variable is to be succeeded at a lag of Δt by each other variable (Figure 1e).
157 Second, we linearly relate this empirical transition matrix with a task-related transition matrix
158 of interest (Figure 1f). This produces a single number that characterizes the extent to which
159 the neural data follow the transition matrix of interest, which we call ‘sequenceness’. Finally,
160 we repeat this entire process for all Δt of interest, yielding a measure of sequenceness at each
161 possible lag between variables, and submit this for statistical inference (Figure 1g).

162

163 Note that, for now, this approach decomposes a sequence (such as $A \rightarrow B \rightarrow C \rightarrow D$) into its
164 constituent transitions and sums the evidence for each transition. Therefore, it does not
165 require that the transitions themselves are sequential: $A \rightarrow B$ and $B \rightarrow C$ could occur at
166 unrelated times, so long as the within-pair time lag was the same. For interested readers, we
167 address how to strengthen the inference by looking explicitly for longer sequences in the
168 Appendix 1: Multi-step sequences.

169

170

171 **Constructing the empirical transition matrix**

172

173 In order to find evidence for state-to-state transitions at some time lag Δt , we could regress a
174 time-lagged copy of one state, X_j , onto another, X_i (omitting residual term ε in all linear
175 equations):

176

$$177 X_j(t + \Delta t) = X_i(t)\beta_{ij}$$

178 (1)

179

180 Instead, TDLM chooses to include all states in the same regression model for important
181 reasons, detailed in section “Moving to multiple linear regression”:

182

$$183 X_j(t + \Delta t) = \sum_{k=1}^n X_k(t)\beta_{kj} \quad (2)$$

184

185 In this equation, the values of all states X_k at time t are used in a single multilinear model to
 186 predict the value of the single state X_j at time $t + \Delta t$.

187
 188 The regression described in Equation 2 is performed once for each X_j , and these equations
 189 can be arranged in matrix form as follows:

$$190$$

$$191 \quad X(\Delta t) = X\beta$$

192 (3)

193
 194 Each row of X is a timepoint, and each of the n columns is a state. $X(\Delta t)$ is the same matrix
 195 as X , but with the rows shifted forwards in time by Δt . β_{ij} is an estimate of the influence of
 196 $X_i(t)$ on $X_j(t + \Delta t)$. β is an $n \times n$ matrix of weights, which we call the *empirical transition*
 197 *matrix*.

198
 199 To obtain β , we invert Equation 3 by ordinary least squares regression.

$$200$$

$$201 \quad \beta = (X^T X)^{-1} X^T X(\Delta t)$$

202 (4)

203 This inversion can be repeated for each possible time lag ($\Delta t = 1, 2, 3, \dots$), resulting in a
 204 separate empirical transition matrix β at every time lag. We call this step the first level
 205 sequence analysis.

206 207 208 **Testing the hypothesized transitions**

209
 210 The first level sequence analysis assesses evidence for all possible state-to-state transitions.
 211 The next step in TDLM is to test for the strength of a particular hypothesized sequence,
 212 specified as a transition matrix, T_F . Therefore, we construct another GLM which relates T_F to
 213 the empirical transition matrix, β . We call this step the second level sequence analysis:

$$214 \quad \beta = \sum_{r=1}^r Z(r) * T_r$$

215 (5)

216 As noted above, β is the empirical transition matrix obtained from first-stage GLM. It has
 217 dimension of n by n , where n is the number of states. Each entry in β reflects the unique
 218 contribution of state i to state j at given time lag. Effectively, the above equation models this
 219 empirical transition matrix β as a weighted sum of *prespecified template matrices*, T_r .
 220 Thus, r is the number of regressors included in the second-stage GLM, and each scalar
 221 valued $Z(r)$ is the weight assigned to the r^{th} template matrix. Put in other words, T_r
 222 constitutes the regressors in the design matrix, each of which has a prespecified template
 223 structure, e.g., T_{auto} , T_{const} , T_F and T_B (Figure 1h).

224
 225 T_F and T_B are the transpose of each other (e.g., red and blue entries in Figure 1b), indicating
 226 transitions of interest in forward and backward direction, respectively. In 1D physical space
 227 T_F and T_B would be the shifted diagonal matrices with ones on the first upper and lower off
 228 diagonals. T_{const} is a constant matrix that models away the average of all transitions,
 229 ensuring that any weight on T_F and T_B reflects its unique contribution. T_{auto} is the identity
 230 matrix, T_{auto} models self-transitions to control for auto-correlation (equivalently, we could
 231 simply omit the diagonal elements from the regression).

232

233 Z is the weights of the second level regression, which is a vector with dimension of r by 1.
234 Each entry in Z reflects the strength of the hypothesised transitions in the empirical ones, i.e.,
235 sequenceness. Repeating the regression of Equation 5 at each time lag ($\Delta t = 1, 2, 3, \dots$)
236 results in time courses of the sequenceness as a function of time lag (e.g., the solid black line
237 in Figure 1f). Z_F, Z_B are the forward and backward sequenceness respectively (e.g., red and
238 blue lines in Figure 1g).

239

240 In many cases, Z_F and Z_B will be the final outputs of a TDLM analysis. However, it may
241 sometimes also be useful to consider the quantity:

242

$$D = Z_F - Z_B \quad (6)$$

244

245 D contrasts forward and backward sequences to give a measure that is positive if sequences
246 occur mainly in a forward direction and negative if sequences occur mainly in a backward
247 direction. This may be advantageous if, for example, Z_F and Z_B are correlated across subjects
248 (due to factors such as subject engagement and measurement sensitivity). In this case, D may
249 have lower cross-subject variance than either Z_F or Z_B , as the subtraction removes common
250 variance.

251

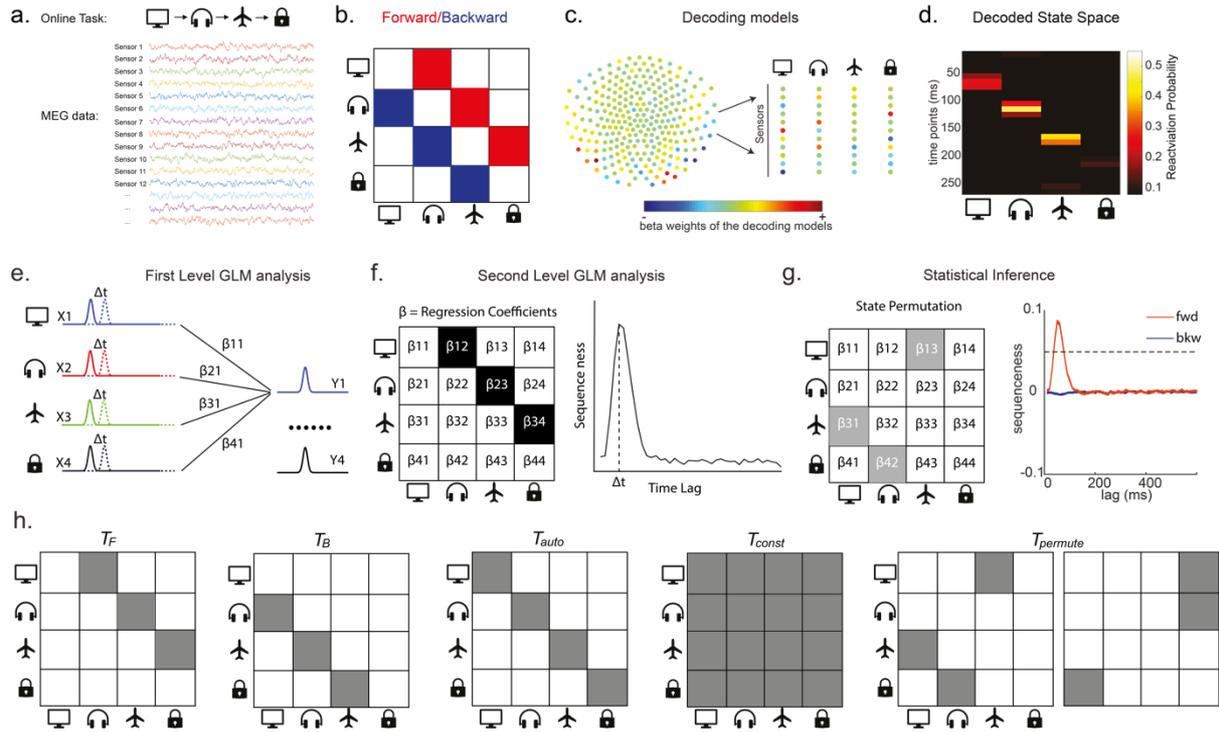
252 Finally, to test for statistical significance, TDLM relies on a nonparametric permutation-
253 based method. The null distribution is constructed by randomly shuffling the identities of the
254 n states many times and re-calculating the second level analysis for each shuffle (Figure 1g).
255 This approach allows us to reject the null hypothesis that there is no relationship between the
256 empirical transition matrix and the task-defined transition of interest. Note that there are
257 many incorrect ways to perform permutations, which permute factors that are not
258 exchangeable under the null hypothesis and therefore lead to false positives. We examine
259 some of these later with simulations and real data. In some cases, it may be desirable to test
260 slightly different hypotheses by using a different set of permutations; this is discussed later.

261

262 If the time lag Δt at which neural sequences exist is not known *a priori*, then we must correct
263 for multiple comparisons over all tested lags. This can be achieved by using the maximum Z_F
264 across all tested lags as the test statistic (see details in section “Correcting for multiple
265 comparisons”). If we choose this test statistic, then any values of Z_F exceeding the 95th
266 percentile of the null distribution can be treated as significant at $\alpha = 0.05$ (e.g., the grey
267 dotted line in Figure 1g).

268

269



270
271

272 **Figure 1. Task design and illustration of TDLM.** **a**, Task design in both simulation and real MEG
 273 data. Assuming there is one sequence, A->B->C->D, indicated by the four objects at the top. During
 274 the task, participants are shown the objects, and asked to figure out a correct sequence for these
 275 objects while undergoing MEG scanning. A snapshot of MEG data is shown below. It is a matrix with
 276 dimensions of sensors by time. **b**, The transitions of interest are shown, with the red and blue entries
 277 indicating transitions in the forward and backward direction respectively. **c**, The first step of TDLM is
 278 to construct decoding models of states from task data, and **(d)** then transform the data (e.g., resting-
 279 state) from sensor space to the state space. TDLM works on the decoded state space throughout. **e**,
 280 The second step of TDLM is to quantify the temporal structure of the decoded states using multiple
 281 linear regressions. The first level GLM results in a state*state regression coefficient matrix (empirical
 282 transition matrix), β at each time lag. **f**, In the second-level GLM, this coefficient matrix is projected
 283 onto the hypothesized transition matrix (black entries), to give a single measure of sequenceness.
 284 Repeating this process for the number of time lags of interest generates sequenceness over time lags
 285 (right panel). **g**, The statistical significance of sequenceness is tested using a nonparametric state
 286 permutation test by randomly shuffling the transition matrix of interest (in grey). To control for
 287 multiple comparisons, the permutation threshold is defined as the 95th percentile of all shuffles on the
 288 maximum value over all tested time lags. **h**, The second level regressors T_{auto} , T_{const} , T_F and T_B , as
 289 well as two examples of the permuted transitions of interest, $T_{permute}$ (for constructing permutation
 290 test), are shown.

291
292
293

294 TDLM STEPS IN DETAIL

295
296
297

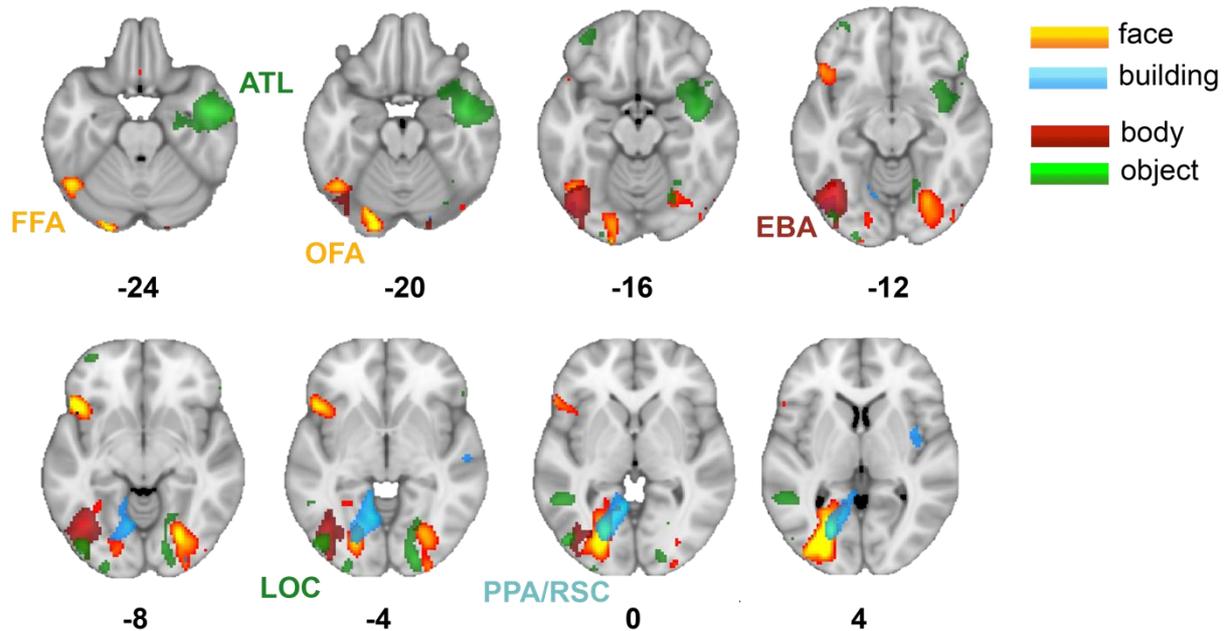
296 Getting the states

298 As described above, the input to TDLM is a set of time series of decoded neural
 299 representations, or states. Here we provide different examples of specific state spaces (X ,
 300 with dimension of time by states) that we have worked with using TDLM.

301
302

302 States as sensory stimuli

303 The simplest case, perhaps, is to define a state in terms of a neural representation of sensory
 304 stimuli, e.g., face, house. To obtain their associated neural representation, we present these
 305 stimuli in a randomized order at the start of a task, and record whole-brain neural activity
 306 using a non-invasive neuroimaging method, e.g., MEG or EEG. We then train a model to
 307 map the pattern of recorded neural activity to the presented image (Figure 1-figure
 308 supplement 1). This could be any of the multitude of available decoding models. For
 309 simplicity we used a logistic regression model throughout.
 310



311
 312
 313 **Figure 1-figure supplement 1. Source localization of stimuli evoked neural activity in MEG.** The
 314 states here are defined in terms of stimuli evoked neural activity. The classifiers are trained at 200 ms
 315 post-stimulus onset. For example, the stimuli are faces, buildings, body parts, and objects. Source
 316 localizing the evoked neural activity, we found that the activation patterns of stimuli in MEG signal
 317 are consistent with those reported in fMRI literature. For faces, activation peaked in a region roughly
 318 consistent with the fusiform face area (FFA) as well as the occipital face area (OFA). Activation for
 319 building stimuli was located between a parahippocampal place area (PPA) and retrosplenial cortex
 320 (RSC), a region also known to respond to scene and building stimuli. Activation for body part stimuli
 321 localised to a region consistent with the extrastriate body area (EBA). Activation for objects was in a
 322 region consistent with an object-associated lateral occipital cortex (LOC) as well as an anterior
 323 temporal lobe (ATL) cluster that may relate to conceptual processing of objects. Those maps are
 324 thresholded to display localized peaks. The full un-thresholded maps can be found at
 325 <https://neurovault.org/collections/6088/>. This is adapted from Wimmer, et al.²².
 326
 327
 328

329 In MEG/EEG, neural activity is recorded by multiple sensor arrays on the scalp. The sensor
 330 arrays record whole-brain neural activity at millisecond temporal resolution. To avoid a
 331 potential selection bias (given the sequence is expressed in time), we choose whole brain
 332 sensor activity at a single time point (i.e., spatial feature) as the training data fed into
 333 classifier training.
 334

335 Ideally, we would like to select a time point where the neural activity can be most truthfully
 336 read out. This can be indexed as the time point that gives the peak decoding accuracy. If the

337 state is defined by the sensory features of stimuli, we can use a classical leave-one-out cross-
338 validation scheme to determine the ability of classifiers to generalise to unseen data of the
339 same stimulus type (decoding accuracy) at each time point (see Appendix 2 for its algorithm
340 box). In essence, this cross-validation scheme is asking whether the classifier trained on this
341 sensory feature can be used to classify the unseen data of the same stimuli (Figure 2a, b).

342

343 After we have identified the peak time point based on the cross validation, we can train the
344 decoding models based on the multivariate sensor data at this given time.

345

346 Specifically, let's denote the training data, M , with dimension of number of observations, b ,
347 by number of sensors, s . The labels, Y , have dimension of b by 1. The aim here is to obtain
348 the classifier weights, W , so that $Y \approx \sigma(MW)$. σ is the logistic sigmoid function.

349

350 Normally we apply L1 regularization on the inference of weights (we will detail the reasons
351 in the later section "Regularization"):

352

$$353 \quad W = \operatorname{argmax}_W [\log(P(Y|M, W)) + b \lambda_{L1} \|W\|_1]$$

354 (7)

355

356 Next, we translate the data at testing time (e.g., during rest), R , from sensor space to the
357 decoded state space:

358

$$359 \quad X = \sigma(RW) \quad (8)$$

360

361 R is the testing data, with dimension of time by sensors. X is the decoded state space, with
362 dimension of time by states.

