

Learning to grasp and extract affordances: the Integrated Learning of Grasps and Affordances (ILGA) model

James Bonaiuto^{1,2,3} · Michael A. Arbib^{2,3,4}

Received: 17 June 2015 / Accepted: 29 October 2015 / Published online: 19 November 2015
© The Author(s) 2015. This article is published with open access at Springerlink.com

Abstract The activity of certain parietal neurons has been interpreted as encoding affordances (directly perceivable opportunities) for grasping. Separate computational models have been developed for infant grasp learning and affordance learning, but no single model has yet combined these processes in a neurobiologically plausible way. We present the Integrated Learning of Grasps and Affordances (ILGA) model that simultaneously learns grasp affordances from visual object features and motor parameters for planning grasps using trial-and-error reinforcement learning. As in the Infant Learning to Grasp Model, we model a stage of infant development prior to the onset of sophisticated visual processing of hand–object relations, but we assume that certain premotor neurons activate neural populations in primary motor cortex that synergistically control different combinations of fingers. The ILGA model is able to extract affordance representations from visual object features, learn motor parameters for generating stable grasps, and generalize its learned representations to novel objects.

Keywords Neural network model · Grasping · Infant development · Affordances

1 Introduction

The notion of affordances as directly perceivable opportunities for action (Gibson 1966) was used to interpret the activity of certain parietal neurons as encoding affordances for grasping in the FARS model of parieto-frontal interactions in grasping (Fagg and Arbib 1998). However, the FARS model “hard-wires” these affordances, whereas our concern is with the development of these affordances and the grasps they afford. While computational models of infant grasp learning (Oztop et al. 2004) and affordance learning (Oztop et al. 2006) have been developed that work in a staged fashion, there do not exist any models that learn affordance extraction and grasp motor programs simultaneously. This model follows from a suggestion of Arbib et al. (2009) and implements a dual learning system that simultaneously learns both grasp affordances and motor parameters for planning grasps using trial-and-error reinforcement learning. As in the Infant Learning to Grasp Model (ILGM, Oztop et al. 2004), we model a stage of infant development prior to the onset of sophisticated visual processing of hand–object relations, but as in the FARS model (Fagg and Arbib 1998), we assume that certain premotor neurons activate neural populations in primary motor cortex that synergistically controls different combinations of fingers. The issue is to understand how different visual patterns can activate the appropriate subset of these neurons. Specifically, the task of ILGA is to learn (i) “affordances,” representations of object features that indicate where it can be grasped, and (ii) motor parameters that can be used to successfully grasp objects based on these representations.

Electronic supplementary material The online version of this article (doi:10.1007/s00422-015-0666-2) contains supplementary material, which is available to authorized users.

✉ James Bonaiuto
bonaiuto@usc.edu

- ¹ Sobell Department of Motor Neuroscience and Movement Disorders, University College London, London WC1N3BG, UK
- ² Neuroscience Program, University of Southern California, Los Angeles, CA 90089-2520, USA
- ³ USC Brain Project, University of Southern California, Los Angeles, CA 90089-2520, USA
- ⁴ Computer Science Department, University of Southern California, Los Angeles, CA 90089-2520, USA

Newborn infants aim their arm movements toward fixated objects (von Hofsten 1982). These early arm movements have been related to the development of object-directed reaching (Bhat et al. 2005), leading to grasping (Bhat and Galloway 2006), the development of which continues throughout childhood (Kuhntz-Buschbeck et al. 1998). Previous relevant models of infant motor development include Berthier's (1996), Berthier et al. (2005) and Caligiore et al.'s (2014) models of learning to reach and the ILGM. The thread shared by these models is reinforcement-based learning of intrinsically motivated goal-directed actions based on exploratory movements, or motor babbling: Movements are generated erratically in response to a target and the mechanisms generating the movements are modified via positive reinforcement (Cangelosi and Schlesinger 2013).

Grasping in development seems to increasingly involve visual information in preprogramming the grasp (Lasky 1977; Lockman et al. 1984; Von Hofsten and Ronnqvist 1988; Clifton et al. 1993; Newell et al. 1993; Witherington 2005). In ILGM, the affordance extraction module only represented the presence, position, or orientation of an object. All fingers were extended in the initial “preshape” portion of the grasp to a maximal aperture. Initially, the enclosure was triggered by the palmar reflex upon object contact. However, each time the result of reflex grasping provided a stable grasp, this grasp was reinforced, and over time, a repertoire developed of situations in which a stable grasp could be elicited—including the appropriate repositioning of the hand to a position from which the final approach of hand to object could be made—without relying on the happenstance of reflex grasping. However, the model did not learn the affordances of objects and a fortiori could not exploit them by preshaping appropriately during the reach to grasp. This shortcoming motivates the Integrated Learning of Grasps and Affordances (ILGA) model which simulates the way in which affordance extraction and grasp specification may be adapted simultaneously. It models the developmental transition to hand preshape based on visual information (Von Hofsten and Ronnqvist 1988; Schettino et al. 2003; Witherington 2005) and utilizes the “virtual finger hypothesis” for hand control during grasping. The virtual finger hypothesis states that grasping involves the assignment of real fingers to the so-called virtual fingers (VFs) or force applicators (Arbib et al. 1985). For example in a power grasp, one virtual finger might be the thumb and the other might be the palm while a precision pinch might oppose the thumb to a virtual finger comprising one or more of the other fingers. The task of grasping is then to preshape the hand according to the selected virtual fingers and the size of the object and bring the opposition axis of the virtual fingers into alignment with the selected object surface opposition axis (grasp affordance). Experimental evidence consistent with this hypothesis, also known as hierarchical control of prehension synergies, has

been found (Smeets and Brenner 2001; Zatsiorsky and Latash 2004; Wings and Santello 2005). These studies suggest that grasp force planning occurs in two hierarchical levels, with virtual finger force planning occurring first and then planning at the level of individual fingers, which provides a common input to motor neurons in various finger muscles (Zatsiorsky and Latash 2004; Wings and Santello 2005). However, note that it has also been suggested that prehension synergies are not well represented in motor cortex and that the role of motor cortex is to modulate subcortically represented synergies in order to allow individual control of the fingers (Mollazadeh et al. 2014).

In the FARS model, ventral premotor region F5 contained populations of neurons *prewired* for each grasp type (precision pinch, power grasp, etc.). Within each population, subpopulations were selective for each phase of the grasp. In the behavioral protocol used in the FARS simulations, this included Set, Extension, Flexion, Hold, and Release; however, only Set, Extension, and Flexion are considered here. Neurons in each Set subpopulation excited neurons in the same subpopulation and inhibited those neurons in Set subpopulations selective for other grasps. This connectivity implemented a winner-take-all dynamic that selected which grasp to perform based on affordance input from the anterior intraparietal area AIP. The grasp was controlled by activating F5 populations that encoded each phase of the grasp (although some neurons might be active across more than one consecutive phase). Feedback projections from F5 neurons to AIP modulated its activity according to the grasp phase.

In FARS, the secondary somatosensory cortex (SII) detected when the hand aperture reached the predicted maximal value for the planned grasp and triggered grasp enclosure by activating the appropriate Flexion subpopulation in F5. ILGM included detection of object contact with the palm in SI that automatically triggered enclosure. These represent feedforward and naïve feedback strategies, respectively. However, a more sophisticated feedback strategy might be used where the detection of the hand approaching the object is used to trigger the grasp enclosure. This is still a feedback strategy in that the relative position of hand and target object is used to trigger the enclosure, but does so before the hand contacts the object as in the naïve feedback strategy.

Before delving further into the anatomy, it is important to note that we are conflating neurophysiological and neuroanatomical data from macaques with behavioral and lesion data from humans. Thus, the areas that ILGA simulates are primarily rooted in macaque neuroanatomy. However, it is widely accepted that the basic mechanisms of the reach and grasp are conserved across monkeys, apes, and humans, and thus the ILGA model, like the FARS model, is assumed to apply equally well to the macaque and (with suitable alignment of homologous brain regions) for apes and humans. The FARS model made an important distinction between the

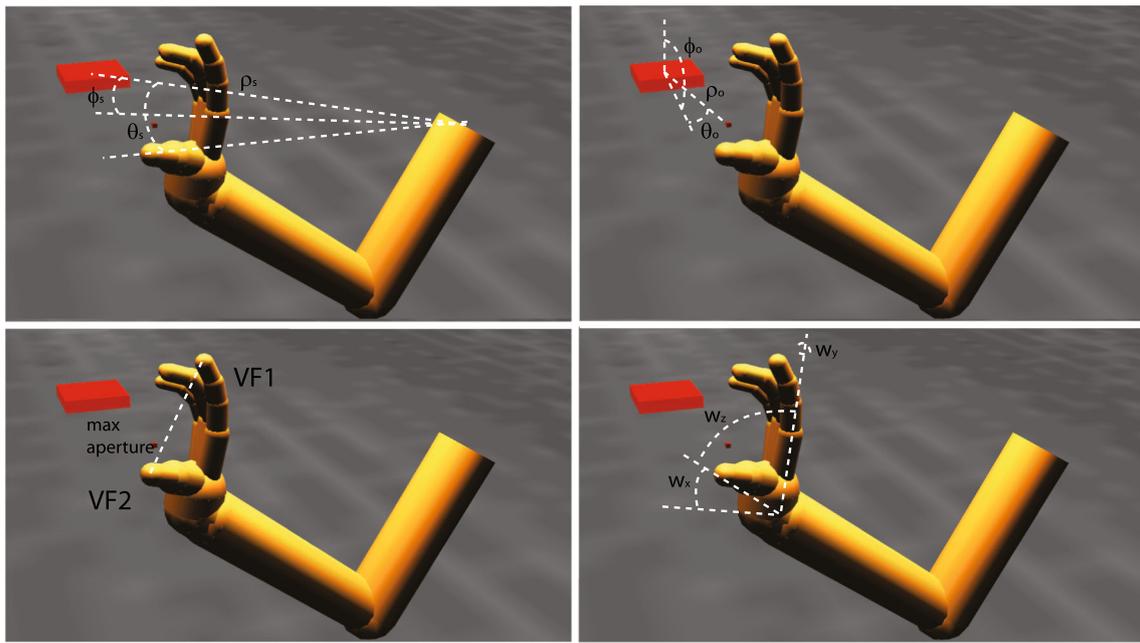


Fig. 1 Reach and grasp parameters encoded by the premotor cortex. The *small cube* denotes the planned reach offset point, and the larger rectangular prism is the object to be grasped. As explained later, area F2 encodes the shoulder-centered object position, $\varphi_s, \theta_s, \rho_s$, F7 encodes the

object-centered reach offset, $\varphi_o, \theta_o, \rho_o$, F5 encodes the VF combination and maximum aperture used for the grasp, and F2/F5 encodes the wrist orientation w_x, w_y, w_z

processing of visual input via the dorsal and ventral streams. The dorsal stream was modeled as extracting the affordances of the attended object, whereas the ventral stream could extract the identity of the object and could thus supply data to prefrontal cortex (PFC) to use in determining which of the affordances would best meet current task demands and working memory. The division of labor can be exemplified by having the ventral stream recognize a mug, recall that it still contains coffee, deciding to drink it, and thus selecting the handle for grasping; but it is up to the dorsal stream to process shape information to extract the affordance of the handle and the motor parameters of the grasp appropriate to that affordance.