363

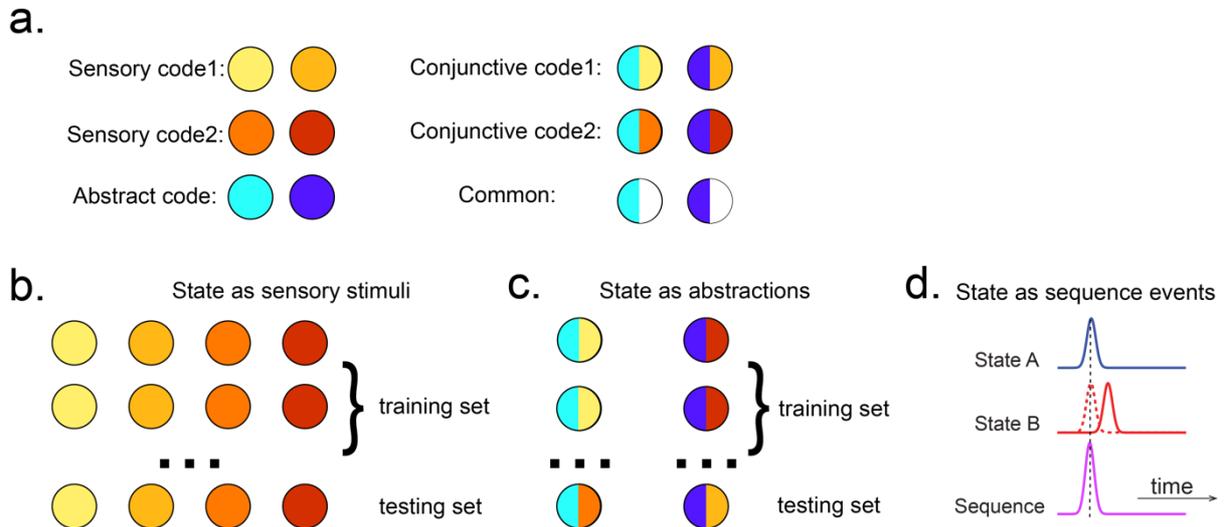
364

365 **States as abstractions**

366 As well as sequences of sensory representations, it is possible to search for replay of more
367 abstract neural representations. Such abstractions might be associated with the presented
368 image (e.g., mammal vs fish), in which case analysis can proceed as above by swapping
369 categories for images²². A more subtle example, however, is where the abstraction pertains
370 to the sequence or graph itself. In space, for example, grid cells encode spatial coordinates in
371 a fashion that abstracts over the sensory particularities of any one environment, and therefore
372 can be reused across environments²⁹. In human studies similar representations have been
373 observed for the location in a sequence^{24,30}. For example, different sequences have shared
374 representations for their second items (Figure 2). These representations also replay
375²⁴. However, to measure this replay we need to train decoders for these abstract
376 representations. This poses a conundrum as it is not possible to elicit the abstract
377 representations in the absence of the concrete sequence (the sensory stimuli). Care is required
378 to ensure that the decoders are sensitive to the abstract code rather than the sensory
379 representations (see Appendix 2 for algorithm box of selecting time point for training abstract
380 code). Useful strategies include training classifiers to generalise across stimulus sets, and
381 ensuring the classifiers are orthogonal to sensory representations (Figure 2-figure supplement
382 1 - details in Liu, et al.²⁴). One way that excludes the possibility of sensory contamination is
383 if the structural representations can be shown to sequence before the subjects have ever seen
384 their sensory correlates²⁴.

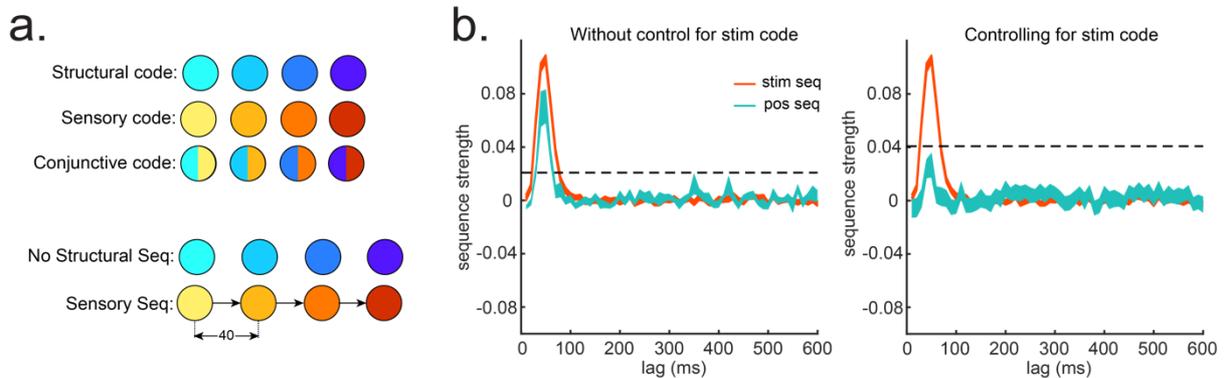
385

386 TDLM can also be used iteratively to ask questions about the ordering of different types of
 387 replay events (Figure 2d). This can provide for powerful inferences about the temporal
 388 organisation of replay, such as the temporal structure between sequences, or the repeating
 389 pattern of the same sequence. This more sophisticated use of TDLM merits its own
 390 consideration and is discussed in the Appendix 3: Sequences of sequences.
 391
 392



393
 394
 395 **Figure 2. Obtaining different state spaces.** **a**, Assuming we have two abstract codes, each abstract
 396 code has two different sensory codes (left panel). The M/EEG data corresponding to each stimulus is
 397 a conjunctive representation of sensory and abstract codes (right panel). The abstract code can be
 398 operationalised as the common information in the conjunctive codes of two stimuli. **b**, Training
 399 decoding models for stimulus information. The simplest state is defined by sensory stimuli. To
 400 determine the best time point for classifier training, we can use a classical leave-one-out cross
 401 validation scheme on the stimuli-evoked neural activity. **c**, Training decoding models for abstracted
 402 information. The state can also be defined as the abstractions. To extract this information, we need to
 403 avoid a confound of sensory information. We can train the classifier on the neural activity evoked by
 404 one stimulus and test it on the other sharing the same abstract representation. If neural activity
 405 contains both a sensory and abstract code, then the only information that can generalize is the
 406 common abstract code. **d**, The state can also be defined as the sequence event itself.

407
 408
 409



410
 411
 412 **Figure 2-figure supplement 1. Sequences of abstract code.** **a**, Illustration of the relationship
 413 between sensory code and (abstract) structural code. Structural code cannot be accessed directly but

414 can be indirectly obtained from the conjunctive code (overlapping representation of sensory and
415 structural code). In this simulation, there is sequence of sensory code but not of structural code. **b**, We
416 show the importance of controlling for sensory (stim) information when looking for sequences of
417 abstract code: If sensory information is not controlled, we would observe significant sequences of
418 structural code, while in fact, it is not present, i.e., false positive.

419

420

421 **Controlling confounds and maximising sensitivity in sequence detection**

422

423 Here, we motivate the key features of TDLM.

424

425 **Temporal correlations**

426

427 In standard linear methods, unmodelled temporal autocorrelation can inflate statistical scores.
428 Techniques such as auto-regressive noise modelling are commonplace to mitigate these
429 effects^{31,32}. However, autocorrelation is a particular burden for analysis of sequences, where
430 it interacts with correlations between the decoded neural variables.

431

432 To see this, consider a situation where we are testing for the sequence $X_i \rightarrow X_j$. TDLM is
433 interested in the correlation between X_i and lagged X_j (see Equation 1). But if the X_i and X_j
434 time series contain autocorrelations, and are also correlated with one another, then $X_i(t)$ will
435 necessarily be correlated with $X_j(t + \Delta t)$. Hence, the analysis will spuriously report
436 sequences.

437

438 Correlations between states are commonplace. Consider representations of visual stimuli
439 decoded from neuroimaging data. If these states are decoded using an n -way classifier
440 (forcing exactly one state to be decoded at each moment), then the n states will be anti-
441 correlated by construction. On the other hand, if states are each classified against a null state
442 corresponding to the absence of stimuli, then the n states will typically be positively
443 correlated with one another.

444

445 Notably, in our case, because these autocorrelations are identical between forward and
446 backward sequences, one approach for removing them is to compute the difference measure
447 described above ($D = Z_F - Z_B$). This works well as shown in Kurth-Nelson, et al.¹¹.
448 However, a downside is it prevents us from measuring forward and backward sequences
449 independently. The remainder of this section considers alternative approaches that allow for
450 independent measurement of forward and backward sequences.

451

452 **Moving to multiple linear regression:** The spurious correlations above are induced because
453 $X_j(t)$ mediates a linear relationship between $X_i(t)$ and $X_j(t + \Delta t)$. Hence, if we knew $X_j(t)$,
454 we can solve the problem by simply controlling for it in a linear regression, as in Granger
455 Causality³³:

456

$$457 X_j(t + \Delta t) = \beta_0 + X_i(t)\beta_{ij} + X_j(t)\beta_{jj} \quad (9)$$

458

459 Unfortunately, we do not have access to the ground truth of X because these variables have
460 been decoded noisily from brain activity. Any error in $X_j(t)$ but not $X_i(t)$ will mean that the
461 control for autocorrelation is imperfect, leading to spurious weight on β_{ij} , and therefore
462 spurious inference of sequences.

463
464 This problem cannot be solved without a perfect estimate of X , but it can be systematically
465 reduced until negligible. It turns out the necessary strategy is simple. We do not know ground
466 truth $X_j(t)$, but what if we knew a subspace that included estimated $X_j(t)$? If we control for
467 that whole subspace, we would be on safe ground. We can get closer and closer to this by
468 including further co-regressors that are themselves correlated with estimated $X_j(t)$ with
469 different errors from ground truth $X_j(t)$. The most straightforward approach is to include the
470 other states of $X(t)$, each of which has different errors, leading to the multiple linear
471 regression of Equation 2.

472
473 Figure 3a shows this method applied to the same simulated data whose correlation structure
474 induces false positives in the simple linear regression of Equation 1, and by the same logic,
475 so too in cross correlation. This is why previous studies based on a cross-correlation^{34,35}
476 cannot look for sequenceness in forward and backward directions separately, but have to rely
477 on their asymmetry. The multiple regression accounts for the correlation structure of the data
478 and allows correct inference to be made. Unlike the simple subtraction method proposed
479 above (Figure 3a, left panel), the multiple regression permits separate inference on forwards
480 and backwards sequences.

481
482 **Oscillations and long timescale autocorrelations:** Equation 2 performs multiple regression,
483 regressing each $X_j(t + \Delta t)$ onto each $X_i(t)$ whilst controlling for all other state estimates at
484 time t . This method works well when spurious relationships between $X_i(t)$ and $X_j(t + \Delta t)$
485 are mediated by the subspace spanned by the other estimated states at time t (in particular
486 $X_j(t)$). One situation in which this assumption might be challenged is when replay is
487 superimposed on a large neural oscillation. For example, during rest (which is often the time
488 of interest in replay analysis), MEG and EEG data often express a large alpha rhythm, at
489 around 10Hz.

490
491 If all states experience the same oscillation at the same phase, the approach correctly controls
492 false positives. The oscillation induces a spurious correlation between $X_i(t)$ and $X_j(t + \Delta t)$
493 but, as before, this spurious correlation is mediated by $X_j(t)$.

494
495 However, this logic fails when states experience oscillations at different phases. This scenario
496 may occur, for example, if we assume there are travelling waves in cortex^{36,37}, because
497 different sensors will experience the wave at different times, and different states have
498 different contributions from each sensor. MEG sensors can be seen as measures of local field
499 potential on the scalp, which contain background neural oscillations. In humans this is
500 dominantly alpha during rest.

501
502 In this case, $X_i(t)$ predicts $X_j(t + \Delta t)$ over and above $X_j(t)$. To see this, consider the
503 situation where Δt is $\frac{1}{4} \tau$ (where τ is the oscillatory period) and the phase shift between $X_i(t)$
504 and $X_j(t)$ is $\pi/2$. Now every peak in $X_j(t + \Delta t)$ corresponds to a peak in $X_i(t)$ but a zero of
505 $X_j(t)$.

506
507 To combat this, we can include phase shifted versions/more timepoints of $X(t)$. If dominant
508 background oscillation is at alpha frequency (e.g., 10Hz), neural activity at time T would be
509 correlated with activity at time $T + \tau$. We can control for that, by including $X(t + \tau)$, as well
510 as $X(t)$ in the GLM (Figure 3b). Here $\tau = 100$ ms, if assuming the frequency is 10Hz.

511 Applying this method to the real MEG data during rest, we see much diminished 10Hz
512 oscillation in sequence detection during rest ²⁴.

513

514 **Spatial correlations**

515

516 As mentioned above, correlations between decoded variables commonly occur. The simplest
517 type of decoding model is a binary classifier that maps brain activity to one of two states.
518 These states will, by definition, be perfectly anti-correlated. Conversely, if separate classifiers
519 are trained to distinguish each state’s representation from baseline (“null”) brain data, then
520 the states will often be positively correlated with each other.

521

522 Unfortunately, positive or negative correlations between states reduces the sensitivity of
523 sequence detection, because it is difficult to distinguish between states within the sequence:
524 collinearity impairs estimation of β in Equation 2. In Figure 3c, we show in simulation that
525 the ability to detect real sequences goes down as the absolute value of a spatial correlation
526 goes up. We took the absolute value here because the direction of correlation is not important,
527 only the magnitude of the correlation matters.

528

529 Ideally, the state decoding models should be as independent as possible. We have suggested
530 the approach of training models to discriminate one state against a mixture of other states and
531 null data ^{24,35}. This mixture ratio can be adjusted. Adding more null data causes the states to
532 be positively correlated with each other, while less null data leads to negative correlation. We
533 adjust the ratio to bring the correlation between states as close to zero as possible. In Figure
534 3d, we show in simulation the ensuing benefit for sequence detection. An alternative method
535 is penalizing covariance between states in the classifier’s cost function ³⁸.

536

537 **Regularization**

538

539 A key parameter in training high dimensional decoding models is the degree of regularization.
540 In sequence analysis, we are often interested in spontaneous reactivation of state
541 representations, as in replay. However, our decoding models are typically trained on
542 stimulus-evoked data, because this is the only time at which we know the ground truth of
543 what is being represented. This poses a challenge in so far as the models best suited for
544 decoding evoked activity at training may not be well suited for decoding spontaneous activity
545 at subsequent tests. Regularising the classifier (for example with an L1 Norm) is a common
546 technique for increasing out-of-sample generalisation (to avoid overfitting). Here it has the
547 added potential benefit of reducing spatial correlation between classifier weights.

548

549 During classifier training, we can impose L1 or L2 constraints over the inference of classifier
550 coefficients, W . This amounts to finding the coefficients, W that maximise the likelihood of
551 the data observations, under the constraint imposed by the regularization term. L1
552 regularization can be phrased as maximising the likelihood, subject to a regularisation penalty
553 on the L1 norm of the coefficient vector:

554

$$555 \quad W = \operatorname{argmax}_W [\log(P(Y|M,W)) + b \lambda_{L1} \|W\|_1]$$

556 (10)

557

558 L2 regression can be viewed as a problem of maximising the likelihood, subject to a
559 regularisation penalty on the L2 norm of the coefficient vector:

560

561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581

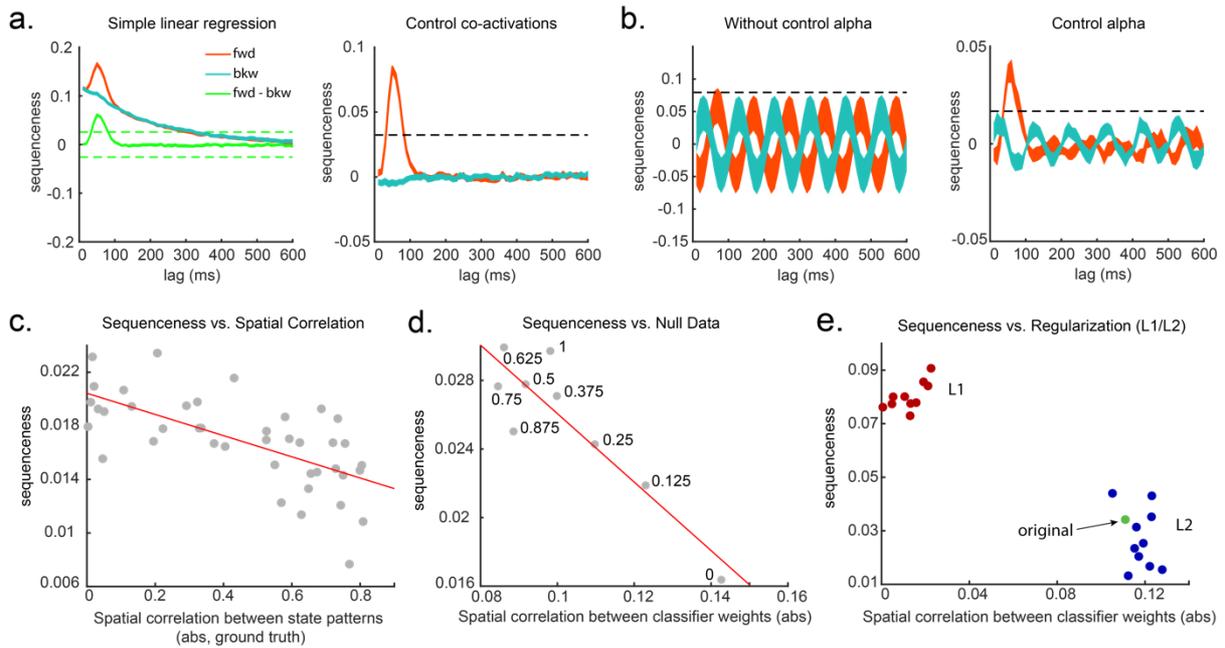
$$W = \operatorname{argmax}_W [\log(P(Y|M, W)) + b \lambda_{L2} \|W\|_2]$$

(11)

Where M is the task data, with dimension of number of observations, b , by number of sensors, s . Y is the label of observations, a vector with dimension of b by 1. $P(Y|M, W) = \sigma(MW)$, and σ is the logistic sigmoid function.

We simulate data with varying numbers of true sequences at 40 ms lag, and we find the beta estimate of sequence strength at 40 ms positively relates to the number of sequences. We also find that L1 weight regularization is able to detect sequences more robustly than L2 regularization, while L2 performs no better than an unregularized model (Figure 3e). The L1 models also have much lower spatial correlation, consistent with L1 achieving better sequence detection by reducing the covariances between classifiers³⁹.

In addition to minimizing spatial correlations, as discussed above, it can also be shown that L1-induced sparsity encodes weaker assumptions about background noise distributions into the classifiers, as compared to L2 regularization³⁹. This might be of special interest to researchers who want to measure replay during sleep. Here, the use of sparse classifiers is helpful as background noise distributions are likely to differ more substantially from the (awake state) training data.



582
583
584
585
586
587
588
589
590
591
592
593
594

Figure 3. Effects of temporal, spatial correlations, and classifier regularization on TDLM. **a**, Simple linear regression or cross-correlation approach relies on an asymmetry of forward and backward transitions; therefore, subtraction is necessary (left panel). TDLM instead relies on multiple linear regression. TDLM can assess forward and backward transitions separately (right panel). **b**, Background alpha oscillations, as seen during rest periods, can reduce sensitivity of sequence detection (left panel), controlling alpha in TDLM helps recover the true signal (right panel). **c**, The spatial correlation between the sensor weights of decoders for each state reduces the sensitivity of sequence detection. This suggests reducing overlapping patterns between states is important for sequence detection. **d**, Adding null data to the training set increases the sensitivity of sequence detection by reducing the spatial correlations of the trained classifier weights. Here the number indicates the ratio between null data and task data. “1” means the same amount of null data and the

595 task data. “0” means no null data is added for training. **e**, L1 regularization helps sequence detection
596 by reducing spatial correlations (all red dots are L1 regularization with a varying parameter value),
597 while L2 regularization does not help sequenceness (all blue dots are L2 regularization with a varying
598 parameter value) as it does not reduce spatial correlations of the trained classifiers compared to the
599 classifier trained without any regularization (green point).

600
601
602
603
604
605
606
607

608 **STATISTICAL INFERENCE**

609
610
611
612

So far, we have shown how to quantify sequences in representational dynamics. An essential final step is assessing the statistical reliability of these quantities.

613
614
615
616
617
618
619

All the tests described in this section evaluate the consistency of sequences *across subjects*. This is very important, because even in the absence of any real sequences of task-related representations, spontaneous neural activity is not random but follows repeating dynamical motifs⁴⁰. Solving this problem requires a randomized mapping between the assignment of physical stimuli to task states. This can be done across subjects, permitting valid inference at the group level.

620
621
622
623
624
625

At the group level, the statistical testing problem can be complicated by the fact that sequence measures do not in general follow a known distribution. Additionally, if a state-to-state lag of interest (Δt) is not known a priori, it is then necessary to perform tests at multiple lags, creating a multiple comparisons problem over a set of tests with complex interdependencies. In this section we discuss inference with these issues in mind.

626
627

626 **Distribution of sequenceness at a single lag**

628
629
630
631
632
633

If a state-to-state lag of interest (Δt) is known a priori then the simplest approach is to compare the sequenceness against zero, for example using either a signed-rank test, or one-sample t test (assuming Gaussian distribution). Such testing assumes the data are centred on zero if there were no real sequences. We show this approach is safe in both simulation (assuming no real sequences) and real MEG data where we know there are no sequences.

634
635
636
637
638
639
640

In simulation, we assume no real sequences, but state time courses are autocorrelated. At this point, there is no systematic structure in the correlation between the neuronal representations of different states (see later for this consideration). We then simply select the 40 ms time lag and compare its sequenceness to zero using either a signed-rank test or one-sample t test. We compare false positive rates predicted by the statistical tests with false positive rates measured in simulation (Figure 4a). We see the empirical false positives are well predicted by theory.

641
642
643
644
645

We have tested this also on real MEG data. In Liu, et al.²⁴ we had one condition where we measured resting activity before the subjects saw any stimuli. Therefore, by definition these sensory stimuli could not be replayed, we can use classifiers from these stimuli (measured later) to test a false positive performance of statistical tests on replay. Note, in our case, each

646 subject saw the same stimuli in a different order. They could not know the correct stimulus
647 order when these resting data were acquired. These data provide a valid null for testing false
648 positives.

649

650 To obtain many examples, we randomly permute the 8 different stimuli 10,000 times and
651 then compare sequenceness (at 40 ms time lag) to zero using either a signed rank test or one-
652 sample t test across subjects. Again, predicted and measured false positive rates match well
653 (Figure 4b, left panel). This holds true across all computed time lags (Figure 4b, right panel).

654

655 An alternative to making assumptions about the form of the null distribution is to compute an
656 empirical null distribution by permutation. Given that we are interested in the sequence of
657 states over time, one could imagine permuting either state identity or time. However,
658 permuting time uniformly will typically lead to a very high incidence of false positives, as
659 time is not exchangeable under the null hypothesis (Figure 4c, blue colour). Permuting time
660 destroys the temporal smoothness of neural data, creating an artificially narrow null
661 distribution^{24,35}. This false positive also exists if we circular shift the time dimension of each
662 state, rather than randomly permuting the state identities. This is because the signal is highly
663 non-stationary. Replays come in bursts, as recently analysed⁶, and this will break a circular
664 shift⁴¹. State permutation, on the other hand, only assumes state identities are exchangeable
665 under the null hypothesis, while preserving the temporal dynamics of the neural data,
666 represents a safer statistical test that is well within 5% false positive rate (Figure 4c, purple
667 colour).

668

669 **Correcting for multiple comparisons**

670

671 If the state-to-state lag of interest is not known, we have to search over a range of time lags.
672 As a result, we then have a multiple comparison problem. Unfortunately, we don't as yet
673 have a good parametric method to control for multiple testing over a distribution. It is
674 possible that one could use methods that exploit the properties of Gaussian Random Fields, as
675 is common in fMRI⁴², but we have not evaluated this approach. Alternatively, we could use
676 Bonferroni correction, but the assumption that each computed time lag is independent is
677 likely false and overly conservative.

678

679 We recommend relying on state-identity based permutation. To control for the family wise
680 error rate (assuming $\alpha = 0.05$), we want to ensure there is a 5% probability of getting the
681 tested sequenceness strength (S_{test}) or bigger by chance in *any* of the multiple tests. We
682 therefore need to know what fraction of the permutations give S_{test} or bigger in any of *their*
683 multiple tests. If any of the sequenceness scores in each permutation exceed S_{test} , then the
684 **maximum** sequenceness score in the permutation will exceed S_{test} , so it is sufficient to test
685 against the maximum sequenceness score in the permutation. The null distribution is
686 therefore formed by first taking the peak of sequenceness across all computed time lags of
687 each permutation. This is the same approach as used for family-wise error correction for
688 permutations tests in fMRI data⁴³, and in our case it is shown to behave well statistically
689 (Figure 4d).

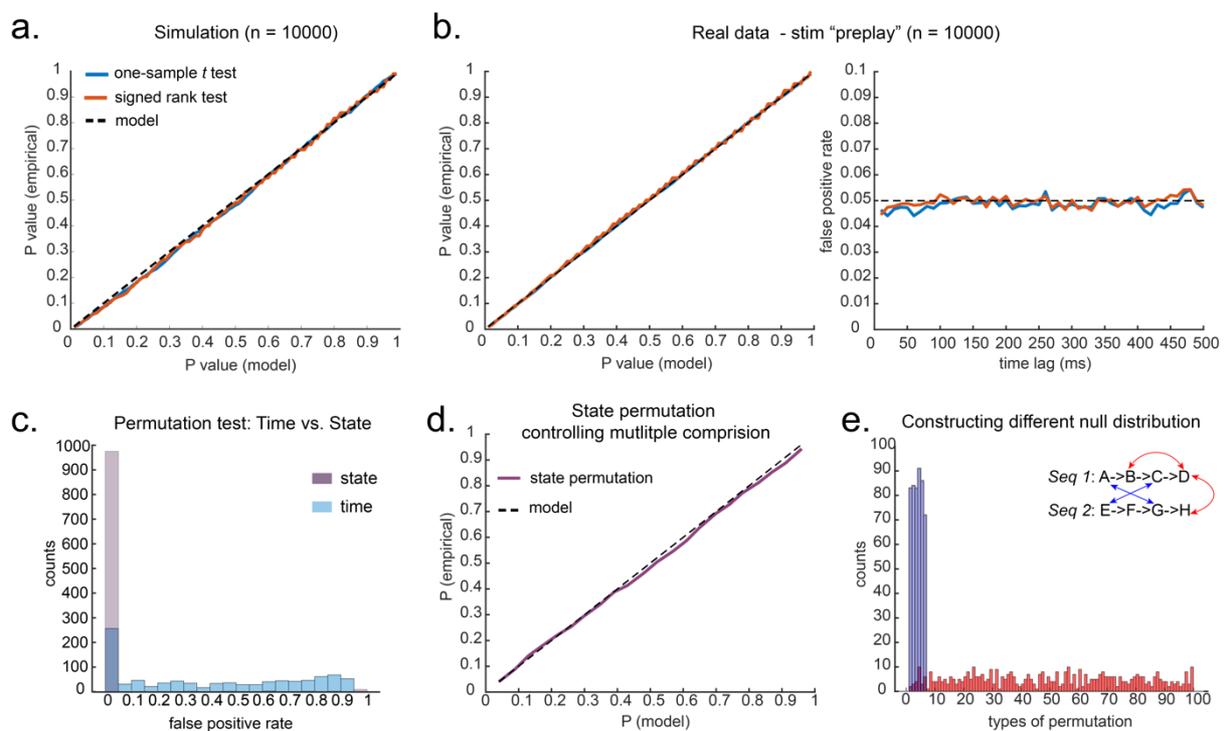
690

691 **What to permute**

692

693 We can choose which permutations to include in the null distribution. For example, consider
694 a task with two sequences, *Seq1*: $A \rightarrow B \rightarrow C \rightarrow D$, and *Seq2*: $E \rightarrow F \rightarrow G \rightarrow H$. We can
695 form the null distribution either by permuting all states (e.g., one permutation might be:

696 $E \rightarrow F \rightarrow A \rightarrow B, H \rightarrow C \rightarrow E \rightarrow D$), as implemented in Kurth-Nelson, et al.³⁵. Alternatively,
 697 we can form a null distribution which only includes transitions between states in different
 698 sequences (e.g., one permutation might be: $D \rightarrow G \rightarrow A \rightarrow E, H \rightarrow C \rightarrow F \rightarrow B$), as
 699 implemented in Liu, et al.²⁴. In each case, permutations are equivalent to the test data under
 700 the assumption that states are exchangeable between positions and sequences. The first case
 701 has the advantage of many more possible permutations, and therefore may make more precise
 702 inferential statements in the tail. The second may be more sensitive in the presence of a signal,
 703 as the null distribution is guaranteed not to include permutations which share any transitions
 704 with the test data (Figure 4e). For example, in the Figure 4e, the blue swaps are the
 705 permutations that only exchange state identity across sequences, as in Liu, et al.²⁴; while the
 706 red swaps are the permutations that permit all possible state identity permutations, as in
 707 Kurth-Nelson, et al.³⁵. Note there are many more different state permutations in red swaps
 708 than in blue swaps. We can make different levels of inferences by controlling the range of the
 709 null distributions in the state permutation tests.
 710



711
 712
 713 **Figure 4. Statistical inference.** **a**, P-P plot of one-sample t test (blue) and Wilcoxon signed rank test
 714 (red) against zero. This is performed in simulated MEG data, assuming auto-correlated state time
 715 courses, but no real sequences. In each simulation, the statistics are done only on sequenceness at 40
 716 ms time lag, across 24 simulated subjects. There are 10,000 simulations. **b**, We have also tested the
 717 sequenceness distribution on real MEG data. Illustrated is the pre-task resting state on 22 subjects
 718 from Liu et. al, where the ground truth is the absence of sequences given the stimuli have not yet been
 719 shown. The statistics are done on sequenceness at 40 ms time lag, across the 22 subjects. There are
 720 eight states. The state identity is randomly shuffled 10,000 times to construct a null distribution. **c**,
 721 Time-based permutation test tends to result in high false positive, while state identity-based
 722 permutation does not. This is done in simulation assuming no real sequences (n=1000). **d**, P-P plot of
 723 state identity-based permutation test over peak sequenceness is shown. To control for multiple
 724 comparisons, the null distribution is formed taking the maximal absolute value over all computed time
 725 lags within a permutation, and the permutation threshold is defined as the 95% percentile over
 726 permutations. In simulation, we only compared the max sequence strength in the data to this
 727 permutation threshold. There are 10,000 simulations. In each simulation, there are 24 simulated
 728 subjects, with no real sequence. **e**, In state-identity based permutation, we can test more specific

729 hypotheses by controlling the null distribution. Blue are the permutations that only exchange state
730 identity across sequences. Red are the permutations that permit all possible state identity permutations.
731 500 random state permutations are chosen from all possible ones. The X axis is the different
732 combinations of the state permutation. It is sorted so that the cross-sequence permutations are in the
733 beginning.

734
735

736

737 **Cautionary note on exchangeability of states after training**

738

739 Until now, all non-parametric tests have assumed that state identity is exchangeable under the
740 null hypothesis. Under this assumption, it is safe to perform state-identity based permutation
741 tests on Z_F and Z_B . In this section, we consider a situation where this assumption is broken.

742

743 More specifically, take a situation where the neural representation of state A and B are related
744 in a systematic way or, in other words, the classifier on state A is confused with state B , and
745 we are testing sequenceness of $A \rightarrow B$. Crucially, to break the exchangeability assumption,
746 representations of A and B have to be systematically more related than other states, e.g., A
747 and D . This cannot be caused by low level factors (e.g., visual similarity) because states are
748 counterbalanced across subjects, so any such bias would cancel at the population level.
749 However, such a bias might be *induced* by task training.

750

751 In this situation, it is, in principle, possible to detect sequenceness of $A \rightarrow B$, even in the
752 absence of real sequences. In the autocorrelation section above, we introduced protections
753 against the interaction of state correlation with autocorrelation. These protections may fail in
754 the current case as we cannot use other states as controls (as we do in the multiple linear
755 regression), because A has systematic relationship with B , but not other states. State
756 permutation will not protect us from this problem because state identity is no longer
757 exchangeable.

758

759 Is this a substantive problem? After extensive training, behavioural pairing of stimuli can
760 indeed result in increased neuronal similarity^{44,45}. These early papers involved long training
761 in monkeys. More recent studies have shown induced representational overlap in human
762 imaging within a single day^{28,46,47}. However, when analysed across the whole brain, such
763 representational changes tend to be localised to discrete brain regions^{48,49}, and as a
764 consequence may have limited impact on whole brain decodeability.

765

766 Whilst we have not yet found a simulation regime in which false positives are found (as
767 opposed to false negatives), there exists a danger in cases where, by experimental design, the
768 states are not exchangeable.

769

770

771

772 **SOURCE LOCALIZATION**

773

774 Uncovering temporal structure of neural representation is important, but it is of interest to ask
775 where in the brain a sequence is generated. Rodent electrophysiology research focuses mainly
776 on the hippocampus when searching for replay. One advantage of whole-brain non-invasive
777 neuroimaging over electrophysiology (despite many known disadvantages, including poor
778 anatomical precision, low signal-noise ratio) is in its ability to examine neural activity in

779 multiple other brain regions. Ideally, we would like a method that is capable of localizing
780 sequences of more abstract representation in brain regions beyond hippocampus²⁴.