Since the distinction was first made between the roles of dorsal and ventral visual streams in grasping (Goodale and Milner 1992; Jeannerod et al. 1994), the dorsal stream has been further subdivided into the dorsal-medial and dorsal-ventral streams (Rizzolatti and Matelli 2003). It has been suggested that the dorsal-medial stream, involving superior parietal and intraparietal regions and the dorsal premotor cortex, controls reaching, while the dorsal-ventral stream, including inferior parietal and intraparietal regions and the ventral premotor cortex, controls grasping (Jeannerod et al. 1995; Wise et al. 1997). The main regions of the dorsal-medial stream seem to include the medial intraparietal area (MIP) and area V6A in the parietal cortex, and area F2 in the dorsal premotor cortex. The spatial dimensions of potential targets such as direction and distance are likely processed independently in

parallel (Battaglia-Mayer et al. 2003). In support of this idea, direction and distance reach errors dissociate (Soechting and Flanders 1989; Gordon et al. 1994) and distance information decays faster than direction information in working memory (McIntyre et al. 1998). Our model therefore dissociates the representation of direction and distance in order to reduce the dimensionality of spatial representations to be associated with motor parameter values. However, we note data suggesting that many neurons are modulated by a combination of both variables (Fu et al. 1993; Messier and Kalaska 2000).

ILGA learns to perform successful grasps by developing mappings between visual object features and various parameters used to control the arm and hand (Fig. 1). The grasp is determined by the selection of virtual fingers, VF1 and VF2, as well as the maximum aperture. While studies of parietal representation of object location test the encoding of the center of a visual target, reach-to-grasp movements direct the wrist to some region offset from this point so that the hand may contact the object’s affordances appropriately. Similar to the ILGM, ILGA learns to select an object-centered reach offset in spherical coordinates, $\varphi_o, \theta_o, \rho_o$. This vector is combined with a shoulder-centered representation of the object center, $\varphi_s, \theta_s, \rho_s$, in order to compute a reach target. Once the wrist reaches this point, the reach controller aims for the center of the object, ensuring that the hand approaches the object from the chosen direction. Along with affordance information, the selected offset and grasp type influence the selection of the wrist orientation, w_x, w_y, w_z .

2 Methods

The simulation environment is composed of

1. the Neural Simulation Language (NSL) simulator interfaced with the Open Dynamics Engine (ODE, <http://www.ode.org>) for physics simulation and Java3D (<http://java3d.java.net>) for visualization,
2. a new model of the primate arm and hand, and
3. the implementation of the ILGA model in NSL.

2.1 The simulation environment

In order to embody models in simulated environments, the Java version of the Neural Simulation Language (NSL, Weitzenfeld et al. 2002) simulator has been extended to include 3D graphics functionality and a physics engine. Utilities were developed to create a simulated 3D environment and embed bodies in this environment with limbs connected by hinge, universal, or ball joints. The 3D graphics function, physics engine, and 3D simulation utilities allow NSL models to control bodies in a simulated 3D world and to receive virtual sensory input from the environment.

2.2 Modeling the primate arm and hand

The arm and hand generate forces on the object via the ODE, which in turn is used to estimate grasp stability. To facilitate translation of ILGA to a robotic implementation (see Sect. 4), we chose to model the arm and hand as realistically as possible within the confines of our simulation engine. We implemented a 22-degree of freedom (DOF) arm/hand model using limb proportions based on those for a 7.5-kg monkey (Chan and Moran 2006). The arm has a ball joint at the shoulder with 3 DOFs, a 1-DOF hinge joint at the elbow, and a 3-DOF ball joint at the wrist (Fig. 2). The fingers each have three joints with 1 DOF for the metacarpophalangeal, proximal interphalangeal, and distal interphalangeal joints, while the thumb has one 2-DOF joint at its base (simplifying the carpometacarpal joint) and a 1-DOF metacarpophalangeal joint.

The forward kinematics of the arm is described by a matrix, known as the Jacobian matrix, which transforms joint angles of the shoulder and elbow into resulting wrist positions. The inversion of the Jacobian can then be used to determine required joint angles for a desired wrist position. The Jacobian matrix for the arm with upper arm length l_1 and forearm length l_2 is given by:

$$\mathbf{J} = \begin{bmatrix} -((s_3c_1 - c_3s_2s_1)c_4 - (-s_3s_1 - c_3s_2c_1)s_4)l_2 - (s_3s_1 + c_3s_2c_1)l_1 & -((c_3c_2c_1)c_4 - (-c_3c_2s_1)s_4)l_2 - c_3c_2s_1l_1 \\ -((-c_3c_1 - s_3s_2s_1)c_4 - (c_3s_1 - s_3s_2c_1)s_4)l_2 - (-c_3s_1 + s_3s_2c_1)l_1 & -((s_3c_2c_1)c_4 - (-s_3c_2s_1)s_4)l_2 - s_3c_2s_1l_1 \\ -((-c_2s_1)c_4 - (-c_2c_1)s_4)l_2 - c_2c_1l_1 & -((-s_2c_1)c_4 - (s_2s_1)s_4)l_2 + s_2s_1l_1 \\ -((c_3s_1 - s_3s_2c_1)c_4 - (c_3c_1 + s_3s_2s_1)s_4)l_2 - (-c_3c_1 - s_3s_2s_1)l_1 & -((s_3s_1 + c_3s_2c_1)s_4 - (s_3c_1 - c_3s_2s_1)c_4)l_2 \\ -((s_3s_1 + c_3s_2c_1)c_4 - (s_3c_1 - c_3s_2s_1)s_4)l_2 - (-s_3c_1 + c_3s_2s_1)l_1 & -((c_3s_1 + s_3s_2c_1)s_4 - (-c_3c_1 - s_3s_2s_1)c_4)l_2 \\ 0 & -((-c_2c_1)s_4 - (-c_2s_1)c_4)l_2 \end{bmatrix}$$

The simulated 3D world uses Java3D to maintain a scene graph—a data structure commonly used in computer games to represent the spatial configuration of objects in a scene. Geometric transformations and compound objects are efficiently handled by associating transformation matrices with graph nodes. These matrices can be transformed in order to move an object and all of its child objects (e.g., moving the elbow moves the arm and the hand).

The Open Dynamics Engine (ODE) is used for the physics simulation. ODE is an open source library for simulation of rigid body physics. It contains several joint types and performs collision detection (and contact force application) with friction. When the engine is initialized, NSL maintains a coupling between it and the Java3D representation. At each time step, the physics engine is polled for the position and orientation of the object which is used to update the object's position and orientation in the Java3D scene graph. Forces can be applied to objects in the scene and torque to objects connected with joints.

where $\sin(\theta_1)$ and $\cos(\theta_1)$ are abbreviated as s_1 and c_1 , $\sin(\theta_2)$ and $\cos(\theta_2)$ as s_2 and c_2 , $\sin(\theta_3)$ and $\cos(\theta_3)$ as s_3 and c_3 , and $\sin(\theta_4)$ and $\cos(\theta_4)$ as s_4 and c_4 . The angle θ_1 is the angle of the shoulder in the x -axis, θ_2 is the shoulder angle in the y -axis, θ_3 is the shoulder angle in the z -axis, and θ_4 is the elbow angle.

Proportional-derivative (PD) controllers are used to control each DOF of the arm and hand. A PD controller adjusts a variable until it reaches a target value and includes gain and damping parameters which can be adjusted to produce fast and smooth trajectories. In ILGA, each PD controller applies torque τ at time t to its controlled joint with angle θ in order to reach to a desired value $\hat{\theta}$:

$$\tau(t) = p(\hat{\theta} - \theta(t)) + d\dot{\theta}(t)$$

where p is the gain and d is a damping parameter (see Online Resource 1, Table 1 for parameter values).