781

782 We want to identify the *time* when a given sequence is very likely to unfold, so we can
783 construct averages of independent data over these times. We achieve this, by transforming
784 from the space of original states, X_{orig} , to the space of sequence events, X_{seq} . First, based on
785 the transition of interest, T , we can obtain the projection matrix, X_{proj} :

786

$$787 \quad X_{proj} = X_{orig} \times T \quad (12)$$

788

789 If we know the state lag within sequence, Δt (e.g., the time lag give rise to the strongest
790 sequenceness) or have it a priori. We can obtain the time lagged matrix, X_{lag} :

791

$$792 \quad X_{lag} = X_{orig}(t - \Delta t) \quad (13)$$

793

794 Then, we obtain state space with sequence event as states by elementwise multiply X_{proj} and
795 X_{lag} :

$$796 \quad X_{seq} = X_{lag} .* X_{proj} \quad (14)$$

797

798 Each element in X_{seq} indicates the strength of a (pairwise) sequence at a given moment in
799 time. At this stage, X_{seq} is a matrix with number of time points as rows (same as X_{orig}), and
800 with number of pairwise sequences (e.g., A->B; B->C; etc) as columns. Now on this matrix,
801 X_{seq} , we can either look for sequences of sequences (see in Appendix 3), or sum over
802 columns (i.e., average over pairwise sequence events), and obtain a score at each timepoint
803 reflecting how likely it is to be a sequence member (Figure 5a).

804

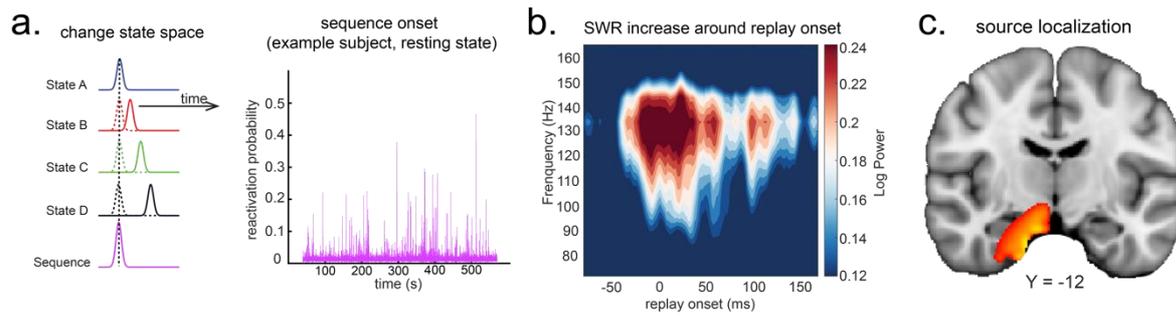
805 We can use this score to construct averages of other variables that might co-vary with replay.
806 For example, if we choose timepoints when this score is high (e.g., 95th) percentile after
807 being low for the previous 100 ms and construct an average time-frequency plot of the raw
808 MEG data aligned to these times, we can reconstruct a time-frequency plot that is, *on average*,
809 associated with replay onset (Figure 5b). Note that although this method assigns a score for
810 individual replay events as an intermediary variable, it results in an *average* measure across
811 many events.

812

813 This approach is similar to spike-triggered averaging^{50,51}. Applying this to real MEG data
814 during rest, we can detect increased hippocampal power at 120-150 Hz, at replay onset
815 (Figure 5b, c). Source reconstruction in the current analysis was performed using linearly
816 constrained minimum variance (LCMV) beamforming, a common method for MEG source
817 localization, but it is known to suffer from distal correlated sources⁵². A better method may
818 be Empirical Bayesian Beamformer for accommodating correlated neural source as a priori⁵³.

819

820



821
822

Figure 5. Source localization of replay onset. **a**, TDLM indexes the onset of a sequence based on the identified optimal state-to-state time lag (left panel). Sequence onset during resting state from one example subject is shown (right panel). **b**, There was a significant power increase (averaged across all sensors), in the ripple frequency band (120-150 Hz), at the onset of replay, compared to the pre-replay baseline (100 to 50 ms before replay). **c**, Source localization of ripple-band power at replay onset revealed significant hippocampal activation (peak MNI coordinate: X = 18, Y = -12, Z = -27). Panel b and c is reproduced from Figure 7 A, C, Liu et al. 2019, Cell, published under the Creative Commons Attribution 4.0 International Public License (CC BY 4.0).

831
832
833
834
835
836

837 TDLM FOR RODENT REPLAY

838
839
840
841
842
843
844
845

So far, we have introduced TDLM in the context of analysing human MEG data. Relatedly, its application on human EEG data was also explored (Appendix 4: Apply TDLM to human whole-brain EEG data). Historically, replay-like phenomena have been predominantly studied in rodents with electrophysiology recordings in the hippocampal formation^{16,17,20}. This raises an interesting question: how does TDLM compare to the existing rodent replay methods; can TDLM be applied to spiking data for detecting rodent replays, and what are the pros and cons? In this section we address this question.

846
847

848 Generality of graph- vs line-based replay methods

849
850
851
852
853
854
855

Given TDLM works on the decoded state space, rather than sensor (with analogy to cell) level, we compared TDLM to rodent methods that work on the posterior decoded position (i.e., state) space, normally referred to as Bayesian-based methods²⁰ (Note that these methods are typically Bayesian in how position is decoded from spikes⁵⁴ but not in how replay is measured from decoded position). Two commonly used methods are Radon transform¹⁶ and linear weighted correlation¹⁷.

856
857
858
859
860
861

Both methods proceed by forming a 2D matrix, where one dimension is the decoded state (e.g., positions on a linear track), and the other dimension is time (note that the decoded state is embedded in 1D). The methods then try to discover if an ordered line is a good description of the relationship between state and (parametric) time. For this reason, we call this family of approaches “line search”.

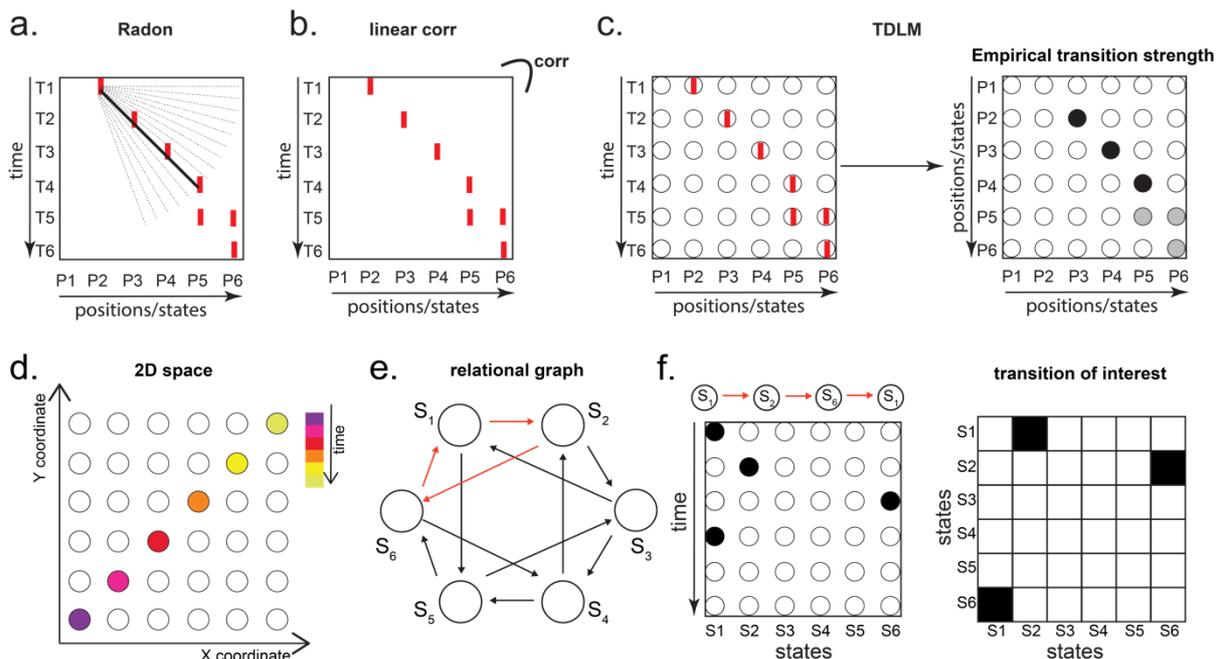
862
863
864

The radon method uses a discrete Radon transform to find the best line in the 2D matrix⁵⁵ and then evaluates the radon integral, which will be high if the data lie on a line (Figure 6a).

865 It compares this to permutations of the same data where the states are reordered ²⁰. The linear
 866 weighted correlation method computes the average correlation between the time and
 867 estimated position in the 1D embedding (Figure 6b). The correlation is non-zero provided
 868 there is an orderly reactivation along the state dimension.

869
 870 Both methods are applied to decoded positions, where they are sorted based on the order in a
 871 linearized state space. TDLM also works on the decoded position space, but instead of
 872 directly measuring the relationship between position and time, it measures the transition
 873 strength for each possible state to state transitions (Figure 6c).
 874

875 This is a key difference between TDLM and these popular existing techniques. To reiterate,
 876 the latter rely on a continuous parametric embedding of behavioural states and the
 877 relationship between this embedding and time (parametrically encoded). TDLM is
 878 fundamentally different as it works on a graph and examines the statistical likelihood of some
 879 transitions happening more than others. This is therefore a more general approach that can be
 880 used for sequences drawn from any graph (e.g., 2D maze, Figure 6d), not just graphs with
 881 simple embeddings (like a linear track). For example, in a non-spatial decision-making task ³⁵,
 882 all states lead to two different states and themselves can be arrived at from two other different
 883 states (Figure 6e). Existing “line search” methods will not work because there is no linear
 884 relationship between time and states (Figure 6f).
 885



886
 887
 888 **Figure 6. TDLM vs. existing rodent replay methods.** **a.** The Radon method tries to find the best
 889 fitting line (solid line) of the decoded positions as a function of time. The red bars indicate strong
 890 reactivation at a given location. **b.** The linear correlation method looks for correlations between time
 891 and decoded position. **c.** The TDLM method, on the other hand, does not directly measure the
 892 relationship between state and time, but quantifies the likelihood of each transition. In the right panel,
 893 likelihood is indicated by darkness of shading. For example, P5 can be followed by either P5 or P6,
 894 making each transition half as likely as the deterministic P4->P5 transition. Later this empirical
 895 transition matrix is compared to a theoretical one, to quantify the extent to which the empirical
 896 transitions fit with a hypothesis. **d.** Sequences in 2D space is in three dimensions, which is hard to
 897 translate into a line search problem, e.g., time*position spaces. **e.** This is the transition matrix used in
 898 Kurth-Nelson, et al. ³⁵, which cannot be translated into a linear state space. The transitions in red are

899 an example of a trajectory. **f.** Putting the example trajectory into the time by state matrix, we can see
900 there is no linear relationship between them (left panel). In TDLM, this is tested by forming a
901 hypothesis regressor in the state-to-state transition matrix (right panel).

902

903

904

905 **Multi-scale TDLM**

906

907 While continuous spaces can be analysed in TDLM by simply chunking the space into
908 discrete states, TDLM in its original form may potentially be less sensitive for such analyses
909 than techniques with build-in assumptions about the spatial layout of the state space, such as
910 the linear relationship between time and reactivated states (Appendix 5 “Less sensitivity of
911 TDLM to skipping sequences”). In essence, because TDLM works on a graph, it has no
912 information about the Euclidean nature of the state space, while techniques that make
913 assumptions about the linear relationship between space and time benefit from these
914 assumptions. For example, detecting state 1 then state 5 then state 10 counts as replay in
915 these techniques, but not in TDLM.

916

917 However, TDLM can be extended to address this problem. For continuous state spaces, we
918 first need to decide how to best discretise the space. If we choose a large scale, we will miss
919 replays that occur predominantly within a spatial bin. If we choose a small scale, we will
920 miss transitions that jump spatial bins. A simple solution is to apply TDLM at multiple
921 different scales and take a (variance-weighted) average of the sequenceness measures at
922 different scales. For example, when measuring replay speed, we can average events that
923 travel 5 cm in 10 ms together with events that travel 10 cm in 20 ms.

924

925 Specifically, to perform multi-scale TDLM, we discretise position bins at multiple widths.
926 This generates rate maps at multiple scales (e.g., 5 cm, 10 cm, 20 cm, 40 cm), and hence a
927 multi-scale state space. For each replay speed of interest, we apply TDLM separately at each
928 scale, and then take a variance-weighted average of replay estimates over all scales.

929

$$930 \quad \beta_M = \frac{\sum_{i=1}^n \beta_i / V_i}{\sum_{i=1}^n 1 / V_i}$$

931 (15)

932

933 Where β_i is the sequence strength of given speed (i.e., state-to-state lag) measured at scale i ,
934 V_i is the variance of its β_i estimator, and n is the number of scales. In the end, statistical
935 testing is performed on the precision weighted averaged sequence strength, β_M , in the same
936 way as we do in the original TDLM.

937

938 It is easy to see why this addresses the potential concerns raised above as some scales will
939 capture the 1->2->3 transitions, whilst others will capture the 1->10->20 transitions: Because
940 the underlying space is continuous, we can average results of the same replay speed together,
941 and this will reinstate the Euclidean assumptions.

942

943

944 **Applying multi-scale TDLM to real rodent data (place cells in CA1)**

945

946 We demonstrate the applicability of multi-scale TDLM by analyzing CA1 place cell spiking
947 data from Ólafsdóttir, et al. ²⁶. In Ólafsdóttir, et al. ²⁶, rats ran multiple laps on a 600 cm Z
948 maze, and were then placed in a rest enclosure for 1.5 hours (Figure 7a). The Z maze consists

949 of 3 tracks, with its ends and corners baited with sweetened rice to encourage running from
950 one end to the other. The animal's running trajectory was linearized, dwell time and spikes
951 were binned into 2 cm bins and smoothed with a Gaussian kernel ($\sigma = 5$ bins). We generated
952 rate maps separately for inbound (track1->track2->track3) and outbound (track3->track2-
953 >track1) running (see details in section "Rodent Replay dataset" in the Methods).

954
955 As in Ólafsdóttir, et al. ²⁶, cells recorded in CA1 were classified as place cells if their peak
956 firing field during track running was above 1 Hz with a width of at least 20 cm (see an
957 example in Figure 7b). The candidate replay events were identified based on multi-unit (MU)
958 activity from place cells during rest time. Periods exceeding the mean rate by 3 standard
959 deviations of MU activity were identified as possible replay events. Events less than 40 ms
960 long, or which included activity from less than 15% of the recorded place cell ensemble, were
961 rejected (see an example of putative replay event in Figure 7c), and the remaining events
962 were labelled putative replay events.

963
964 We analyzed data from one full recording session (track running for generating rate map,
965 post-running resting for replay detection) from Rat 2192 reported in Ólafsdóttir, et al. ²⁶.
966 Following the procedure described above, we identified 58 place cells, and 1183 putative
967 replay events. Replay analysis was then performed on the putative replay events, separately
968 for inbound and outbound rate maps given the same position has a different decoded state
969 depending on whether it was during an outbound or inbound run.

970
971 A forward sequence is characterised by states from the outbound map occurring in the
972 outbound order, or states from the inbound map occurring in the inbound order. Conversely, a
973 backward sequence is when states from the inbound map occur in the outbound order or
974 states from the outbound map occur in the inbound order. Candidate events were decoded
975 based on a rate map, transforming the ncells * ntime to nstates * ntime. Each entry in this
976 state space represents the posterior probability of being in this position at a given time.
977 Replay analysis was performed solely on this decoded state space.

978
979 Note, TDLM is applied directly to the concatenated rather than individual replay events. This
980 is because TDLM is a linear modelling framework. Applying TDLM on each single replay
981 event, and then averaging the beta estimates (appropriately weighted by the variances) is
982 equivalent to running TDLM once on the concatenated replay events. It quantifies the
983 average amount of replay across many events, this is different compared to existing replay
984 methods that focus on single replay events. Because TDLM addresses statistical questions in
985 linear modelling, it does not require secondary statistics to ask whether the "counts" of
986 individual events are more likely than chance, or more likely in one situation than another.

987
988 During the whole sleep period, TDLM identified a significant forward sequence for the
989 outbound map with a wide speed range around from 1 to 10 m/s (Figure 7d, left panel),
990 consistent with recent findings from Denovellis, et al. ⁵⁶ on varying replay speed (similar
991 results were obtained for inbound map, not shown here for simplicity). In our analysis, the
992 fastest speed is up to 10 m/s, which is around 20X faster than its free running speed,
993 representing approximately half a track-arm in a typical replay event, consistent with
994 previous work ^{10,16,57,58}.

996 **Second order inferences**

997 As pointed out by van der Meer, et al. ¹⁹, there are two types of statistical questions: a "first-
 998 order" sequence question, which concerns whether an observed sequenceness is different
 999 from random (i.e., do replays exist?); and a "second-order" question, which requires a
 1000 comparison of sequenceness across conditions (i.e., do replays differ?). Because it is
 1001 embedded in a linear regression framework, TDLM is ideally placed to address such
 1002 questions. There are two ways of asking such questions in linear modelling - **Contrasts** and
 1003 **Interactions**. We explain them with examples here.

1004 1005 **Linear contrasts**

1006 After fitting a regression model, resulting in coefficients for different regressors, we can test
 1007 hypotheses about these coefficients by constructing linear combinations of the coefficients
 1008 that would be zero under the null hypothesis. For example, if we want to test whether effect
 1009 **A** is greater than effect **B** then we can compute the linear contrast **A - B** (which would be
 1010 zero under the null hypothesis) and perform statistics on this new measure. If we want to test
 1011 whether replay increases linearly over 5 conditions [A, B, C, D, E], we can compute
 1012 the linear contrast $-2*A - B + 0*C + D + 2*E$, (which would be zero under the null
 1013 hypothesis) and perform statistics on this new measure. Statistics (within or across animals)
 1014 can operate with these contrasts in exactly the same way as with the original coefficients
 1015 from the linear model. Here we demonstrate this by showing in our example data set that
 1016 there was a greater preponderance for forward than backward replay. We construct the
 1017 contrast (Forwards - Backwards) and test it against zero using a multiple-comparison-
 1018 controlled permutation test (Figure 7d, right panel, pink line). By constructing a different
 1019 contrast (Forwards + Backwards), we can also show that the total replay strength across both
 1020 types of replays was significant (Figure 7d, right panel, green line).