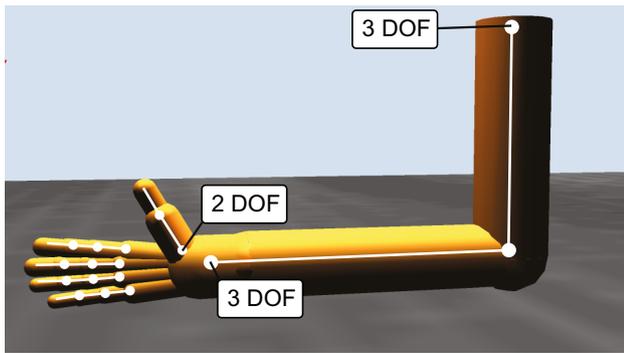


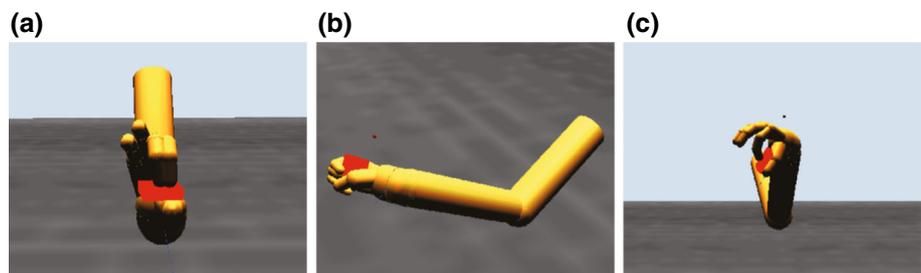
Fig. 2 Unless specified, each joint has 1 DOF. The simulated arm/hand has a total of 22 DOFs

The torque from the joint controllers is applied to the arm and hand segments in the physics model, which updates the position of all of the rigid bodies in the simulated world including the object. Stable grasps are evaluated by monitoring contact points between the fingers and the object in order to provide the reinforcement signal. Figure 3 provides three examples of the arm and hand positioned to grasp an object. In each case, we show a configuration that ILGA can achieve after it has learned to extract affordances from an object and pair them with appropriate grasps.

2.3 Integrated Learning of Grasping and Affordances

The main modules in ILGA are the Feature and Affordance Extraction, Reach and Grasp Planning, and Primary Motor modules (Fig. 4). The Feature Extraction module represents the metric features of graspable objects such as their location, orientation, shape, and oriented surfaces. This information is passed on to the Affordance Extraction module which learns to combine this information into representations of affordances for grasping the object. These representations are used by the Reach and Grasp Planning modules to select motor parameters for reaching to and grasping the object. The Primary Motor module decodes these parameters and controls the reach and grasp movements. Grasp success is monitored by the primary somatosensory (S1) module and is used to generate a reinforcement signal which modifies the connections between the Feature and Affordance Extraction

Fig. 3 Examples of the model successfully grasping the same object with different types of grasps. **a** Tripod, **b** power, **c** precision pinch



modules and those between the Affordance Extraction and Reach and Grasp Planning modules.

The Affordance Extraction module receives basic object information such as location, size, shape, and orientation in the form of population codes from the Feature Extraction module. The Affordance Extraction module contains a self-organizing feature map (SOM) with its learning rate modulated by a global reinforcement signal. SOMs use unsupervised learning (although the version used in ILGA is semi-supervised since the learning rate is modulated by a reinforcement signal) to map a high-dimensional vector space onto a lower-dimensional space. The resulting map preserves topological relationships between vectors (i.e., similar vectors in the high-dimensional input space are mapped onto nearby vectors in the low-dimensional output space, albeit with some possible discontinuities) and identifies dimensions in the input space with the highest variance. The preservation of topological relationships allows the network to generalize to objects it has never seen before but nonetheless elicit activation patterns overlapping those generated by objects in the training set. They are similar to other dimensionality reduction methods such as multidimensional scaling (MDS) and principal components analysis (PCA, Yin 2008), and it has been argued that a similar mechanism organizes representations of high-dimensional spaces in the cerebral cortex (Durbin and Mitchison 1990). Our addition of modulation by a reinforcement signal causes the network to preferentially represent input vectors that are used to generate grasps which result in a positive reward signal by effectively increasing the variance in these regions of input space, a method comparable to the reinforcement-driven dimensionality reduction model (Bar-Gad et al. 2003). The result of training is a network that can extract combinations of object features that afford successful grasps (affordances).

Neurons in the Affordance Extraction module activate dynamic neural fields (DNFs, Amari and Arbib 1977; Erlhagen and Schoner 2002) in the Reach and Grasp Planning modules that select parameters for grasping such as reach offset, wrist rotation, grasp type, and maximum aperture. DNFs utilize cooperation and competition between neurons depending on their preferred stimulus values. In their most basic form, DNFs implement a winner-take-all (WTA) process, resulting in a population code centered on the cell

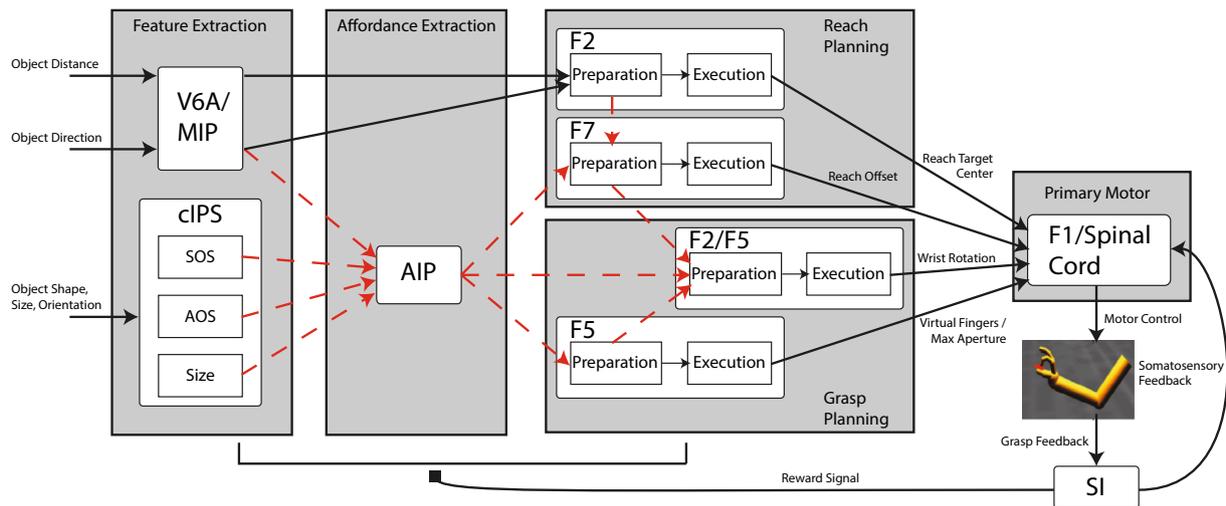


Fig. 4 An overview of the ILGA model. Connections modifiable by reinforcement learning are shown in dashed lines. The parietal regions MIP and V6A provide the premotor region F2 with object position information to plan the reach. V6A, MIP, and the cIPS populations project to AIP, which projects to the signal-related populations of the other premotor regions. Each premotor region selects a value for the parameter it encodes and projects to the primary motor region F1 which controls the movement. The reach planning is performed by area F2 which represents the center of the target object and F7 which selects the

reach offset. Grasp planning is performed by area F5 which selects the virtual fingers and maximum grasp aperture and F2/F5 which selects the wrist rotation. Grasp feedback is returned to somatosensory area S1 which provides the reinforcement signal to the model and somatosensory feedback to F1. Preparation-related premotor populations plan the reach and excite corresponding execution-related premotor populations. Each execution-related premotor population additionally receives tonic inhibitory inhibit (not shown) that is released when a go signal is detected, triggering the performance of the reach and grasp

with the highest mean input. In ILGA, we use one-, two-, and three-dimensional DNFs as WTA networks to select grasp parameters. Each neuron in every DNF has a preferred stimulus value for each dimension of the encoded parameter. In general, the preferred values of each unit could be set arbitrarily, but we set them in a regular fashion such that the population defines a grid in stimulus space (one-, two-, or three-dimensional depending on the dimensionality of the DNF).

Below, we specify the equations used by the Primary Motor module (F1/Spinal Cord in Fig. 4) to decode the activities of the Reach and Grasp Planning module DNFs to obtain the values of various parameters used to plan the reach and grasp movement. Note that the Primary Motor module is *not* modeled as a neural network, but supplies the inputs needed by the physics simulator to apply torques to the joints of the hand and arm. Due to noise in each layer, a random grasp plan can be generated with a small probability, which encourages exploration of the parameter space. Noise levels were set empirically to optimize the trade-off between exploration and exploitation.

Reinforcement learning modifies the connections to and from Affordance Extraction module, resulting in (i) affordances: representations in of combinations of object features relevant for grasping, and (ii) grasp plans: connection weights between the Affordance Extraction and Reach and Grasp Planning modules that bias selection of appropriate grasp motor parameters. The realistic physics simulator we use

allows evaluation of grasp stability based on whether or not the grasp can be maintained. Grasps which do not apply forces to appropriate contact points on the object will cause the object to rotate and slip from the hand's grasp. Positive reinforcement is given by the realization of a stable grasp of the target object, and negative reinforcement is given for grasps that do not contact the object or are unstable enough to allow the object to slip from the hand. The result is that Affordance Extraction neurons are shaped to provide “better” affordance input for the Reach and Grasp Planning module, which in turn expands the repertoire of grasp actions providing more data points for Affordance Extraction learning. Eventually, a stable state is reached when affordance representations are nearly static and grasp performance reaches an upper limit. When this dual learning system stabilizes, the model is endowed with a set of affordance extraction and robust grasp planning mechanisms.

2.4 Reinforcement

Rather than relying on basic extrinsic reward signals from primary reinforcers, it has been suggested that hierarchies of skills could be learned based on intrinsic reward signals that reinforce unexpected salient events (Chentanez et al. 2004). Infants as young as 3–4 months old have experience with objects being placed in their grasp and seem intrinsically motivated to reach for and grasp objects themselves (Thelen et al. 1993). The ILGM thus posited an intrinsic

“joy of grasping” as the reward stimulus generated from sensory feedback resulting from the stable grasp of an object. We use the same signal to train the connection weights in this model using reinforcement learning (Sutton and Barto 1998). However, this model uses a more realistic physics simulator than the ILGM, taking into account not only kinematics but also dynamics. This makes motor control a much more difficult task, but simplifies grasp stability evaluation (see below). Another consequence is that since the object can be moved, hand–object collision can knock the object out of reach, making successful grasps much less likely to occur by chance during trial-and-error learning. In order to increase the probability of successful grasps, we pre-train the connection weights that determine the direction of hand approach to the object and the wrist orientation using a more basic reinforcement signal, what may be called the “joy of palm contact.” After pretraining these connection weights, the majority of attempted grasps make at least transient palm contact with the object, increasing the number of stable grasps during the next stage of training. We therefore model palm contact as an unexpected and intrinsically rewarding salient event during early training. As the infant becomes more proficient in reaching for and orienting their hand toward the object, we suggest that palm contact is less novel and therefore less intrinsically rewarding. At this stage, stable grasps are the primary intrinsically rewarding events.

The ILGM used a kinematic simulator that did not handle dynamics and therefore had to use an ad hoc scheme to estimate grasp stability. Since we use a physics simulator that handles rigid body dynamics including friction, grasp stability evaluation is more direct. We ran all simulations reported here with gravity turned off and the object suspended at various locations in order to simplify control of the arm and hand. Even without gravity, the physics simulation will cause the object to slip from the hand’s grasp if the grasp is unstable. The simulator informs the model of the list of contact points between the hand and the object. If two contact points are achieved that define an opposition axis (a vector connecting them in space) that passes through the object, and these contact points are maintained for 2 seconds of simulation time, the grasp is declared successful. Note that none of the other modeled regions have any notion of contact points—the grasp is planned and controlled in an open loop manner. However, contact point feedback could be used to learn internal models for feedback-based grasp control (see Sect. 4).

2.5 Primary motor module: reach and grasp generation

With the Primary Motor module (F1 and the Spinal Cord in Fig. 4), we leave the domain of neural networks: The module decodes the motor parameters for the reach and grasp

from the activities of the premotor populations and directs the actual movement by setting the joint angle targets of the PD controllers for each DOF. While future versions of the model may implement this module in a more neurobiologically plausible way, we chose to implement it using techniques from robotics in order to simplify the neural components of the model and focus on parietal and premotor neural activity. Although non-neural, the Primary Motor module serves to achieve the planned grasp end state in order to evaluate the success of the chosen parameter values in producing a stable grasp.

The wrist rotation, reach, and grasp components of the movement are handled by separate controllers (Fig. 5). The reach and grasp components are coupled by starting the pre-shape phase of the grasp once the reach target has been determined, and triggering the enclose phase once the hand reaches a certain distance from the object or achieves palm contact. Palm or inner thumb contact is also used as a signal to stop the reach controller at the current wrist position. Each controller decodes the input it receives from premotor execution-related populations, transforms this input along with some proprioceptive or tactile signals, and sets joint angle targets for PD controllers that apply torque to each joint.

Each neuron in every DNF has a preferred stimulus value for each dimension (\hat{x} , \hat{y} , \hat{z}) of the parameter it encodes. Reach and grasp parameter values were decoded from the activity of each DNF using the center-of-mass technique (Wu et al. 2002). Since noise can greatly bias this form of decoding in small populations, we only include the activities of neurons that pass a threshold, ξ (set to 0.01 in these simulations). For a one-dimensional DNF, the encoded value, x , was estimated as:

$$x = \frac{\sum_i (\mathbf{f}(i) \hat{x}(i))}{\sum_i \mathbf{f}(i)}$$

where the sums are over all neurons i with an activation greater than or equal to the threshold, ξ . Similarly, the encoded x , y value was estimated from a two-dimensional DNF as:

$$x = \frac{\sum_i \left(\sum_j (\mathbf{f}(i, j) \hat{x}(i)) \right)}{\sum_i \left(\sum_j \mathbf{f}(i, j) \right)},$$

$$y = \frac{\sum_j \left(\sum_i (\mathbf{f}(i, j) \hat{y}(j)) \right)}{\sum_i \left(\sum_j \mathbf{f}(i, j) \right)}$$

and the x , y , z value encoded by a three-dimensional DNF as:

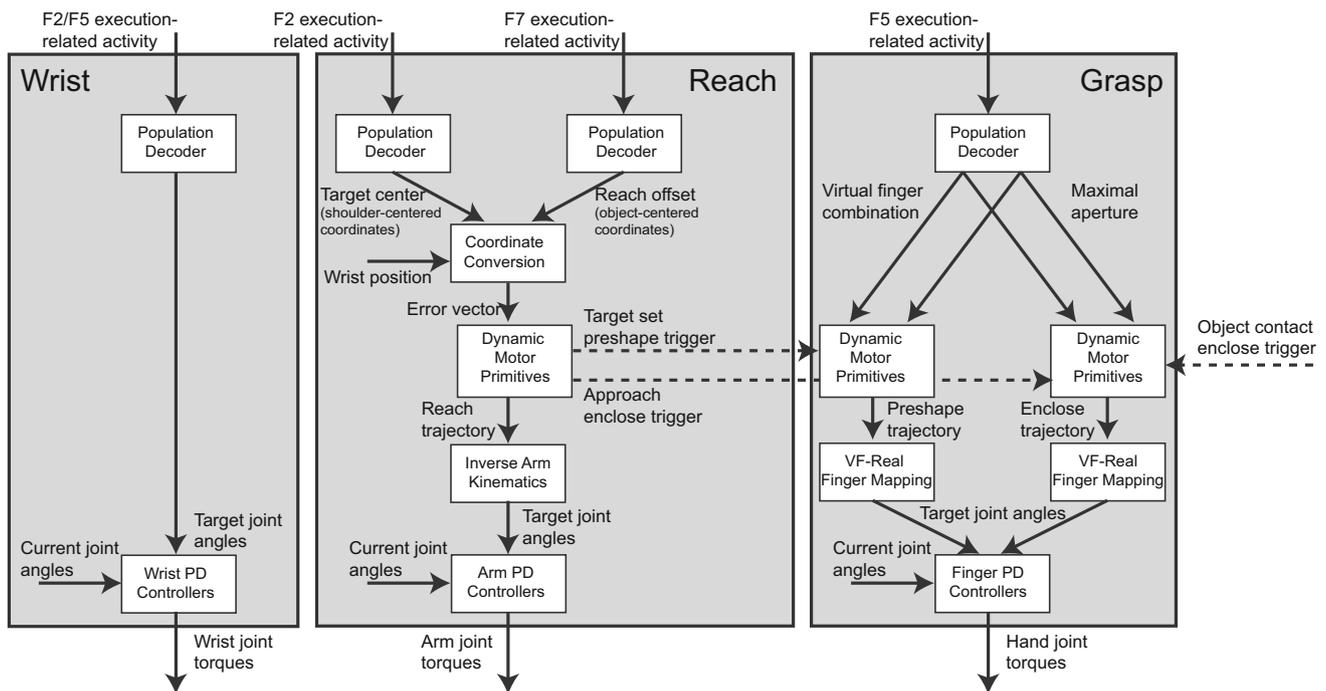


Fig. 5 Wrist, reach, and grasp motor controllers. Each uses population decoders to decode reach and grasp parameter values from premotor inputs and set joint angle targets for PD controllers which move the limbs by applying torque to the joints. The reach motor controller combines the shoulder-centered object position, object-centered reach offset, and current wrist position to compute a wrist error vector. The error vector is used to set goal values for dynamic movement primitives,

which generate a reach trajectory for the wrist. An inverse arm kinematics module computes target joint angles for each target wrist position. The grasp motor controller contains dynamic movement primitives for the preshape and enclose phases that are triggered by reach and tactile events. These dynamic movement primitives generate normalized trajectories for each virtual finger that are converted into target joint angles by VF → real finger mapping modules

$$x = \frac{\sum_i \left(\sum_j \left(\sum_k \left(\mathbf{f}(i, j, k) \hat{\mathbf{x}}(i) \right) \right) \right)}{\sum_i \left(\sum_j \left(\sum_k \mathbf{f}(i, j, k) \right) \right)},$$

$$y = \frac{\sum_i \left(\sum_j \left(\sum_k \left(\mathbf{f}(i, j, k) \hat{\mathbf{y}}(i) \right) \right) \right)}{\sum_i \left(\sum_j \left(\sum_k \mathbf{f}(i, j, k) \right) \right)},$$

$$z = \frac{\sum_i \left(\sum_j \left(\sum_k \left(\mathbf{f}(i, j, k) \hat{\mathbf{z}}(i) \right) \right) \right)}{\sum_i \left(\sum_j \left(\sum_k \mathbf{f}(i, j, k) \right) \right)}$$

The simplest controller, the *wrist rotation controller*, decodes the target wrist angles from its premotor input and passes them as target angles to the wrist PD controllers. The *reach controller* combines the output of F2 encoding the center of the object, with the output of F7 encoding the object-centered reach offset into a wrist-centered error vector defining the initial target of the reach. The reach controller couples a trajectory planning mechanism (dynamic movement primitives, DMPs, Ijspeert et al. 2002) with an inverse arm kinematics module. Given a reach target location, the reach planner uses DMPs to generate a trajectory of desired wrist locations to

reach it starting from the current wrist position. DMPs can generate arbitrary trajectories and dynamically adapt to new goals. They are defined by the following differential equation:

$$\dot{v} = uK \left(\frac{\sum_i \psi_i(u) (c_i + x_0)}{\sum_i \psi_i(u)} - x \right) + (1 - u) K (g - x) - Dv$$

where x is the current value of the controlled variable (the position of the wrist in this case), x_0 is the initial value (the starting position of the wrist), v is the current target velocity of the variable, c_i are equilibrium points of linear acceleration fields with nonlinear basis functions ψ_i , g is the goal value (the target reach position), K and D are gain and damping parameters, and u is a phase variable which can be used to scale the duration of the movement. DMPs therefore generate a trajectory from x_0 to g that can be straight or parameterized to take any arbitrary path (which we do not exploit in this model but leave open the possibility for future work). In the reach module, the output of the DMP, $\hat{x}(t)$, is used as an actual target position for the wrist that is input to the inverse kinematics controller at each time step.