1021 1022 **Interactions**

1023 A second method for performing second order tests is to introduce them into the linear
 1024 regression as interaction terms, and then perform inference on the regression weights for
 1025 these interactions. This means changing equation 2 to include new regressors. For example, if
 1026 interested in how reactivations change over time, one could build new regressors
 1027 ($Xtime_k(t)$), obtained by element-wise multiplying the state regressor, e.g., $X_k(t)$ with time
 1028 indices ($Xtime_k(t) = X_k(t) * time$). Now the first level GLM is constructed as (omitting
 1029 residual term ϵ , same as equation 2):

$$1030$$

$$1031 X_j(t + \Delta t) = \sum_{k=1}^n X_k(t)\beta_{kj} + Xtime_k(t)\beta t_{kj} \quad (16)$$

$$1032$$

1033 Example regressors in the design matrix can be seen in Figure 7e below. The first regressor,
 1034 $X_k(t)$, is one of the state reactivation regressors used in standard TDLM. The second
 1035 regressor, $Xtime_k(t)$, is this same as $X_k(t)$ multiplied by time. (There are k regressors of
 1036 each form in regressor matrix.) Here, we chose to demean the time regressor before the
 1037 interaction, so the early half of the regressor is negative, and the late half is positive. This has
 1038 no effect on the regression coefficients of the interaction term but, by rendering the
 1039 interaction approximately orthogonal to $X_k(t)$, it makes it possible to estimate the main
 1040 effect and the interaction in the same regression.

1041
 1042 Note that the interaction regressor is **orthogonal** to the state reactivation regressor, so it will
 1043 have no effect on the first order regression terms. If we include such regressors for all states,
 1044 then we can get two measures for each replay direction (sequence effect and time effect). The
 1045 first tells us the average amount of replay throughout the sleep period (first order). The

1046 second tells us whether replay increases or decreases as time progresses through the sleep
 1047 period (second order).

1048

1049 **Orthogonal tests in regions of interest**

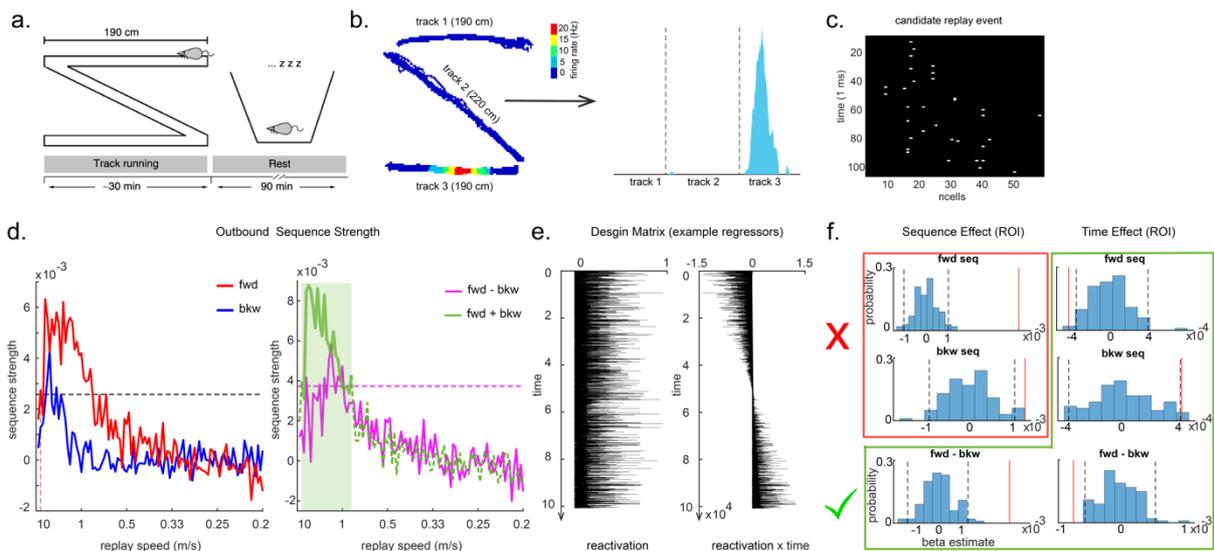
1050 When examining Forward-Backward replay above, we did separate inference for each replay
 1051 speed, and then performed multiple comparison testing using the max-permutation method
 1052 (see statistical inference section above). We now take the opportunity to introduce another
 1053 method common in human literature.

1054

1055 To avoid such multiple comparison correction, it is possible to select a “Region of Interest”
 1056 (ROI), average the measure in question over that ROI, and perform inference on this average
 1057 measure. Because we are now only testing one measure, there is no multiple comparison
 1058 problem. Critical in this endeavour, however, is that we do not use the measure under test, or
 1059 anything that correlates with that measure as a means to define the ROI. This will induce a
 1060 selection bias⁵⁹. In the example in Figure 7f, we have used the average replay (Forwards
 1061 +Backwards) to select the ROI. We are interested in speeds in which there is detectable
 1062 replay on average across both directions and the whole sleep period (Figure 7d, right panel,
 1063 green shaded area). If we select our ROI in this way, we cannot perform unbiased inference
 1064 on first order Forwards or Backwards replay because Forwards and Backwards regressors
 1065 correlate with their sum (Figure 7d, statistical inference in the red rectangle is
 1066 biased). However, we **can** perform unbiased inference on several second order effects (Figure
 1067 7d, statistical inference in the green rectangle). We can test (Forwards - Backwards)
 1068 assuming the difference of terms is orthogonal to their sum (as it is in this case). Further we
 1069 can test any interaction with time, because the ROI is defined on the average over time, and
 1070 the interaction looks for *differences* as a function of time. When we perform these tests in our
 1071 example dataset (Figure 7d, green rectangle), we confirm that there are more forward than
 1072 backward replay on average. We further show that forward replay is decreasing with time
 1073 during sleep, and that backward replay is increasing with time. Their difference (Forwards -
 1074 Backwards) is also significant.

1075

1076



1077
 1078

1079 **Figure 7. TDLM applied to real rodent data.** **a**, The experimental design of Ólafsdóttir, et al.²⁶.
 1080 Rats ran on Z maze for 30 mins, followed by 90 min rest. **b**, An example rate map for a place cell is
 1081 shown. The left panel shows its spatial distribution on the Z maze, and the right panel is its linearized
 1082 distribution. **c**, An example of a candidate replay event (spiking data) is shown. **d**, Sequence strength

1083 as a function of replay speed is shown for the outbound rate map. Black dotted line is the permutation
1084 threshold after controlling for multiple comparisons. Left panel: forward sequence (red) and backward
1085 sequence (blue). The red dotted line indicates the fastest replay speed that is significant – 10 m/s.
1086 Right panel: forward – backward sequence. The pink dotted line indicates the multiple comparison
1087 corrected permutation threshold for the replay difference. The green line is the sum of sequence
1088 strength between forward and backward direction. The solid line (with green shading) indicates the
1089 significant replay speeds (0.88 – 10 m/s) after controlling for multiple comparisons. We use this as a
1090 ROI to test for time varying effect on replay in panel f. **e**, Illustration of two exemplar regressors in
1091 the design matrix for assessing time effect on replay strength. The “reactivation” regressor is a lagged
1092 copy of reactivation strength of given position and is used to obtain sequence effect. The “reactivation
1093 x time” regressor is the elementwise multiplication between this position reactivation and time (z-
1094 scored), it explicitly models the effect of time on sequence strength. Both regressors are demeaned. **f**,
1095 Beta estimate of the sequence effect (left panel), as well as time modulation effect on sequence (right
1096 panel) in the ROI are shown. Negative value indicates replay strength decreases over time, while
1097 positive value means replay increases as a function of sleep time. The statistical inference is done
1098 based on a permutation test. The two black dotted lines in each panel indicate the 2.5th and 97.5th
1099 percentile of the permutation samples, respectively. The red solid line indicates the true beta estimate
1100 of the effect. Note there is a selection bias in performing statistical inference on forward and
1101 backward sequence strength (red rectangle) within this ROI, given the sum of forward and backward
1102 sequence is correlated with either forward, or backward sequence alone. There is no selection bias in
1103 performing statistics on the difference of sequence effects, or effects relating to time (green rectangle).

1104
1105
1106
1107
1108

1109 In addition to the time varying effect, we can also test the spatial modulation effect, i.e., how
1110 replay strength (at the same replay speed) change as a function of its spatial content. For
1111 example, is replay stronger for transitions in the start of maze, compared to the end of the
1112 track. As an illustrative example, we have used the same ROI defined above, and test the
1113 spatial modulation effect on forward replay. Note this test of spatial modulation effect is also
1114 unbiased from the overall strength of forward replay, and thereby no selection bias in this
1115 ROI, as well.

1116 For visualization purposes, we have first plotted the estimated strength for each pairwise
1117 forward sequence (Figure R8a), separately within each scale (from 1 to 4, with increasing
1118 spatial scales). The pairwise sequences are ordered from the start of the maze to the end of
1119 the maze. Alongside the pair-wise sequence plot, we have plotted the mean replay strength
1120 over all possible pairwise transitions (in red), in comparison to the mean of all control
1121 transitions (in grey. As expected, they are all around 0). Note that we cannot perform
1122 inference on the difference between the red and grey bars here because they have been
1123 selected from a biased ROI. It is simply for illustration purposes. We have therefore put them
1124 in red squares to match Figure 7f.

1125

1126 To formally test the spatial modulation effect, we can use the exact same approach as
1127 outlined above in the **linear contrasts** section. Here, we test a linear increase or decrease
1128 across different *transitions*. We take the linear contrast weight vector, c ([-2,-1,0,1,2] for the
1129 largest scale, [-3:3] for the next scale, [-5:5] for the next scale, [-12:12] for the smallest scale)
1130 and multiply these by the beta estimates of the transitions:

1131

1132

1133

1134

$$contrast = c^T \beta$$

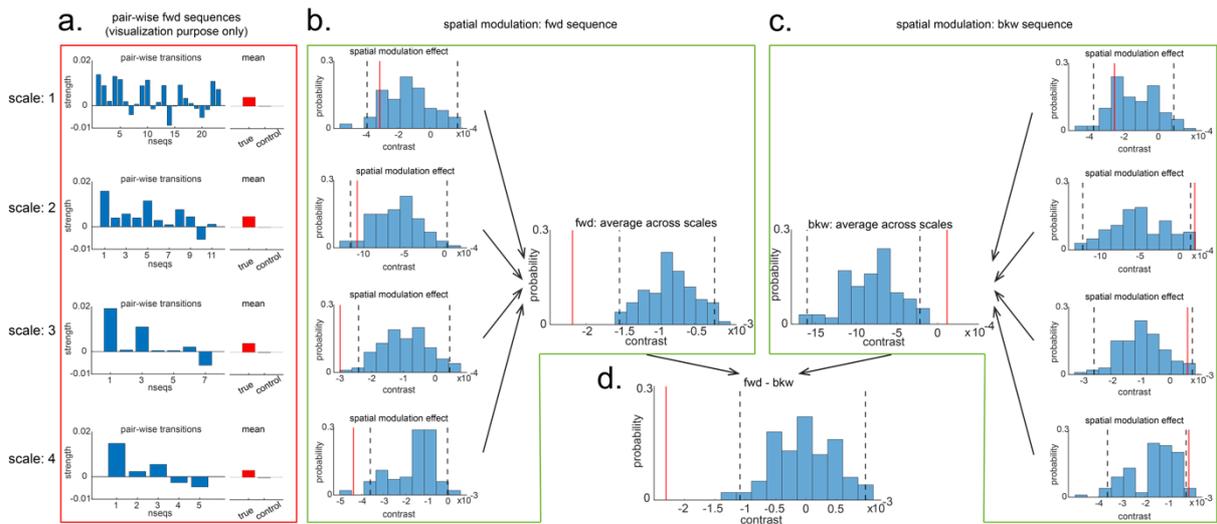
(17)

1135 If this new measure, *contrast*, is different from zero, then there is a linear increase/decrease
 1136 from one end of the track to the other. Note that this new contrast is no longer biased by the
 1137 ROI selection as each transition contributed equally to the ROI selection, but we are now
 1138 comparing between transitions. Inference on this contrast is therefore valid. We have
 1139 therefore put them in green boxes to match Figure 7f (Figure 8b, c).

1140
 1141 Within the larger two scales, these contrasts are significantly negative (tested against
 1142 permutations in exactly the same way as the “mean” contrasts). Since we are still in the linear
 1143 domain, we can now just average these contrasts across the 4 scales and get a single measure
 1144 for spatial modulation of replay. This average measure is significantly negative (Figure 8b).
 1145 Hence, on average, forward replay is stronger at the beginning of the track.

1146
 1147 We can do the same thing for backward replay. We found a opposite pattern, i.e., strength of
 1148 backward replay is stronger at the end of the track, and similarly, it is not significant in the
 1149 smallest scale, and become significant in the largest scale, and also significant on average
 1150 across all scales (Figure 8c). Again, since we are in the linear domain, we can further contrast
 1151 these contrasts, asking if this effect is different for forwards and backward replay. We found
 1152 the difference is indeed significant (Figure 8d). This set of results is consistent with previous
 1153 rodent literature⁶⁰. Note we would like to stress again, that this analysis is not about a single
 1154 replay event but is testing for average differences across all replay events.

1155
 1156



1157
 1158
 1159 **Figure 8. Pairwise sequence & spatial modulation effect.** **a.** Within each scale, strengths of each
 1160 pairwise forward sequences in the ROI (significant replay speeds, cf. Figure R7d, green shading) are
 1161 ordered from the start of maze to the end of the maze; alongside that, the mean sequence strength
 1162 across all of these valid pairwise transitions is plotted (red) in comparison to the mean of all control
 1163 transitions (grey). This is for visualization purpose only and is included in the red rectangle. **b.** The
 1164 contrast defining a linear change in forward sequenceness across the track (spatial modulation)
 1165 is shown (red line), both separately for each scale, and average across scales, and compared to
 1166 permutations. On average, forward replay is stronger at the beginning of the track. **c.** Same as panel b,
 1167 but this is for the backward sequences. Unlike forward replay, backward replay is stronger at the end
 1168 of the track. Note, both panel b and c are about spatial modulation effect, which is orthogonal to
 1169 overall sequence strength, allowing valid inference. They are therefore included in green boxes. **d.**
 1170 The difference of this spatial modulation effect between forward and backward sequence is also
 1171 significant. The black dotted lines indicate the 2.5th and 97.5th percentile of the permutation samples.
 1172 The red solid line indicates the estimate of the true contrast effect.

1173
1174
1175
1176
1177
1178
1179
1180
1181

Notably, extra care needs to be exercised for second-order questions (compared to first order ones). Problems can emerge due to biases in second order inference, such as in behavioral sampling (e.g., track 1 may be experienced more than track 2 during navigation. This creates a bias when evaluating replay in track 1 vs. track 2 during rest). Such issues are real but can be finessed by experimental design considerations of a sort commonly applied in the human literature. For example:

1182
1183
1184
1185
1186
1187

- (1) Ensure that biases that might occur within subjects will not occur consistently in the same direction *across subjects* (e.g., by randomising stimuli across participants).
- (2) Compare across conditions in each subject.
- (3) Perform a random effects inference across the population, by comparing against the between-subject variance.

1188
1189
1190
1191
1192

Such approaches are not yet common in rodent electrophysiology and may not be practical in some instances. In such cases, it remains important to be vigilant to guard against these biases with TDLM as with other techniques. If these approaches are feasible, the machinery for computing second-order inferences is straightforward in a linear framework like TDLM.

1193
1194

1195 **GENERALITY OF TDLM**

1196
1197
1198
1199
1200
1201
1202
1203
1204
1205

We have now discussed the applicability of TDLM in relation to human MEG, as well as in rodent electrophysiology (with comparisons to standard replay detection methods). A preliminary attempt at detecting replay in human EEG was also shown in the Appendix 4. We believe this establishes TDLM as a domain-general sequence analysis method: TDLM works at the level of decoded state space, rather than the sensor/cell level of the data. It can be applied to a wide range of data types and settings in both humans and rodents, stimulating cross-fertilization across disciplines. It is based on the GLM framework, and this lends it flexibility for regressing out potential confounds while offering an intuitive understanding of the overall approach.

1206
1207

In this section, we discuss the generality of TDLM.

1208
1209
1210
1211
1212
1213
1214

States: TDLM assesses the statistical likelihood of certain transitions on a graph. In its original form, TDLM works on discrete states (i.e., nodes in the graph). Continuous spaces can be incorporated by chunking them into discrete spaces. Furthermore, by averaging the same replay speeds measured at multiple scales of discretization (see section “TDLM FOR RODENT REPLAY”), the statistical benefits of an assumption of a Euclidean geometry can be recovered.

1215
1216
1217
1218
1219
1220
1221
1222

Time length: The longer the time length, the more accurate the estimates in TDLM. This is because TDLM assesses sequence evidence based on a GLM framework, where time length is the sample size. Higher sample size will lead to more accurate estimates. In the case of rodent analysis, we recommend applying TDLM to aggregated replay events rather than to a single event because this results in 1) more time samples for estimation; 2) more activated states in the analysis time framework. Unlike other techniques which search for a single replay in a single event, this aggregation can be implemented without losing generality, as

1223 TDLM is able to handle multiple sequences in the same data with respect to different
1224 directions, contents or speeds. Furthermore, by aggregating linearly across all replay events
1225 of the same condition, it provides a natural measure for comparing replay strength, speed and
1226 direction across different experimental conditions.

1227
1228 TDLM has already proved important in human experiments where complex state-spaces have
1229 been used^{22,24,25,35}. We expect this generality will also be important as rodent replay
1230 experiments move beyond 1D tracks, for example to foraging in 2D, or in complex mazes.

1231 1232 1233 **DISCUSSION**

1234
1235 TDLM is a domain general analysis framework for capturing sequence regularity of neural
1236 representations. It is developed on human neuroimaging data, and can be extended to other
1237 data sources, including rodent electrophysiology recordings. It offers hope for cross-species
1238 investigations on replay (or neural sequences in general), and potentially enable studies of
1239 complex tasks in both human and animals, e.g., complex 2D maze in rodents.

1240
1241 TDLM adds a new analysis toolkit to the replay field. It is especially suited for summarising
1242 replay strength across many events, for comparing replay strength between conditions, and
1243 for analysing replay strength in complex behavioural paradigms. Its linear modelling nature
1244 makes it amenable to standard statistical tests, and thereby allows wide use across task,
1245 modality, and species. Unlike alternative tools, we have not shown TDLM applied to
1246 individual replay events.

1247
1248 The temporal dynamics of neural states have been studied previously with MEG^{40,61}.
1249 Normally such states are defined by common physiological features (e.g., frequency,
1250 functional connectivity) during rest, and termed resting state networks (e.g., default mode
1251 network⁶²). However, these approaches remain agnostic about the *content* of neural activity.
1252 The ability to study the temporal dynamics of representational content permits richer
1253 investigations into cognitive processes⁶, as neural states can be analysed in the context of
1254 their roles with respect to a range of cognitive tasks.

1255
1256 Reactivation of neural representations have also been studied previously⁶³ using approaches
1257 similar to the decoding step of TDLM, or multivariate pattern analysis (MVPA)⁶⁴. This has
1258 proven fruitful in revealing mnemonic functions⁴⁷, understanding sleep⁶⁵, and decision-
1259 making⁶⁶. However, classification alone does not reveal the rich *temporal structures* of
1260 reactivation dynamics. We have described the application of TDLM mostly during off-task
1261 state in this paper. However, the very same analysis can be applied to on-task data, to test for
1262 cued sequential reactivation²², or sequential decision-making⁶⁷. For example, the ability to
1263 detect sequences on-task allows us to tease apart clustered from sequential reactivation,
1264 where this may be important for dissociating decision strategies³⁴ and their individual
1265 differences^{22,67}. TDLM, therefore, may allow testing of neural predictions from process
1266 models such as reinforcement learning during task performance⁶⁸, which have proved hard to
1267 probe previously in humans²²⁻²⁵.

1268
1269 In the human neuroimaging domain, we have mainly discussed the application of TDLM
1270 with respect to MEG data. In the Appendix 4, we show TDLM also works well with EEG
1271 data. This is not surprising given EEG and MEG are effectively measuring the same neural
1272 signature, namely local field potential (or associated magnetic field) on the scalp. We do not

1273 have suitable fMRI data to test TDLM. However, related work has suggested it might be
1274 possible to measure sequential reactivation using fMRI ⁶⁹, but particular methodological
1275 caveats need to be considered (e.g., a bias from last events due to slow hemodynamic
1276 response) ⁷⁰. We believe TDLM can deal with this, given it models out non-specific
1277 transitions, although further work is needed. In future, we consider it will be useful to
1278 combine the high temporal resolution available in M/EEG and the spatial precision available
1279 in fMRI to probe region specific sequential computations.

1280

1281 In the rodent electrophysiology domain, we show what TDLM (its multi-scale version) has to
1282 offer uniquely compared to existing rodent replay methods. Most importantly, TDLM works
1283 on an arbitrary graph and its generality makes replay studies in complex mazes possible. Its
1284 linear framework makes the assessment of time varying effect on replay (Figure 7), or other
1285 second-order sequence questions straightforward. In future work, a promising direction will
1286 be to further separate process noise (e.g., intrinsic variability within sequences) and
1287 measurement noise (e.g., noise in MEG recording). This might be achieved by building latent
1288 state-space models as have explored recently in rodent community ^{18,56}.

1289

1290 Together, we believe TDLM opens doors for novel investigations of human cognition,
1291 including language, sequential planning and inference in non-spatial cognitive tasks ^{34,35}, as
1292 well as complicated tasks in rodents, e.g., forging in 2D mazes. TDLM is particularly suited
1293 to test specific neural predictions from process models, such as reinforcement learning. We
1294 hope TDLM can promote an across species synthesis between experimental and theoretical
1295 neuroscience and, in so doing, shed novel light on neural computation.

1296

1297 **METHODS**

1298

1299 **Simulating MEG data**

1300

1301 We simulate the data so as to be akin to human MEG.

1302

1303 *Task data for obtaining state patterns*

1304

1305 We generate ground truth multivariate patterns (over sensors) of states. We then add random
1306 gaussian noise on the ground truth state patterns to form the task data. We train a logistic
1307 regression classifier on the task data so as to obtain a decoding model for each of the state
1308 patterns. Later we use this decoding model to transform the resting-state data from sensor
1309 space (with dimension of time by sensors) to the state space (with dimension of time by
1310 states).

1311

1312 *Rest data for detecting sequences*

1313

1314 First, to imitate temporal autocorrelations and spatial correlations commonly seen in human
1315 neuroimaging data, we generate the rest data using an auto-aggressive model with
1316 multivariate (over sensors) gaussian noise and add a dependence among sensors. In some
1317 simulations, we also add a rhythmic oscillation (e.g., 10Hz).

1318

1319 Second, we inject a sequence of state patterns in the rest data. The sequences follow the
1320 ground truth of state transitions of interest. The state-to-state time lag is assumed to follow a
1321 gamma distribution. We vary the number of sequences to be injected in the rest data to
1322 control the strength of sequences.

1323
1324 Lastly, we project the rest data to the decoding model of states obtained from the task data.
1325 TDLM will then work on the decoded state space.

1326
1327 An example of the Matlab implementation is called “Simulate_Replay” from the Github link:
1328 <https://github.com/yunzheliu/TDLM>

1329
1330

1331 **Human Replay dataset**

1332
1333

Task design

1334

1335 Participants were required to perform a series of tasks with concurrent MEG scanning (see
1336 details in Liu, et al. ²⁴). The functional localizer task was performed before the main task and
1337 was used to train a sensory code for eight distinct objects. Note, the participants were
1338 provided with no structural information at the time of the localizer. These decoding models,
1339 trained on the functional localizer task, capture a sensory level neural representation of
1340 stimuli (i.e., stimulus code). Following that, participants were presented with the stimuli and
1341 were required to unscramble the “visual sequence” into a correct order, i.e., the “unscrambled
1342 sequence” based on a structural template they had learned the day before. After that,
1343 participants were given a rest for 5 mins. At the end, stimuli were presented again in random
1344 order, and participants were asked to identify the true sequence identity and structural
1345 position of the stimuli. Data in this session are used to train a structural code (position and
1346 sequence) for the objects.

1347
1348

MEG data Acquisition, Pre-processing and Source Reconstruction

1349
1350

We follow the same procedure that has been reported in Liu, et al. ²⁴. We have copied it here
1351 for references.

1352
1353

“MEG was recorded continuously at 600 samples/second using a whole-head 275-channel axial
1354 gradiometer system (CTF Omega, VSM MedTech), while participants sat upright inside the scanner.
1355 Participants made responses on a button box using four fingers as they found most comfortable. The
1356 data were resampled from 600 to 100 Hz to conserve processing time and improve signal to noise
1357 ratio. All data were then high-pass filtered at 0.5 Hz using a first-order IIR filter to remove slow drift.
1358 After that, the raw MEG data were visually inspected, and excessively noisy segments and sensors
1359 were removed before independent component analysis (ICA). An ICA (FastICA,
1360 <http://research.ics.aalto.fi/ica/fastica>) was used to decompose the sensor data for each session into 150
1361 temporally independent components and associated sensor topographies. Artefact components were
1362 classified by combined inspection of the spatial topography, time course, kurtosis of the time course
1363 and frequency spectrum for all components. Eye-blink artefacts exhibited high kurtosis (>20), a
1364 repeated pattern in the time course and consistent spatial topographies. Mains interference had
1365 extremely low kurtosis and a frequency spectrum dominated by 50 Hz line noise. Artefacts were then
1366 rejected by subtracting them out of the data. All subsequent analyses were performed directly on the
1367 filtered, cleaned MEG signal, in units of femtotesla.

1368
1369

All source reconstruction was performed in SPM12 and FieldTrip. Forward models were generated on
1370 the basis of a single shell using superposition of basis functions that approximately corresponded to
1371 the plane tangential to the MEG sensor array. Linearly constrained minimum variance beamforming ⁷¹,
1372 was used to reconstruct the epoched MEG data to a grid in MNI space, sampled with a grid step of 5
1373 mm. The sensor covariance matrix for beamforming was estimated using data in either broadband
1374 power across all frequencies or restricted to ripple frequency (120-150 Hz). The baseline activity was

1375 the mean neural activity averaged over -100 ms to -50 ms relative to sequence onset. All non-
1376 artefactual trials were baseline corrected at source level. We looked at the main effect of the
1377 initialization of sequence. Non-parametric permutation tests were performed on the volume of interest
1378 to compute the multiple comparison (whole-brain corrected) P-values of clusters above 10 voxels,
1379 with the null distribution for this cluster size being computed using permutations (n = 5000
1380 permutations).”

1381
1382
1383
1384

1385 **Rodent Replay dataset**

1386
1387
1388

1387 *Data description*

1389 This data is from Ólafsdóttir, et al. ²⁶. We analyzed one full recording session data (track
1390 running for generating rate map, post-running resting for replay detection) from Rat 2192.

1391

1392 *Task description*

1393

1394 In Ólafsdóttir, et al. ²⁶, rats ran multiple laps on a Z maze, and were then placed in a rest
1395 enclosure. The two parallel sections of the Z (190 cm each) were connected by a diagonal
1396 section (220 cm). Animals were pretrained to run on the track. At the recording session, rats
1397 were placed at one end of the Z-track. The ends and corners of the track were baited with
1398 sweetened rice to encourage running from one end to the other. In each session rats
1399 completed 20 full laps (30–45 min). Following the track session, rats were placed in the rest
1400 enclosure for 1.5 hour.

1401

1402 *Preprocessing*

1403

1404 Following Ólafsdóttir, et al. ²⁶, when generating rate maps we excluded data from both the
1405 ends and corners because the animals regularly performed non-perambulatory behaviors there.
1406 Periods when running speed was less than 3 cm/s were also excluded. Running trajectories
1407 were then linearized, dwell time and spikes were binned into 2 cm bins and smoothed with a
1408 Gaussian kernel ($\sigma = 5$ bins). We generated rate maps separately for inbound (track1->track2-
1409 >track3) and outbound (track3->track2->track1) running.

1410

1411 As in Ólafsdóttir, et al. ²⁶, cells recorded in CA1 were classified as place cells if their peak
1412 firing field during track running was above 1 Hz and at least 20 cm wide. The candidate
1413 replay events were identified based on multi-unit (MU) activity from place cells during rest
1414 time. Only periods exceeding the mean rate by 3 stand deviation of MU activity were
1415 identified as putative replay events. Events less than 40 ms long or which included activity
1416 from less than 15% of the recorded place cell ensemble were rejected.

1417

1418 We analyzed data from one full recording session (track running for generating rate map,
1419 post-running resting for replay detection) of Rat 2192 reported in Ólafsdóttir, et al. ²⁶.
1420 Following the procedure described above, we have identified 58 place cells, and 1183
1421 putative replay events. Replay analysis was then performed on the putative replay events,
1422 separately for inbound and outbound rate maps.

1423

1424 **Code availability**

1425 Source code of TDLM can be found at <https://github.com/yunzheliu/TDLM>.

1426 **Data availability**

1427 Data are also available at <https://github.com/yunzheliu/TDLM>.

1428

1429 **Acknowledgement**

1430 We thank Matthew A. Wilson for help with rodent theta sequence analysis. We thank Elliott
1431 Wimmer and Toby Wise for helpful discussion and generous sharing of their data. We thank
1432 Matt Nour for helpful comments on a previous version of the manuscript. Y.L. is also
1433 grateful for the unique opportunity provided by the Brains, Minds & Machines Summer
1434 Course. We acknowledge fundings from the Open Research Fund of the State Key
1435 Laboratory of Cognitive Neuroscience and Learning to Y.L., Wellcome Trust Investigator
1436 Award (098362/Z/12/Z) to R.J.D., Wellcome Trust Senior Research Fellowship
1437 (104765/Z/14/Z), and Principal Research Fellowship (219525/Z/19/Z), together with a James
1438 S. McDonnell Foundation Award (JSMF220020372), to T.E.J.B; and Wellcome Trust Senior
1439 Research Fellowship (212281/Z/18/Z) to C.B. Both Wellcome Centres are supported by core
1440 funding from the Wellcome Trust: Wellcome Centre for Integrative Neuroimaging
1441 (203139/Z/16/Z), Wellcome Centre for Human Neuroimaging (091593/Z/10/Z). The Max
1442 Planck UCL Centre is a joint initiative supported by UCL and the Max Planck Society.

1443

1444

1445 **Appendix 1: Multi-step sequences**

1446

1447 TDLM can be used iteratively. One extension of TDLM of particular interest is: multi-step
1448 sequences. It asks about a consistent regularity among multiple states.

1449

1450 So far, we introduced methods for quantifying the extent to which the state-to-state transition
1451 structure in neural data matches a hypothesized task-related transition matrix. An important
1452 limitation of these methods is that they are blind to hysteresis in transitions. In other words,
1453 they cannot tell us about multi-step sequences. In this section, we describe a methodological
1454 extension to measure evidence for sequences comprising more than one transition: for
1455 example, $A \rightarrow B \rightarrow C$.

1456

1457 The key ingredient is controlling for shorter sub-sequences (e.g., $A \rightarrow B$ and $B \rightarrow C$), in order
1458 to find evidence unique to a multi-step sequence of interest.

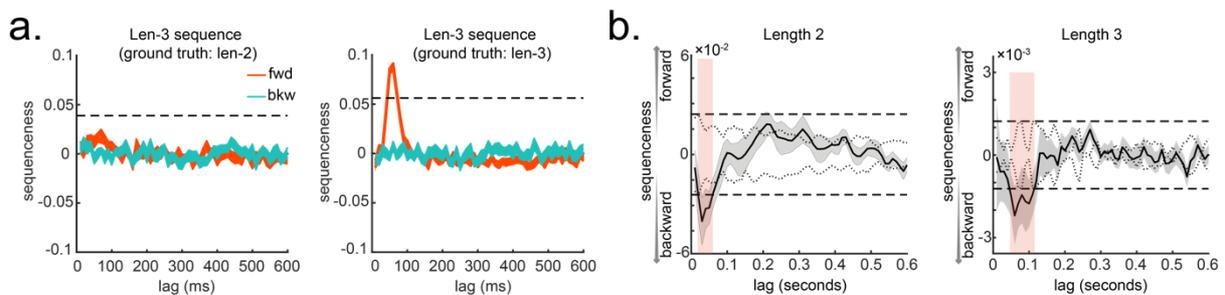
1459

1460 Assuming constant state-to-state time lag, Δt , between A and B, and between B and C. We
1461 can create a new state space AB, by shifting B up Δt , and elementwise multiply it with state
1462 A. This new state AB measures the reactivation strength of $A \rightarrow B$, with time lag Δt . In the
1463 same way, we can create a new state space, BC, AC, etc. Then we can construct the same
1464 first level GLM on the new state space. For example, if we want to determine the evidence of
1465 $A \rightarrow B \rightarrow C$ at time lag Δt , we can regress AB onto state time course C, at each Δt (cf.
1466 Equation 1). But we want to know the unique contribution of AB to C. More specifically, we
1467 want to test if the evidence of $A \rightarrow B \rightarrow C$ is stronger than $X \rightarrow B \rightarrow C$, where X is any other
1468 state but not A. Therefore, similar to Equation 2, we need to control CB, DB, when looking
1469 for evidence of AB of C. Applying this method, we show TDLM successfully avoids false

1470 positives arising out of strong evidence for shorter length (see simulation results in Appendix
 1471 1-figure 1a, see results obtained on human neuroimaging data in Appendix 1-figure 1b). This
 1472 process can be generalized to any number of steps.
 1473

1474 TDLM, in its current form, assumes a constant intra-sequence state-to-state time lag. If there
 1475 is a variability between state transitions TDLM can still cope, but not very elegantly. Assume
 1476 there is a three states sequence, $A \rightarrow B \rightarrow C$, with intra-sequence variance. TDLM will need
 1477 to test all possible combinations of state-to-state time lags in $A \rightarrow B$ and $B \rightarrow C$. If there are n
 1478 number of time lag of interest in either of the two transitions, TDLM will then have to test
 1479 n^2 possible time lag combinations. This is a large search space and one that increases
 1480 exponentially as a function of the length of a sequence.
 1481

1482 We note this analysis is different from a typical rodent replay analysis which assesses the
 1483 overall evidence for a sequence length^{16,17}. TDLM asks if there is more evidence for $A \rightarrow B \rightarrow C$,
 1484 above and beyond evidence for $B \rightarrow C$, for example. However, if the main question of
 1485 interest is “do we have evidence of $A \rightarrow B \rightarrow C$ in general”, as normally is the case in the
 1486 rodent replay analysis^{16,17}, then we should not control for shorter lengths. Instead, we can
 1487 simply average the evidence together, as implemented in Kurth-Nelson, et al.³⁵.
 1488



1489 **Appendix 1-figure 1. Extension to TDLM: Multi-step sequences.** **a**, TDLM can quantify not only
 1490 pair-wise transition, but also longer length sequences. It does so by controlling for evidence of shorter
 1491 length to avoid false positives. **b**, Method applied to human MEG data, incorporating control of both
 1492 alpha oscillation and co-activation for both length-2 and length-3 sequence length. Dashed line
 1493 indicates the permutation threshold. This is reproduced from Figure 3 A, C, Liu et al. 2019, Cell,
 1494 published under the Creative Commons Attribution 4.0 International Public License (CC BY 4.0).
 1495
 1496
 1497
 1498
 1499
 1500
 1501
 1502
 1503
 1504
 1505
 1506
 1507
 1508
 1509
 1510
 1511
 1512
 1513

1514
1515
1516
1517
1518
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538

Appendix 2: Pseudocode of sensory code and abstract code cross-validations

In the consideration of the formatting, we have attached the Latex-based algorithm box in picture form.