Given a desired wrist location, the inverse arm kinematics controller computes the required wrist displacement and then uses the pseudo-inverse of the Jacobian to compute the required joint rotations to bring the wrist to that position. The body's Jacobian matrix describes how changes in shoulder ($\theta_1, \theta_2, \theta_3$) and elbow (θ_4) angles result in changes in the wrist's 3D position (x, y, z):

$$\begin{bmatrix} \dot{x} \\ \dot{y} \\ \dot{z} \end{bmatrix} = \mathbf{J} \begin{bmatrix} \dot{\theta}_1 \\ \dot{\theta}_2 \\ \dot{\theta}_3 \\ \dot{\theta}_4 \end{bmatrix}$$

The inverse of the Jacobian matrix then describes how much each joint must rotate in order to effect a desired wrist displacement. The Jacobian is not invertible, so we use the pseudo-inverse:

$$\mathbf{J}^+ = \mathbf{J}^T (\mathbf{J}\mathbf{J}^T)^{-1}$$

Each required joint rotation is used to input the target joint angle into the PD controller for that DOF:

$$\hat{\theta} = \theta + \mathbf{J}^+ \begin{bmatrix} \dot{x} \\ \dot{y} \\ \dot{z} \end{bmatrix}$$

where $\dot{x}, \dot{y}, \dot{z}$ describe the desired wrist displacement.

The *grasp controller* controls the timing of the preshape and enclosure phases of the grasp. Depending on the phase of the grasp, the controller translates the selected virtual finger combination and maximum aperture into final target joint angles for each finger (see Online Resource 1, Table 2). Each virtual finger combination is associated with a preshape hand configuration with certain finger angles that can be modulated by the maximum aperture parameter, and a set of fingers to control during the enclose phase. The possible virtual finger combinations define the following grasps: precision pinch (index finger and thumb extended then enclosed), tripod grasp (index and middle fingers and thumb extended then enclosed), whole-hand prehension (all fingers and thumb extended then enclosed), and side grasps (all fingers enclosed and thumb extended then enclosed).

The grasp controller uses DMPs to control the timing of the preshape and enclosure phases of the grasp, but here the DMP output is interpreted as a normalized timing signal, rather than a physical target value. This is accomplished by setting the initial input to the DMP at 0 and the target to 1, and using its output at each time step to interpolate between the current finger joint angles and the final target angles in order to generate targets for the PD controllers. The preshape DMP is triggered as soon as a reach target is set, while the enclose DMP is triggered once the wrist reaches a certain threshold

distance from the object, κ , or once the palm contacts the object, whichever happens first. ILGA thus captures a period of development where the innate infant grasp reflex is present, but some predictive hand enclosure during grasping starts to emerge.

2.6 Parietal module: feature/affordance extraction

The populations in the Feature Extraction and Affordance Extraction modules are based on findings from a series of primate single-unit recording studies (Sakata et al. 1998; Galletti et al. 2003) in parietal cortex. These experiments found that neurons in the anterior intraparietal sulcus (AIP) are responsive to 3D features of objects relevant for manipulation (Murata et al. 2000). Neurons in the caudal intraparietal sulcus (cIPS) are selective for objects and their surfaces at preferred orientations. Subsets of these neurons have been described as axis orientation selective (AOS) and surface orientation selective (SOS, Taira et al. 1990). It has been suggested that the regions V6A and MIP are involved in encoding the direction of movement required to bring the arm to potential reach targets (Rizzolatti et al. 1998; Galletti et al. 2003). Both cIPS and V6A project to AIP (Shipp et al. 1998; Nakamura et al. 2001). Based on these studies, in ILGA the areas cIPS, MIP, and V6A extract object features and location and V6A/MIP and cIPS project this information to AIP for grasp affordance extraction.

The neurophysiological experiments upon which the model parietal regions are based used simple objects (cube, cylinder, sphere, etc.) as stimuli. However, more complex objects such as hammers or coffee cups contain multiple affordances. We suggest that the dorsal visual stream analyzes objects in terms of a set of multiple affordances and accurately represents the metrics of different object components to provide data essential for grasp affordance extraction (Fig. 6).

2.6.1 V6A/MIP

The model region V6A/MIP represents the shoulder-centered direction and distance of each object component in spherical coordinates as two- and one-dimensional population codes, respectively. The module provides this information to parietal area AIP for affordance extraction and premotor area F2 for programming the reach. Area V6A contains mostly visual cells (Galletti et al. 1997; Rizzolatti et al. 1998) with receptive fields covering the whole visual field and representing each portion of it multiple times (Galletti et al. 1993, 1999). The so-called real-position cells are able to encode the spatial location of objects in the visual scene with visual receptive fields that remain anchored despite eye movements (Galletti et al. 1993). Intermingled with real-position

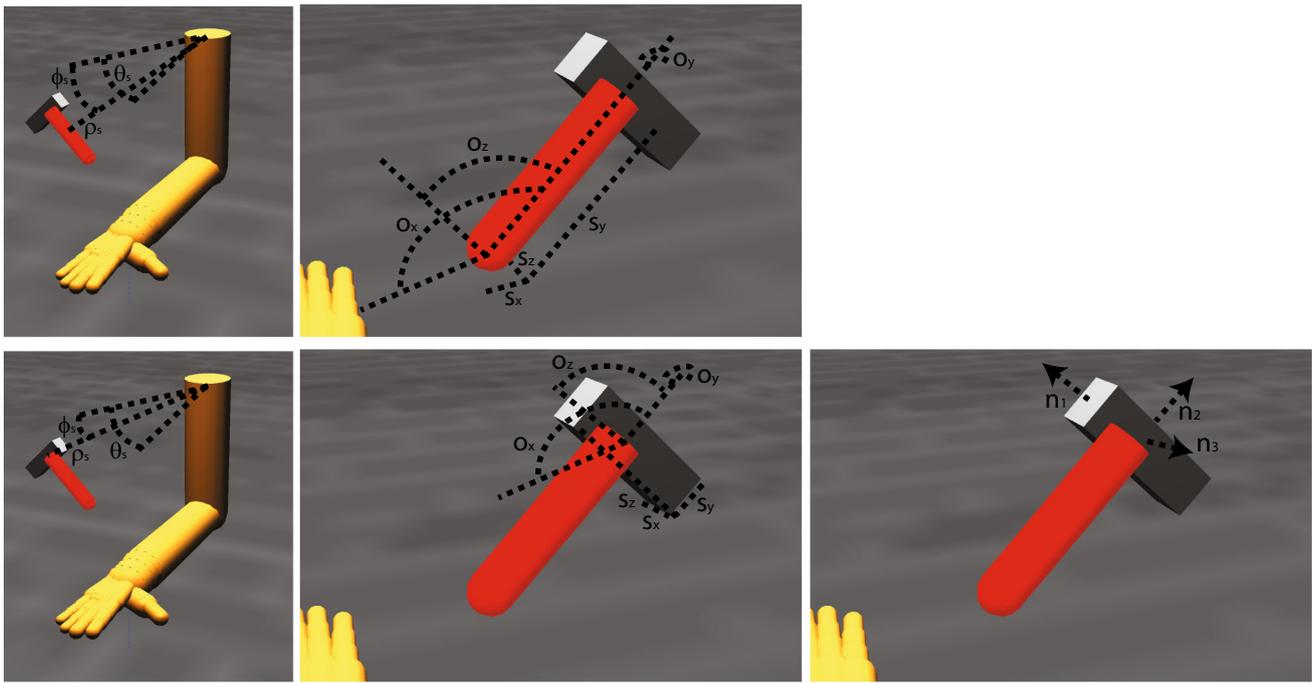


Fig. 6 Object component variables represented in each parietal region for the handle (top row) and head (bottom row) of a hammer. The areas V6A and MIP represent the shoulder-centered direction of the center of the primitive, φ_s and θ_s , and the shoulder-centered distance, ρ_s , and

cIPS represents the object component's orientation, o_x , o_y , o_z , size, s_x , s_y , s_z , and orientation of surface normal vectors (n_1 , n_2 , and n_3 in this case)

cells are retinotopic cells, whose visual receptive fields shift with gaze, suggesting that the region is involved in converting coordinates from retinotopic to head- or body-centered reference frames (Galletti et al. 1993). Many neurons in V6A only respond when the arm is directed toward a particular region of space (Galletti et al. 1997; Fattori et al. 2005). Lesions of the region result in misreaching with the contralateral arm (Battaglini et al. 2002). It has thus been suggested that the region is involved in encoding the direction of movement in at least head-centered coordinates (but we use shoulder-centered coordinates for the single arm in the model) required to bring the arm to potential reach targets (Rizzolatti et al. 1998; Galletti et al. 2003). Area V6A receives input from central and peripheral visual field representations in V6 (Shipp et al. 1998) and projects to the premotor area F2 (Matelli et al. 1998; Shipp et al. 1998; Luppino et al. 2005).

Distance is represented in a one-dimensional population code, with each unit having a preferred distance, \hat{d}_s , uniformly distributed between 0 and 1 m. The activity of each unit, i , at time t is given by a Gaussian population code over the unit's preferred distance and the actual shoulder-centered distance, d_s , to each object component, p :

$$\mathbf{DIST}(i, t) = \sum_p \left(e^{-\frac{(\hat{d}_s(i) - d_s(p, t))^2}{2\sigma_{\text{DIST}}^2}} \right) + \varepsilon_{\text{DIST}}$$

where σ_{DIST} is the population code width and $\varepsilon_{\text{DIST}}$ is a noise term. For all codes, we add a Gaussian noise term in order to simulate stochasticity in neural activity. A table of all parameter values is given in Online Resource 1, Table 3. Direction is encoded in a two-dimensional population code with each unit selective for both particular azimuth and elevation values of the shoulder-centered object direction:

$$\mathbf{DIR}(i, j, t) = \sum_p \left(e^{-\frac{(\hat{\theta}_s(i, j) - \theta_s(p, t))^2 + (\hat{\varphi}_s(i, j) - \varphi_s(p, t))^2}{2\sigma_{\text{DIR}}^2}} \right) + \varepsilon_{\text{DIR}}$$

where $\theta_s(p, t)$ is the azimuth angle, $\varphi_s(p, t)$ is the elevation angle of object component p at time t in a shoulder-centered reference frame, σ_{DIR} is the population code width, and ε_{DIR} is a noise term. Each unit of the population, i, j , has preferred angles, $\hat{\theta}_s, \hat{\varphi}_s$, with $\hat{\theta}_s$ uniformly distributed between 0 and π , and $\hat{\varphi}_s$ uniformly distributed between $-\pi$ and 0.

2.6.2 cIPS

The model region cIPS contains three populations that represent the object orientation, size, and visible surface normal vectors as three-dimensional population codes. The caudal intraparietal sulcus (cIPS) is a region located the caudal

part of the lateral bank and fundus of the intraparietal sulcus (Shikata et al. 1996). The cIPS receives input mainly from V3A, whose neurons are sensitive to binocular disparity and have small, retinotopic receptive fields (Sakata et al. 2005), and projects primarily to the anterior intraparietal sulcus (AIP, Nakamura et al. 2001). Neurons in area cIPS have large receptive fields (10–30 degrees in diameter) with no retinotopic organization (Tsutsui et al. 2005). Two functional classes of neurons in area cIPS have been described: surface-orientation-selective (SOS) neurons that are selective to the orientation of flat surfaces, and axis-orientation-selective (AOS) neurons that respond best to an elongated object whose principal axis is oriented in a particular direction. Both types of neurons respond best to binocular stimuli (Sakata et al. 1997) and are spatially intermingled (Nakamura et al. 2001). Muscimol-induced inactivation of this region disrupts performance on a delayed match-to-sample task with oriented surfaces using perspective and disparity cues (Tsutsui et al. 2001, 2005). Both types of cells include some neurons that are selective for the object’s dimensions (Kusunoki et al. 1993; Sakata et al. 1998). Again, these neurons have only been tested with simple objects. We suggest that they actually encode the features of object components that comprise complex objects. To simplify the model, we include AOS and SOS cells as well as one population that encodes the size of each object component.

Axis-orientation-selective (AOS) cells prefer bars tilted in the vertical, horizontal, or sagittal planes (Sakata et al. 1998, 1999). Some are selective for shape (rectangular versus cylindrical) and probably represent surface curvature (Sakata et al. 2005). Their discharge rate increases monotonically with object length, and their width response curve is monotonically decreasing in the 2–32 cm range. It is thought that these cells integrate orientation and width disparity cues to represent principal axis orientation (Sakata et al. 1998).

Surface-orientation-selective (SOS) cells are tuned to the surface orientation in depth of flat and broad objects (Shikata et al. 1996; Sakata et al. 1997, 1998, 1999). These cells respond to a combination of monocular and binocular depth cues (texture and disparity gradient cues) in representing surface orientation (Sakata et al. 2005). Neurons sensitive to multiple depth cues are widely distributed and spatially intermingled with those sensitive to only one depth cue (Tsutsui et al. 2005).

We model AOS cells as two subpopulations—one selective for rectangular and one for cylindrical objects. Each subpopulation is a three-dimensional population code, with each neuron i, j, k , selective for a combination of the components of the object’s main axis orientation, o_x, o_y , and o_z . When the object is cylindrical, the activity of each unit in the cylindrical AOS population is given by:

$$\begin{aligned} & \text{CYL}(i, j, k, t) \\ &= \sum_p \left(e^{-\frac{(\hat{o}_x(i, j, k) - o_x(p, t))^2 + (\hat{o}_y(i, j, k) - o_y(p, t))^2 + (\hat{o}_z(i, j, k) - o_z(p, t))^2}{2\sigma_{\text{CIPS}}^2}} \right) \\ &+ \varepsilon_{\text{CIPS}} \end{aligned}$$

where $\hat{o}_x(i, j, k)$, $\hat{o}_y(i, j, k)$, and $\hat{o}_z(i, j, k)$ are the preferred orientations of the unit i, j, k in the x, y , and z dimensions. When the object is not cylindrical, each unit’s activity in the cylindrical AOS population is given by the noise term. Given a rectangular object, the activity of each unit in the rectangular AOS population is defined by:

$$\begin{aligned} & \text{RECT}(i, j, k, t) \\ &= \sum_p \left(e^{-\frac{(\hat{o}_x(i, j, k) - o_x(p, t))^2 + (\hat{o}_y(i, j, k) - o_y(p, t))^2 + (\hat{o}_z(i, j, k) - o_z(p, t))^2}{2\sigma_{\text{CIPS}}^2}} \right) \\ &+ \varepsilon_{\text{CIPS}} \end{aligned}$$

Similarly, the noise term determines the activation of these units when the object is not rectangular.

We model the SOS population as a noisy three-dimensional Gaussian population code over the normal vector ($\mathbf{n}_f, \mathbf{n}_f, \mathbf{n}_f$) of each visible surface, f , of a rectangular object component, p :

$$\begin{aligned} & \text{SOS}(i, j, k, t) \\ &= \sum_p \left(\sum_f \left(e^{-\frac{(\hat{n}_x(i, j, k) - \mathbf{n}_f(p, x, t))^2 + (\hat{n}_y(i, j, k) - \mathbf{n}_f(p, y, t))^2 + (\hat{n}_z(i, j, k) - \mathbf{n}_f(p, z, t))^2}{2\sigma_{\text{CIPS}}^2}} \right) \right) \\ &+ \varepsilon_{\text{CIPS}} \end{aligned}$$

Each unit has a preferred value in each dimension ($\hat{n}_x, \hat{n}_y, \hat{n}_z$) uniformly distributed between -1 and 1 . If the object is not rectangular, the activity of each unit in these populations is determined by the noise term.

The size population, S , represents the size of an object component p in each dimension (s_x, s_y, s_z), with the activity of each unit, i, j, k , given by:

$$\begin{aligned} & \text{S}(i, j, k, t) \\ &= \sum_p \left(e^{-\frac{(\hat{s}_x(i, j, k) - s_x(p, t))^2 + (\hat{s}_y(i, j, k) - s_y(p, t))^2 + (\hat{s}_z(i, j, k) - s_z(p, t))^2}{2\sigma_{\text{CIPS}}^2}} \right) \\ &+ \varepsilon_{\text{CIPS}} \end{aligned}$$

where σ_{CIPS} is the population code width and $\varepsilon_{\text{CIPS}}$ is a noise term.

2.6.3 AIP

The anterior intraparietal area AIP is located on the lateral bank of the anterior intraparietal sulcus and contains visually response neurons selective for 3D features

of objects, motor-dominant neurons that only respond during grasping, and visuomotor neurons that are activated by grasping and modulated by sight of the object (Sakata et al. 1998). The region receives its main input from area cIPS (Nakamura et al. 2001), but also receives input from V6A (Shipp et al. 1998) and projects most strongly to the premotor region F5 (Borra et al. 2007).

In ILGA, the AIP module receives input from each population of the V6A/MIP and cIPS modules. While the distance to the object is important for parameterizing the reach as well as the coordination of the reach and grasp movements, it is not important for specifying the grasp itself, and therefore we do not model a connection between the distance population of V6A/MIP and AIP. The SOM is a toroidal grid of 40×40 units, with the input vector, \mathbf{I} , constructed by concatenating the activity vectors of each V6A/MIP and cIPS population into one vector, which is then normalized. The activity of each AIP unit with indices i and j is given by:

$$\mathbf{AIP}(i, j, t) = \mathbf{W}_{\text{AIP}}(i, j, t)\mathbf{I}(t) + \varepsilon_{\text{AIP}}$$

The weights \mathbf{W}_{AIP} are initialized to small random values. Weight training uses a form of competitive learning. This learning algorithm modifies the connection weights of neurons in a neighborhood surrounding the neuron with the greatest activity level. The size of this neighborhood gets smaller over time and in ILGA, both the neighborhood size and learning rate are increased by the presence of the reinforcement signal. At the beginning of training, the neighborhood is broad and the learning rate is high. This causes the self-organization to take place on the global scale. As training progresses and the neighborhood size decreases, the weights converge to local estimate of the training input vectors. The modulation of the learning rate and neighborhood function by the global reinforcement signal ensures that input vectors that are used to plan stable grasps become represented by more units in the SOM at the expense of input vectors that result in failed grasps.

Given the input vector, \mathbf{I} , the AIP unit with the most similar weight vector is determined as the best matching unit (BMU). The similarity metric we used was the Euclidean distance between the vectors. The weights of the BMU and all units within its “neighborhood” are adjusted in the direction of the input vector:

$$\begin{aligned} \mathbf{W}_{\text{AIP}}(i, j, t + 1) = & \mathbf{W}_{\text{AIP}}(i, j, t) \\ & + \Theta(i, j, T)\alpha(t, T)(\mathbf{I}(t) \\ & - \mathbf{W}_{\text{AIP}}(i, j, t)) \end{aligned}$$

where Θ is the reinforcement-dependent neighborhood function, T is the current training epoch, and α is the reinforcement-dependent learning rate. The neighborhood

function is a Gaussian function over the Euclidean distance, β , between the indices of the neuron i, j and those of the BMU. The Gaussian is truncated at a certain radius, r , that also defines its spread:

$$\Theta(i, j, t) = \begin{cases} e^{-\frac{(\beta-0)^2}{2r(T)^2}} & \text{if } \beta < r(T) \\ 0 & \text{otherwise} \end{cases}$$

The radius of the neighborhood function shrinks over the duration of training, but is expanded by the global reinforcement signal, $rs(t)$:

$$r(t, T) = r_0 e^{-\frac{T}{\lambda}} + rs(t)$$

where r_0 is the initial radius and the parameter λ determines the rate at which the radius decreases (see Online Resource 1, Table 3 for parameter values). The learning rate also decreases over the duration of training:

$$\alpha(t, T) = \alpha_0 e^{-\frac{T}{\lambda}} + rs(t)$$

where α_0 is the initial learning rate.

2.7 Premotor module: Reach and Grasp Planning

Each region in the Reach and Grasp Planning module contains one population of preparation-related cells and one population of execution-related cells. Execution-related cells discharge on movement onset, while signal-related cells show anticipatory activity prior to the start of movement. These broad categories of cells have been found in several premotor areas (Kurata 1994; Wise et al. 1997; Cisek and Kalaska 2002). In this model, preparation-related cells receive external input and project topologically to execution-related cells with hard-wired, fixed connections. Execution-related cells additionally receive tonic inhibition which is released when the go signal is detected, ensuring that the movement does not begin until the signal is observed. The basal ganglia are typically implicated in disinhibition of planned movements (Kropotov and Etlinger 1999) and reinforcement learning (Barto 1995). While models of the basal ganglia exist that could provide the tonic inhibition and reinforcement signals in ILGA (Gurney et al. 2001), we simply provide these inputs procedurally.