Algorithm 1: hold one out cross validation to compute classification accuracy. Here N is number of trials, D is number of data dimensions, and P is number of classes:

Algorithm 1: Hold one out cross validation

input : Data set $\mathcal{D} = \{X_i, Y_i\}_{i=1}^N (X_i \in \mathbb{R}^D; Y_i \in \mathbb{Z}_2^P)$
output: Cross validated classification accuracy $\{a \in \mathbb{R} : 0 \leq a \leq 1\}$
 Randomly split \mathcal{D} into $K = \frac{N}{P}$ equally sized subsets, $\mathcal{D} = \{\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_K\}$ such that each \mathcal{D}_i contains a single random sample from each class in \mathcal{Y} ;
for k *in* K **do**
 Create a training dataset $\mathcal{T}_k = \{\mathcal{D}_i : i \neq k\}$;
 Train a logistic regression classifier β_k on \mathcal{T}_k ;
 Compute classification accuracy a_k of β_k on \mathcal{D}_k ;
end
 Compute mean accuracy $a = \frac{1}{K} \sum_{k=1}^K a_k$

Algorithm 2: test a classifier’s abstraction ability across different datasets with some common structure.

Algorithm 2: Classifier Abstraction

input : Data set $\mathcal{D} = \{X_i, Y_i\}_{i=1}^N (X_i \in \mathbb{R}^D; Y_i \in \{A, B, C, D, A', B', C', D'\})$
output: Abstraction accuracy $\{a \in \mathbb{R} : 0 \leq a \leq 1\}$
 Partition \mathcal{D} into two subsets each of which exclusively contain trials from one or other structure sequence: $\mathcal{D}_1 = \{X_i, Y_i\}_{i=1}^N (X_i \in \mathbb{R}^D; Y_i \in \{A, B, C, D\})$ and $\mathcal{D}_2 = \{X_i, Y_i\}_{i=1}^N (X_i \in \mathbb{R}^D; Y_i \in \{A', B', C', D'\})$;
for k *in* $\{1, 2\}$ **do**
 Train a logistic regression classifier β_k on \mathcal{D}_k ;
 Compute classifier predictions p_k of β_k on \mathcal{D}_{3-k} ;
 Compute abstraction accuracy a_k as proportion of samples for which the prediction p_k correctly identifies the sequence location (eg A predicted for A') ;
end
 Compute mean abstraction accuracy $a = \frac{1}{2} \sum_{k=1}^2 a_k$

1539
 1540
 1541
 1542
 1543
 1544
 1545
 1546
 1547
 1548
 1549
 1550

Appendix 3: Sequences of sequences

1551
 1552
 1553
 1554
 1555
 1556

We have detailed the use of either sensory or abstract representations as the states in TDLM. We now take a step further and use sequences themselves as states. Using this kind of hierarchical analysis, we can search for sequences of sequences. This is useful because it can reveal temporal structure not only within sequence, but also between sequences. The

1557 organization between sequences is of particular interest for revealing neural computations.
1558 For example, the forward and backward search algorithms hypothesized in planning and
1559 inference⁷² can be cast as sequences of sequences problem: the temporal structure of forward
1560 and backward sequence. This can be tested by using TDLM iteratively.

1561
1562 To look for sequences between sequences we need first to define sequences as new states. To
1563 do so, the raw state course, for example, state B needs to be shifted up by the empirical
1564 within-sequence time lag Δt (determined by the two-level GLM described above), to align
1565 with the onset of state A, if assuming sequence $A \rightarrow B$ exist (at time lag Δt). Then, we can
1566 elementwise multiply the raw state time course A with the shifted time course B, resulting in
1567 a new state AB. Each entry in this new state time course indicates the reactivation strength of
1568 sequence AB at a given time.

1569
1570 The general two-level GLMs framework still applies, but now with one important caveat. The
1571 new sequence state (e.g., AB) is defined based on the original states (A and B), and where we
1572 are now interested in a reactivation regularity, i.e., sequence, between sequences, rather than
1573 the original states. We need therefore to control for the effects of the original states.
1574 Effectively, this is like controlling for main effects (e.g., state A and shifted state B) when
1575 looking for their interaction (sequence AB). TDLM achieves this by including time lagged
1576 original state regressors A, B, in addition to AB, in the first level GLM sequence analysis.

1577
1578 Specifically, let's assume the sequence state matrix is X_{seq} , after transforming the original
1579 state space to sequence space based on the empirical within-sequence time lag Δt_w . Each
1580 column at X_{seq} is sequence state, denoted by S_{ij} , which indicates the strength of sequence $i \rightarrow$
1581 j reactivation. The raw state i is X_i , and the shifted raw state j is X_{jw} (by time lag Δt_w).

1582
1583 In the first level GLM, TDLM ask for the strength of a unique contribution of sequence state
1584 S_{ij} to S_{mn} while controlling for original states (X_i and X_{jw}). For each sequence state ij , at
1585 each possible time lag Δt , TDLM estimated a separate linear model:

$$1586 \quad S_{mn} = X_i(\Delta t)\beta_i + X_{jw}(\Delta t)\beta_j + S_{ij}(\Delta t)\beta_{ij}(\Delta t)$$

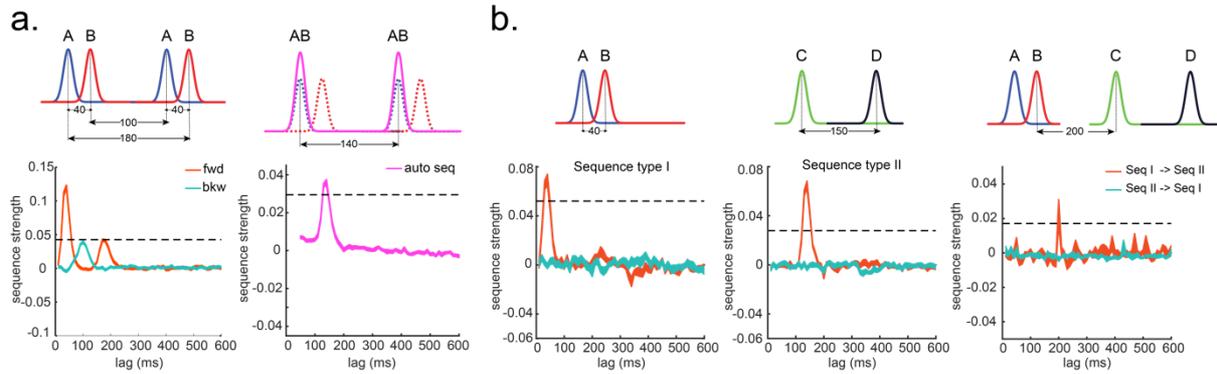
1587
1588 (13)

1589
1590 Repeat this process for each sequence state separately at each time lag, resulting a sequence
1591 matrix β_{seq} .

1592
1593 At the 2nd level GLM, TDLM asks how strong the evidence for a sequence of interest is,
1594 compared to sequences that have the same starting state or end state at each time lag. This 2nd
1595 level GLM will be the same as the equation 5, but with additional regressors to control for
1596 sequences that share the same start or end state. In simulation we demonstrate, applying this
1597 method, that TDLM can uncover hierarchical temporal structure: state A is temporally
1598 leading state B with 40 ms lag, and the sequence A->B tends to repeat itself with a 140 ms
1599 gap (Appendix 3-figure 1a). One interesting application of this is to look for theta sequence
1600⁷³⁻⁷⁵. One can think of theta sequence, a well-documented phenomenon during rodent spatial
1601 navigation, as a neural sequence repeating itself in theta frequency (6 - 12 Hz).

1602
1603 In addition to looking for temporal structure of the same sequence, the method is equally
1604 suitable when searching for temporal relationships between different sequences in a general
1605 form. For example, assuming two different types of sequences, one sequence type has a

1606 within-sequence time lag at 40 ms; while the other has a within-sequence time lag at 150 ms
 1607 (Appendix 3-figure 1b, left and middle panel); and there is a gap of 200 ms between the two
 1608 types of sequences (Appendix 3-figure 1b, right panel). These time lags are set arbitrarily for
 1609 illustration purposes. TDLM can accurately capture the dynamics both within and between
 1610 the sequences, supporting a potential for uncovering temporal relationships between
 1611 sequences in general under the same framework.
 1612
 1613



1614
 1615
 1616 **Appendix 3-figure 1. Sequences of sequences.** **a**, TDLM can also be used iteratively to capture the
 1617 repeating pattern of a sequence event itself. Illustration in the top panel describes the ground truth in
 1618 the simulation. Intra-sequence temporal structure (right) and inter-sequence temporal structure (right)
 1619 can be extracted simultaneously. **b**, Temporal structure between and within different sequences.
 1620 Illustration of two sequence types with different state-to-state time lag within sequence, and a
 1621 systematic gap between the two types of sequences on top. TDLM can capture the temporal structures
 1622 both within (left and middle panel) and between (right panel) the two sequence types.
 1623

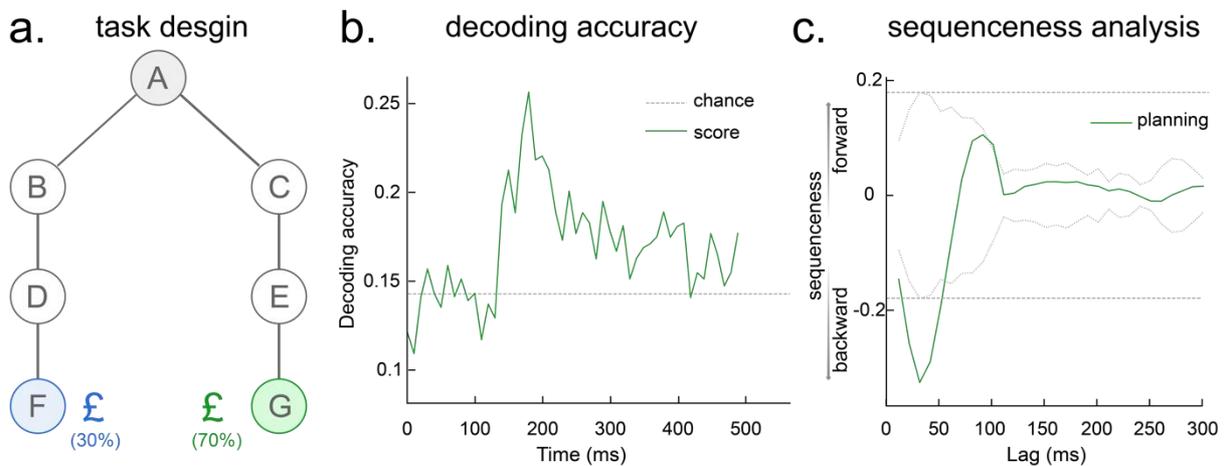
1624
 1625
 1626
 1627
 1628
 1629
 1630
 1631
 1632
 1633
 1634
 1635
 1636
 1637
 1638
 1639 **Appendix 4: Apply TDLM to human whole-brain EEG data**
 1640

1641 An autocorrelation is commonplace in neuroimaging data, including EEG and fMRI. TDLM
 1642 is designed to specifically take care of this confound and, on this basis, we should be able to
 1643 work with EEG and fMRI data. We do not have the suitable fMRI data available to test

1644 TDLM but are interested to investigate this in more depth in our future work. We had
1645 collected EEG data from one participant to test whether TDLM would *just* work.

1646
1647 The task was designed to examine online sequential replay in online decision-making, by Dr.
1648 Toby Wise. This is a ‘T-maze’ like task, where a participant needs to choose a left or right
1649 path based on the value received at the end of the path. We could decode 7 objects well on
1650 the whole-brain EEG data using just raw amplitude (same with our MEG-based analysis),
1651 and could detect fast backward sequenceness (peaked at 30 ms time lag) during
1652 choice/planning time (Appendix 4-figure 1), similar to our previous MEG findings³⁵. As this
1653 result is from one subject, we are cautious about making an excessive claim, but nevertheless
1654 we believe the data show the approach is highly promising for EEG data.

1655
1656



1657
1658

1659 **Appendix 4-figure 1. Sequence detection in EEG data (from one participant).** **a**, Task design.
1660 At each trial, the participant starts at state A, and he/she needs to select either “BDF” or “CEG” path,
1661 based on the final reward receipt at terminal states F and G. All seven states are indexed by pictures. **b**,
1662 The leave-one-out cross-validated decoding accuracy is shown, with a peak at around 200 ms after
1663 stimulus onset, similar to our previous MEG findings. **c**, TDLM method is then applied on the
1664 decoded state time course where we find a fast backward sequenceness that conforms to task structure.
1665 Shown here is a subtraction between forward and backward sequenceness, where a negative
1666 sequenceness indicates stronger backward sequence replay. The dotted line is the peak of the absolute
1667 state-permutation at each time lag, the dotted line the max over all computed state time lags, thereby
1668 controlling for multiple comparisons. This is the same statistical method used in our previous
1669 empirical work, and in the current paper. These EEG sequence results replicate our previous MEG-
1670 based findings based on analyses at planning/decision time (see Figure 3 in Kurth-Nelson et al., 2016,
1671 and also see Figure 3f in Liu et al., 2019).

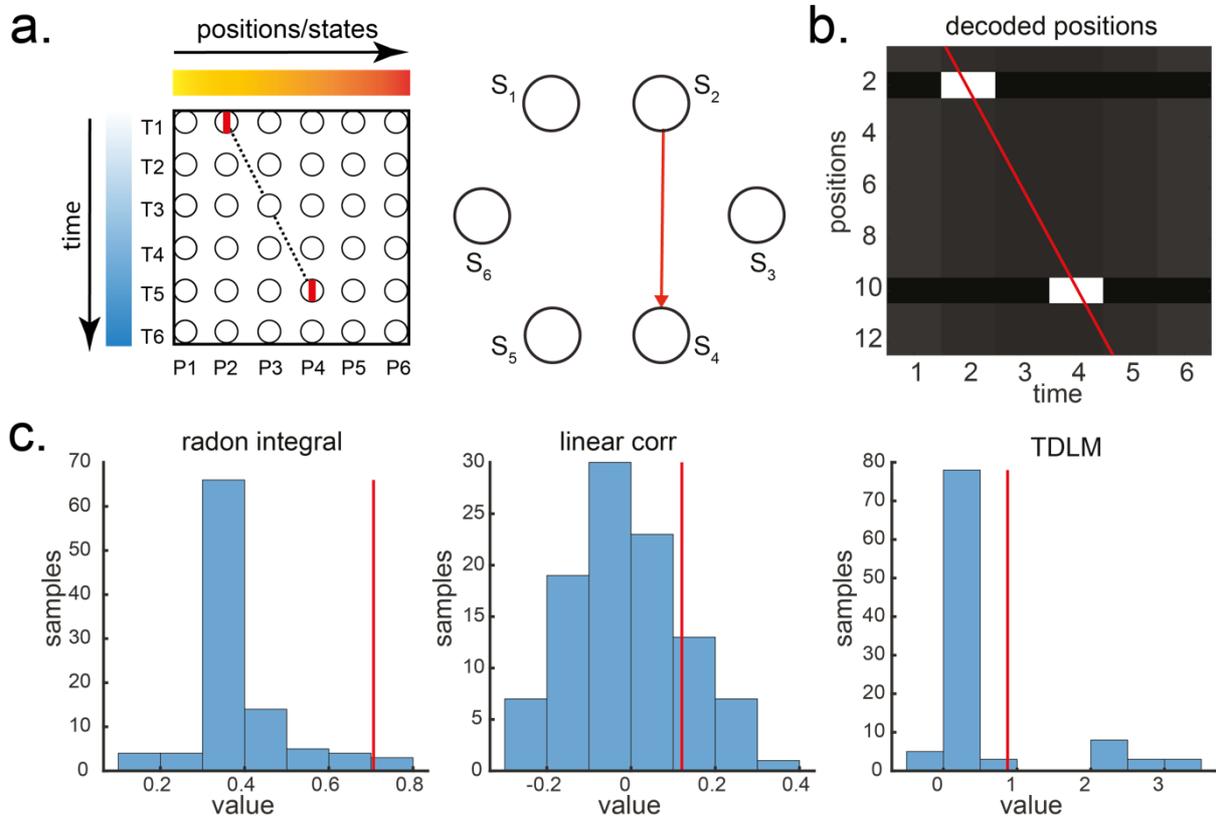
1672
1673
1674
1675
1676
1677
1678

1679 Appendix 5: Less sensitivity of TDLM to skipping sequences

1680

1681 In a linear track where replays only go in a single direction, it is possible that TDLM is less
1682 sensitive compared to the linear correlation or the Radon method, given the latter assumes a
1683 parametric relationship between space and time. For example, if only the first and last state

1684 are activated, but not the intermediate states, the existing methods **will** report replay, but
 1685 TDLM **will not**, because in existing methods **space and time are parametric quantities**
 1686 (Appendix 5-figure 1). In contrast, TDLM only knows about transitions on a graph.
 1687
 1688



1689
 1690

1691 **Appendix 5-figure 1. Parametric relationship between space and time vs. graph transitions.** **a,**
 1692 A scheme for the decoded time by position space is shown (left). Both Radon and linear weighted
 1693 correlation methods aim to capture a parametric relationship between space and time. TDLM, on the
 1694 other hand, tries to capture transitions in a graph (shown in right, with the red indicating the transition
 1695 of interest). **b,** A decoded time by position matrix from simulated spiking data. **c,** Replay analysis
 1696 using all three methods on this decoded position matrix. TDLM is less sensitive compared to existing
 1697 “line search” methods, like radon or linear correlation. The red line indicates the true sequence
 1698 measure from each of these methods. The bar plots are permutation samples by randomly shuffling
 1699 the rate maps.

1700
 1701
 1702
 1703
 1704
 1705
 1706
 1707
 1708

1709 **REFERENCES**

1710
 1711 1 Haxby, J. V., Connolly, A. C. & Guntupalli, J. S. Decoding neural representational
 1712 spaces using multivariate pattern analysis. *Annual Review of Neuroscience* **37**, 435-
 1713 456 (2014).

1714 2 Kriegeskorte, N., Mur, M. & Bandettini, P. A. Representational similarity analysis-
1715 connecting the branches of systems neuroscience. *Frontiers in Systems Neuroscience*
1716 **2**, 4 (2008).

1717 3 Barron, H. C., Garvert, M. M. & Behrens, T. E. Repetition suppression: a means to
1718 index neural representations using BOLD? *Philosophical Transactions of the Royal*
1719 *Society B: Biological Sciences* **371**, 20150355 (2016).

1720 4 Smith, S. M. *et al.* Correspondence of the brain's functional architecture during
1721 activation and rest. *Proceedings of the National Academy of Sciences* **106**, 13040-
1722 13045 (2009).

1723 5 Tavor, I. *et al.* Task-free MRI predicts individual differences in brain activity during
1724 task performance. *Science* **352**, 216-220 (2016).

1725 6 Higgins, C. *et al.* Replay bursts in humans coincide with activation of the default
1726 mode and parietal alpha networks. *Neuron* **109**, 882-893. e887 (2021).

1727 7 Wilson, M. A. & McNaughton, B. L. Reactivation of hippocampal ensemble
1728 memories during sleep. *Science* **265**, 676-679 (1994).

1729 8 Skaggs, W. E. & McNaughton, B. L. Replay of neuronal firing sequences in rat
1730 hippocampus during sleep following spatial experience. *Science* **271**, 1870-1873
1731 (1996).

1732 9 Louie, K. & Wilson, M. A. Temporally structured replay of awake hippocampal
1733 ensemble activity during rapid eye movement sleep. *Neuron* **29**, 145-156 (2001).

1734 10 Lee, A. K. & Wilson, M. A. Memory of sequential experience in the hippocampus
1735 during slow wave sleep. *Neuron* **36**, 1183-1194 (2002).

1736 11 Foster, D. J. Replay comes of age. *Annual Review of Neuroscience* **40**, 581-602
1737 (2017).

1738 12 Ólafsdóttir, H. F., Bush, D. & Barry, C. The Role of Hippocampal Replay in Memory
1739 and Planning. *Current Biology* **28**, R37-R50 (2018).

1740 13 Pfeiffer, B. E. The content of hippocampal “replay”. *Hippocampus* **30**, 6-18 (2020).

1741 14 Carr, M. F., Jadhav, S. P. & Frank, L. M. Hippocampal replay in the awake state: a
1742 potential substrate for memory consolidation and retrieval. *Nature Neuroscience* **14**,
1743 147 (2011).

1744 15 Lisman, J. *et al.* Viewpoints: how the hippocampus contributes to memory, navigation
1745 and cognition. *Nature Neuroscience* **20**, 1434-1447 (2017).

1746 16 Davidson, T. J., Kloosterman, F. & Wilson, M. A. Hippocampal replay of extended
1747 experience. *Neuron* **63**, 497-507 (2009).

1748 17 Grosmark, A. D. & Buzsáki, G. Diversity in neural firing dynamics supports both
1749 rigid and learned hippocampal sequences. *Science* **351**, 1440-1443 (2016).

1750 18 Maboudi, K. *et al.* Uncovering temporal structure in hippocampal output patterns.
1751 *eLife* **7**, e34467 (2018).

1752 19 van der Meer, M. A., Kemere, C. & Diba, K. Progress and issues in second-order
1753 analysis of hippocampal replay. *Philosophical Transactions of the Royal Society B:*
1754 *Biological Sciences* **375**, 20190238 (2020).

1755 20 Tingley, D. & Peyrache, A. On the methods for reactivation and replay analysis.
1756 *Philosophical Transactions of the Royal Society B: Biological Sciences* **375**,
1757 20190231 (2020).

1758 21 Rosenberg, M., Zhang, T., Perona, P. & Meister, M. Mice in a labyrinth: Rapid
1759 learning, sudden insight, and efficient exploration. *bioRxiv* (2021).

1760 22 Wimmer, G. E., Liu, Y., Vehar, N., Behrens, T. E. J. & Dolan, R. J. Episodic memory
1761 retrieval success is associated with rapid replay of episode content. *Nature*
1762 *Neuroscience* **23**, 1025–1033 (2020).

- 1763 23 Nour, M. M., Liu, Y., Arumuham, A., Kurth-Nelson, Z. & Dolan, R. Impaired neural
1764 replay of inferred relationships in schizophrenia. *Cell in press* (2021).
- 1765 24 Liu, Y., Dolan, R. J., Kurth-Nelson, Z. & Behrens, T. E. J. Human replay
1766 spontaneously reorganizes experience. *Cell* **178**, 640-652 (2019).
- 1767 25 Liu, Y., Mattar, M. G., Behrens, T. E., Daw, N. D. & Dolan, R. J. Experience replay
1768 is associated with efficient nonlocal learning. *Science* **372** (2021).
- 1769 26 Ólafsdóttir, H. F., Carpenter, F. & Barry, C. Coordinated grid and place cell replay
1770 during rest. *Nature Neuroscience* **19**, 792 (2016).
- 1771 27 Barron, H. C., Mars, R. B., Dupret, D., Lerch, J. P. & Sampaio-Baptista, C. Cross-
1772 species neuroscience: closing the explanatory gap. *Philosophical Transactions of the*
1773 *Royal Society B: Biological Sciences* **376**, 20190633 (2021).
- 1774 28 Kurth-Nelson, Z., Barnes, G., Sejdinovic, D., Dolan, R. & Dayan, P. Temporal
1775 structure in associative retrieval. *eLife* **4**, e04919 (2015).
- 1776 29 Fyhn, M., Hafting, T., Treves, A., Moser, M.-B. & Moser, E. I. Hippocampal
1777 remapping and grid realignment in entorhinal cortex. *Nature* **446**, 190 (2007).
- 1778 30 Dehaene, S., Meyniel, F., Wacongne, C., Wang, L. & Pallier, C. The neural
1779 representation of sequences: from transition probabilities to algebraic patterns and
1780 linguistic trees. *Neuron* **88**, 2-19 (2015).
- 1781 31 Colclough, G. L., Brookes, M. J., Smith, S. M. & Woolrich, M. W. A symmetric
1782 multivariate leakage correction for MEG connectomes. *Neuroimage* **117**, 439-448
1783 (2015).
- 1784 32 Deodatis, G. & Shinozuka, M. Auto-regressive model for nonstationary stochastic
1785 processes. *Journal of engineering mechanics* **114**, 1995-2012 (1988).
- 1786 33 Eichler, M. Granger causality and path diagrams for multivariate time series. *Journal*
1787 *of Econometrics* **137**, 334-353 (2007).
- 1788 34 Eldar, E., Bae, G. J., Kurth-Nelson, Z., Dayan, P. & Dolan, R. J.
1789 Magnetoencephalography decoding reveals structural differences within integrative
1790 decision processes. *Nature Human Behaviour* **2**, 670-681 (2018).
- 1791 35 Kurth-Nelson, Z., Economides, M., Dolan, Raymond J. & Dayan, P. Fast Sequences
1792 of Non-spatial State Representations in Humans. *Neuron* **91**, 194-204 (2016).
- 1793 36 Lubenov, E. V. & Siapas, A. G. Hippocampal theta oscillations are travelling waves.
1794 *Nature* **459**, 534-539 (2009).
- 1795 37 Wilson, H. R., Blake, R. & Lee, S.-H. Dynamics of travelling waves in visual
1796 perception. *Nature* **412**, 907-910 (2001).
- 1797 38 Weinberger, K. Q., Blitzer, J. & Saul, L. K. in *Advances in neural information*
1798 *processing systems*. 1473-1480.
- 1799 39 Higgins, C. *Uncovering temporal structure in neural data with statistical machine*
1800 *learning models*, University of Oxford, (2019).
- 1801 40 Vidaurre, D., Smith, S. M. & Woolrich, M. W. Brain network dynamics are
1802 hierarchically organized in time. *Proceedings of the National Academy of Sciences*
1803 **114**, 12827-12832 (2017).
- 1804 41 Harris, K. D. Nonsense correlations in neuroscience. *bioRxiv* (2020).
- 1805 42 Worsley, K. J. *et al.* A unified statistical approach for determining significant signals
1806 in images of cerebral activation. *Human brain mapping* **4**, 58-73 (1996).
- 1807 43 Nichols, T. E. Multiple testing corrections, nonparametric methods, and random field
1808 theory. *Neuroimage* **62**, 811-815 (2012).
- 1809 44 Messinger, A., Squire, L. R., Zola, S. M. & Albright, T. D. Neuronal representations
1810 of stimulus associations develop in the temporal lobe during learning. *Proceedings of*
1811 *the National Academy of Sciences* **98**, 12239-12244 (2001).

1812 45 Sakai, K. & Miyashita, Y. Neural organization for the long-term memory of paired
1813 associates. *Nature* **354**, 152-155 (1991).

1814 46 Barron, H. C., Dolan, R. J. & Behrens, T. E. Online evaluation of novel choices by
1815 simultaneous representation of multiple memories. *Nature Neuroscience* **16**, 1492
1816 (2013).

1817 47 Wimmer, G. E. & Shohamy, D. Preference by association: how memory mechanisms
1818 in the hippocampus bias decisions. *Science* **338**, 270-273 (2012).

1819 48 Schapiro, A. C., Rogers, T. T., Cordova, N. I., Turk-Browne, N. B. & Botvinick, M.
1820 M. Neural representations of events arise from temporal community structure. *Nature*
1821 *Neuroscience* **16**, 486 (2013).

1822 49 Garvert, M. M., Dolan, R. J. & Behrens, T. E. A map of abstract relational knowledge
1823 in the human hippocampal–entorhinal cortex. *eLife* **6**, e17086 (2017).

1824 50 Sirota, A. *et al.* Entrainment of neocortical neurons and gamma oscillations by the
1825 hippocampal theta rhythm. *Neuron* **60**, 683-697 (2008).

1826 51 Buzsáki, G. & Vanderwolf, C. H. Cellular bases of hippocampal EEG in the behaving
1827 rat. *Brain Research Reviews* **6**, 139-171 (1983).

1828 52 Hincapié, A.-S. *et al.* The impact of MEG source reconstruction method on source-
1829 space connectivity estimation: a comparison between minimum-norm solution and
1830 beamforming. *Neuroimage* **156**, 29-42 (2017).

1831 53 O'Neill, G. C. *et al.* Testing covariance models for MEG source reconstruction of
1832 hippocampal activity. *bioRxiv* (2021).

1833 54 Zhang, K., Ginzburg, I., McNaughton, B. L. & Sejnowski, T. J. Interpreting neuronal
1834 population activity by reconstruction: unified framework with application to
1835 hippocampal place cells. *Journal of Neurophysiology* **79**, 1017-1044 (1998).

1836 55 Toft, P. A. The Radon transform-theory and implementation. (1996).

1837 56 Denovellis, E. L. *et al.* Hippocampal replay of experience at real-world speeds.
1838 *bioRxiv* (2020).

1839 57 Karlsson, M. P. & Frank, L. M. Awake replay of remote experiences in the
1840 hippocampus. *Nature Neuroscience* **12**, 913 (2009).

1841 58 Nádasdy, Z., Hirase, H., Czurkó, A., Csicsvari, J. & Buzsáki, G. Replay and time
1842 compression of recurring spike sequences in the hippocampus. *Journal of*
1843 *Neuroscience* **19**, 9497-9507 (1999).

1844 59 Kriegeskorte, N., Simmons, W. K., Bellgowan, P. S. & Baker, C. I. Circular analysis
1845 in systems neuroscience: the dangers of double dipping. *Nature Neuroscience* **12**, 535
1846 (2009).

1847 60 Diba, K. & Buzsáki, G. Forward and reverse hippocampal place-cell sequences during
1848 ripples. *Nature Neuroscience* **10**, 1241 (2007).

1849 61 Baker, A. P. *et al.* Fast transient networks in spontaneous human brain activity. *eLife*
1850 **3**, e01867 (2014).

1851 62 Raichle, M. E. *et al.* A default mode of brain function. *Proceedings of the National*
1852 *Academy of Sciences* **98**, 676-682 (2001).

1853 63 Tambini, A. & Davachi, L. Awake Reactivation of Prior Experiences Consolidates
1854 Memories and Biases Cognition. *Trends in Cognitive Sciences* (2019).

1855 64 Norman, K. A., Polyn, S. M., Detre, G. J. & Haxby, J. V. Beyond mind-reading:
1856 multi-voxel pattern analysis of fMRI data. *Trends in Cognitive Sciences* **10**, 424-430
1857 (2006).

1858 65 Lewis, P. A. & Durrant, S. J. Overlapping memory replay during sleep builds
1859 cognitive schemata. *Trends in Cognitive Sciences* **15**, 343-351 (2011).

1860 66 Schuck, Nicolas W., Cai, Ming B., Wilson, Robert C. & Niv, Y. Human Orbitofrontal
1861 Cortex Represents a Cognitive Map of State Space. *Neuron* **91**, 1402-1412 (2016).

1862 67 Eldar, E., Lièvre, G., Dayan, P. & Dolan, R. J. The roles of online and offline replay
1863 in planning. *eLife* **9**, e56911 (2020).
1864 68 Dayan, P. & Daw, N. D. Decision theory, reinforcement learning, and the brain.
1865 *Cognitive, Affective, & Behavioral Neuroscience* **8**, 429-453 (2008).
1866 69 Schuck, N. W. & Niv, Y. Sequential replay of nonspatial task states in the human
1867 hippocampus. *Science* **364**, eaaw5181 (2019).
1868 70 Wittkuhn, L. & Schuck, N. W. Dynamics of fMRI patterns reflect sub-second
1869 activation sequences and reveal replay in human visual cortex. *Nature*
1870 *Communications* **12**, 1-22 (2021).
1871 71 Van Veen, B. D., Van Drongelen, W., Yuchtman, M. & Suzuki, A. Localization of
1872 brain electrical activity via linearly constrained minimum variance spatial filtering.
1873 *IEEE Transactions on biomedical engineering* **44**, 867-880 (1997).
1874 72 Penny, W. D., Zeidman, P. & Burgess, N. Forward and backward inference in spatial
1875 cognition. *PLoS computational biology* **9** (2013).
1876 73 Mehta, M., Lee, A. & Wilson, M. Role of experience and oscillations in transforming
1877 a rate code into a temporal code. *Nature* **417**, 741 (2002).
1878 74 McNaughton, B. L., Battaglia, F. P., Jensen, O., Moser, E. I. & Moser, M.-B. Path
1879 integration and the neural basis of the 'cognitive map'. *Nature Reviews Neuroscience* **7**,
1880 663 (2006).
1881 75 Buzsáki, G. & Moser, E. I. Memory, navigation and theta rhythm in the hippocampal-
1882 entorhinal system. *Nature Neuroscience* **16**, 130 (2013).
1883
1884