The Reach and Grasp Planning module contains several subpopulations based on various regions of dorsal and ventral premotor cortex including F2, F5, and F7. Each of these populations selects grasp motor parameters based on input from AIP and V6A/MIP. In ILGA, areas F2 and F7 are mainly involved with specifying the reach, with F2 selecting the center of the object component to reach to, and area F7 selecting an object-centered offset from that center for the reach target

(Fig. 1). Area F5 selects the grasp type and maximal aperture, and F2/F5 selects the wrist orientation. All the motor parameters handled by the premotor cortex modules in both ILGM and ILGA are kinematics parameters—dynamics is completely ignored by these modules and handled entirely by the primary motor module.

Although there is evidence which supports the role of the primary motor cortex in both kinematic and dynamic encoding (Kalaska 2009), it has been shown that the evidence in favor of kinematic encoding could be an epiphenomenon of multidimensional muscle force coding (Todorov 2000). Lesions of the premotor cortex result in deficits in movement kinematics (Freund 1990; Gallese et al. 1994). However, neurons in dorsal and ventral premotor cortex have been found that correlate with movement dynamics variables (Xiao et al. 2006) and a series of studies has shown that the motor cortex may be involved in transforming kinematic variables from extrinsic to intrinsic reference frames (Takei et al. 1999, 2001, 2003). Here, we make the simplifying assumption that the premotor cortex specifies movement kinematic variables which are translated into muscle forces by the primary motor cortex and spinal cord.

In this model, each preparation- and execution-related population is simulated as a DNF. Each DNF contained leaky integrator neurons with sigmoidal transfer functions. Each neuron consisted of a membrane potential variable, u , and a firing rate variable, f . The membrane potential was computed by integrating the weighted input over time with input from other neurons in the DNF:

$$\tau \frac{du}{dt} = -u + h + \mathbf{IN} + \mathbf{f} * \mathbf{W}_{\text{DNF}} + \epsilon_{\text{DNF}}$$

where \mathbf{u} and \mathbf{f} are the population membrane potentials and firing rates, h is the baseline activation (see Online Resource 1, Table 4 for parameter values), \mathbf{IN} is the weighted input, $*$ is the convolution operator, \mathbf{W}_{DNF} is the winner-take-all (WTA) weight kernel defined below, and ϵ_{DNF} is a noise term. The firing rate, f , is then a sigmoid function of the membrane potential:

$$\mathbf{f} = \frac{1}{1 + e^{-\beta(u-u_0)}}$$

where β and u_0 are parameters of the sigmoid.

The weight kernel was set as a Gaussian that was negative except for a center peak, implementing the WTA functionality. The kernel was set to be twice as large as the population to ensure global competition. For one-dimensional DNFs, the weight kernel was given by:

$$\mathbf{W}_{\text{DNF}}(i) = w_{\text{excite}} e^{-\frac{(i-N/2)^2}{2\sigma_{\text{DNF}}^2}} - w_{\text{inhibit}}$$

where w_{excite} is the height of the peak of the Gaussian, w_{inhibit} is the level of inhibition, N is the size of the population, and σ_{DNF} is the width of the Gaussian. Similarly, the two-dimensional weight kernel was defined as:

$$\mathbf{W}_{\text{DNF}}(i, j) = w_{\text{excite}} e^{-\frac{((i-N/2)^2 + (j-N/2)^2)}{2\sigma_{\text{DNF}}^2}} - w_{\text{inhibit}}$$

and the three-dimensional DNF was given by:

$$\mathbf{W}_{\text{DNF}}(i, j, k) = w_{\text{excite}} e^{-\frac{((i-N/2)^2 + (j-N/2)^2 + (k-N/2)^2)}{2\sigma_{\text{DNF}}^2}} - w_{\text{inhibit}}$$

Every DNF in each modeled region used the same parameters (Table 3), which were determined empirically.

The tonic inhibitory input to each execution-related population, GP , was set to 10 before the go signal was detected and 0 once it appeared. Therefore, preparation-related cells in ILGA plan the movement, while the activation of execution-related cells triggers its onset. Reinforcement learning is applied to the afferent connection weights of the preparation-related cells, using the activity of the corresponding execution-related population as an eligibility trace. Eligibility traces are commonly used in reinforcement learning in order to assign credit to the appropriate connection weight for delayed reward (Singh and Sutton 1996). This is typically a decaying copy of the activated neurons, but since the delay between preparation-related cell activity and the achievement of a stable grasp that elicits a reward can be quite long, we use the activity of corresponding execution-related cells as the eligibility trace.

2.7.1 F2

Within the premotor cortex, the caudal portion F2 most likely codes reach movements in a shoulder-centered reference frame (Caminiti et al. 1991; Rizzolatti et al. 1998; Cisek and Kalaska 2002). Many of the cells in F2 have broad directional tuning, and their population activity appears to encode a vector representing the direction of arm movement and not the position of the end target (Weinrich and Wise 1982; Caminiti et al. 1991). The region was first defined by Matelli et al. (1985) and was later subdivided into the F2 dimple (F2d) and ventrorostral (F2vr) subregions (Matelli et al. 1998). Visual inputs to area F2 come mainly from the superior parietal lobe (Johnson et al. 1993; Caminiti et al. 1996). The subregion F2vr receives projections from area V6A (Shipp and Zeki 1995) and the medial intraparietal area MIP (Matelli et al. 1998; Shipp et al. 1998; Marconi et al. 2001). The main output of F2 projects to F1 (Dum and Strick 2005).

Area F2 contains a rostro-caudal gradient of cell types with preparation-related cells found predominantly in F2vr

and execution-related cells located in F2d, the caudal portion adjacent to F1 (Tanne et al. 1995; Johnson et al. 1996). Preparation-related cells are 43 % of F2 neurons and respond to the visual target for reaching, while execution-related cells have changes in activity that are synchronized with the onset of movement (Weinrich and Wise 1982). Some execution-related cells are only active after the go signal and these are more common caudally (Crammond and Kalaska 2000). This categorization of cells seems to correspond to a similar modality-based classification used by Fogassi et al. (1999) and Raos et al. (2004), which describes cells as purely motor, visually modulated, or visuomotor. Purely, motor cells are not affected by object presentation or visual feedback of the hand, visually modulated cells discharge differentially when reaching in the light vs. dark, and visuomotor cells discharge during object fixation without movement. Most of visually modulated or visuomotor cells are in F2vr (Fogassi et al. 1999) and therefore likely correspond to the preparation-related cells described by Crammond and Kalaska (2000). Our model thus subdivides F2 into rostral and caudal regions (F2vr and F2d, respectively) and simplifies the distribution of cell types by confining preparation-related cells to the rostral region and execution-related cells to the caudal region.

Most cells in F2 are sensitive to amplitude and direction, with very few cells sensitive to only amplitude (Fu et al. 1993; Messier and Kalaska 2000). However, muscimol inactivation caused increases in directional errors when conditional cues are presented, but amplitude and velocity were unchanged (Kurata and Hoffman 1994). Neurons in the dorsal premotor cortex have more recently been shown to encode the relative position of the eye, hand, and goal (Pesaran et al. 2006), but we do not vary the eye position in these simulations and this influence is therefore constant. We thus decode the output of F2d as a population code with each cell having a preferred spherical coordinate in a shoulder-centered reference frame. Note that the issue of which reference frame is used in the reach circuit is still debated.

We model the F2vr region as two DNFs encoding the shoulder-centered direction (**F2vrDIR**) and distance (**F2vrRAD**) of the target object in spherical coordinates. The input to each DNF is given by:

$$\begin{aligned} \mathbf{IN}_{\text{F2vrDIR}}(t) &= \mathbf{DIR}(t)\mathbf{W}_{\text{DIR}\rightarrow\text{F2}} + \boldsymbol{\varepsilon}_{\text{F2}} \\ \mathbf{IN}_{\text{F2vrRAD}}(t) &= \mathbf{DIST}(t)\mathbf{W}_{\text{DIST}\rightarrow\text{F2}} + \boldsymbol{\varepsilon}_{\text{F2}} \end{aligned}$$

where the matrices $\mathbf{W}_{\text{DIR}\rightarrow\text{F2}}$ and $\mathbf{W}_{\text{DIST}\rightarrow\text{F2}}$ define the weights of the projections from V6A/MIP to F2. Since we assume that reaching ability has already developed, these weights are not subject to learning and set according to the following rule:

$$\mathbf{W}(i, j) = 3\mathbf{I}$$

where \mathbf{I} is the identity matrix. This results in F2vr faithfully selecting the center of the object component (as signaled by V6A/MIP) as the position from which to calculate the final target for the wrist using the object-centered reach offset. The F2d region is similarly modeled as two DNFs that each receive excitatory input from F2vr and tonic inhibitory input, *GP*:

$$\begin{aligned} \mathbf{IN}_{\text{F2dDIR}}(t) &= \mathbf{F2vrDIR}(t)\mathbf{W}_{\text{F2}\rightarrow\text{F2}} + \mathbf{GP}(t) + \boldsymbol{\varepsilon}_{\text{F2}} \\ \mathbf{IN}_{\text{F2dRAD}}(t) &= \mathbf{F2vrRAD}(t)\mathbf{W}_{\text{F2}\rightarrow\text{F2}} + \mathbf{GP}(t) + \boldsymbol{\varepsilon}_{\text{F2}} \end{aligned}$$

The weight matrices between preparation- and execution-related premotor populations, $\mathbf{W}_{\text{F2}\rightarrow\text{F2}}$, were not subject to learning and were set as follows:

$$\mathbf{W}(i, j) = 2\mathbf{I}$$

2.7.2 F7

While there does not appear to be direct evidence for a population of premotor neurons encoding an object-centered reach offset, there is some suggestion that such a representation does exist and may be located in the dorsal premotor cortex. The rostral portion of the dorsal premotor cortex, area F7 (approximately equal to PMdr, Wise et al. 1997), can be separated into the dorso-rostral supplementary eye field (SEF) and a lesser-known ventral region. The SEF is known to contain neurons which encode space in an object-centered reference frame (Olson and Gettner 1995), but the region is implicated in control of eye movements. While the properties of ventral F7 are not well known, it does contain neurons related to arm movements (Fujii et al. 1996, 2002), receives the same thalamic input as the arm region of F6, and receives input from the same region of the superior temporal sulcus that projects to F2vr (Rizzolatti and Luppino 2001). The ventral portion of F7 may therefore be a likely candidate for the location of population of neurons encoding reach targets in an object-centered frame of reference.

We model F7 as a preparation- and execution-related population. The preparation-related population consists of two DNFs encoding the object-centered reach offset in spherical coordinates (**F7sDIR** encoding azimuth and elevation, and **F7sRAD** encoding the radius). Inputs to the F7 module come from F2, signaling the center of the object component, and AIP, providing an affordance representation. The input to each preparation-related DNF is given by:

$$\begin{aligned} \mathbf{IN}_{\text{F7sDIR}}(t) &= \mathbf{AIP}(t)\mathbf{W}_{\text{AIP}\rightarrow\text{F7DIR}}(t) \\ &\quad + \mathbf{F2vrDIR}(t)\mathbf{W}_{\text{F2}\rightarrow\text{F7}} + \boldsymbol{\varepsilon}_{\text{F7}} \end{aligned}$$

$$\mathbf{IN}_{\text{F7sRAD}}(t) = \mathbf{AIP}(t)\mathbf{W}_{\text{AIP}\rightarrow\text{F7RAD}}(t) + \boldsymbol{\varepsilon}_{\text{F7}}$$

The execution-related population also contains two DNFs, each corresponding to one DNF in the preparation-related population. The input to each execution-related DNF is given by:

$$\begin{aligned} \mathbf{IN}_{F7eDIR}(t) &= \mathbf{F7sDIR}(t)\mathbf{W}_{F7 \rightarrow F7}(t) + \mathbf{GP}(t) + \boldsymbol{\epsilon}_{F7} \\ \mathbf{IN}_{F7eRAD}(t) &= \mathbf{F7sRAD}(t)\mathbf{W}_{F7 \rightarrow F7}(t) + \mathbf{GP}(t) + \boldsymbol{\epsilon}_{F7} \end{aligned}$$

The connection weights between F2 and AIP and the preparation-related F7 populations were initialized to small random values and subject to learning using a variant of the REINFORCE rule (Sutton and Barto 1998) which is Hebbian for positive reward values and anti-Hebbian for negative ones:

$$\begin{aligned} \mathbf{W}_{AIP \rightarrow F7DIR}(a, b, i, j, t + 1) &= \mathbf{W}_{AIP \rightarrow F7DIR}(a, b, i, j, t) \\ &+ \alpha_{F7rs}(t) (\mathbf{AIP}(a, b, t) \mathbf{F7eDIR}(i, j, t)) \\ \mathbf{W}_{AIP \rightarrow F7RAD}(a, b, i, t + 1) &= \mathbf{W}_{AIP \rightarrow F7RAD}(a, b, i, t) \\ &+ \alpha_{F7rs}(t) (\mathbf{AIP}(a, b, t) \mathbf{F7eRAD}(i, t)) \\ \mathbf{W}_{F2 \rightarrow F7}(a, b, i, j, t + 1) &= \mathbf{W}_{F2 \rightarrow F7}(a, b, i, j, t) \\ &+ \alpha_{F7rs}(t) (\mathbf{F2dDIR}(a, b, t) \mathbf{F7eDIR}(i, j, t)) \end{aligned}$$

The outputs of the execution-related populations are used as the eligibility traces since in general the object may not be visible at the end of the grasp and preparation-related cells may not be active anymore.

2.7.3 F5

Many neurons in premotor area F5 fire in association with specific types of manual action, such as precision grip, finger prehension, and whole-hand prehension (Rizzolatti and Camarda 1988) as well as tearing and holding. Some neurons in F5 discharge only during the last part of grasping; others start to fire during the phase in which the hand opens and continue to discharge during the phase when the hand closes; finally a few discharge prevalently in the phase in which the hand opens. Grasping appears therefore to be coded by the joint activity of populations of neurons, each controlling different phases of the motor act. Raos et al. (2006) found that F5 neurons selective for both grip type and wrist orientation maintained this selectivity when grasping in the dark. Simultaneous recording from F5 and F1 showed that F5 neurons were selective for grasp type and phase, while an F1 neuron might be active for different phases of different grasps (Umiltà et al. 2007). This suggests that F5 neurons encode a high-level representation of the grasp motor schema, while F1 neurons (or, at least, some of them) encode the component movements or components of a population code for muscle activity of each grasp phase.

We model F5 as a preparation- and execution-related population, each containing a one-dimensional DNF for each

VF combination with neurons in each DNF selective for maximum grasp aperture. In this module, in addition to the WTA dynamic within DNFs, every unit in a DNF laterally inhibits every other unit in the other DNFs, so that inter-DNF competition selects a VF combination, while intra-DNF competition selects a maximum aperture. The possible VF combinations are index finger pad–thumb pad (precision grasp), index+middle finger pads–thumb pad (tripod grasp), inner fingers–palm (power grasp), and thumb pad–side of index finger (side grasp). The maximal aperture is encoded as a normalized value from 0 to 1 that is transformed into target finger joint angles by the grasp motor controller (see Online Resource 1, Table 2).

Inputs to the F5 module come from the AIP module, and therefore F5 selects an appropriate grasp based on the learned affordance representation. The inputs to each preparation-related population, **F5sPREC**, **F5sTRI**, **F5sPOW**, **F5sSIDE** for the precision, tripod, power, and side grasps, respectively, are given by:

$$\begin{aligned} \mathbf{IN}_{F5sPREC}(t) &= \mathbf{AIP}(t)\mathbf{W}_{AIP \rightarrow F5sPREC}(t) \\ &- \mathbf{W}_{F5s \rightarrow F5s} \sum_i (\mathbf{F5sTRI}(i, t) \\ &+ \mathbf{F5sPOW}(i, t) + \mathbf{F5sSIDE}(i, t)) + \boldsymbol{\epsilon}_{F5} \\ \mathbf{IN}_{F5sTRI}(t) &= \mathbf{AIP}(t)\mathbf{W}_{AIP \rightarrow F5sTRI}(t) \\ &- \mathbf{W}_{F5s \rightarrow F5s} \sum_i (\mathbf{F5sPREC}(i, t) \\ &+ \mathbf{F5sPOW}(i, t) + \mathbf{F5sSIDE}(i, t)) + \boldsymbol{\epsilon}_{F5} \\ \mathbf{IN}_{F5sPOW}(t) &= \mathbf{AIP}(t)\mathbf{W}_{AIP \rightarrow F5sPOW}(t) \\ &- \mathbf{W}_{F5s \rightarrow F5s} \sum_i (\mathbf{F5sPREC}(i, t) \\ &+ \mathbf{F5sTRI}(i, t) + \mathbf{F5sSIDE}(i, t)) + \boldsymbol{\epsilon}_{F5} \\ \mathbf{IN}_{F5sSIDE}(t) &= \mathbf{AIP}(t)\mathbf{W}_{AIP \rightarrow F5sSIDE}(t) \\ &- \mathbf{W}_{F5s \rightarrow F5s} \sum_i (\mathbf{F5sPREC}(i, t) \\ &+ \mathbf{F5sTRI}(i, t) + \mathbf{F5sPOW}(i, t)) + \boldsymbol{\epsilon}_{F5} \end{aligned}$$

where $\mathbf{W}_{F5s \rightarrow F5s}$ is the inhibitory connection weight between DNFs, set to .25 in these simulations. The inputs to the execution-related populations, **F5ePREC**, **F5eTRI**, **F5ePOW**, **F5eSIDE**, are given by:

$$\begin{aligned} \mathbf{IN}_{F5ePREC}(t) &= \mathbf{F5sPREC}(t)\mathbf{W}_{F5s \rightarrow F5e}(t) + \mathbf{GP}(t) + \boldsymbol{\epsilon}_{F5} \\ \mathbf{IN}_{F5eTRI}(t) &= \mathbf{F5sTRI}(t)\mathbf{W}_{F5s \rightarrow F5e}(t) + \mathbf{GP}(t) + \boldsymbol{\epsilon}_{F5} \\ \mathbf{IN}_{F5ePOW}(t) &= \mathbf{F5sPOW}(t)\mathbf{W}_{F5s \rightarrow F5e}(t) + \mathbf{GP}(t) + \boldsymbol{\epsilon}_{F5} \\ \mathbf{IN}_{F5eSIDE}(t) &= \mathbf{F5sSIDE}(t)\mathbf{W}_{F5s \rightarrow F5e}(t) + \mathbf{GP}(t) + \boldsymbol{\epsilon}_{F5} \end{aligned}$$

The $\mathbf{W}_{F5s \rightarrow F5e}$ weights were set just as the $\mathbf{W}_{F2 \rightarrow F2}$ and $\mathbf{W}_{F7 \rightarrow F7}$ weights, and the connection weights $\mathbf{W}_{AIP \rightarrow F5sPREC}$, $\mathbf{W}_{AIP \rightarrow F5sTRI}$, $\mathbf{W}_{AIP \rightarrow F5sPOW}$, and $\mathbf{W}_{AIP \rightarrow F5sSIDE}$ were initialized to small random values and subject to learning using the

REINFORCE learning rule (as the connections to F7, Sutton and Barto 1998):

$$\mathbf{W}_{\text{AIP} \rightarrow \text{F5PREC}}(a, b, i, t + 1) = \mathbf{W}_{\text{AIP} \rightarrow \text{F5PREC}}(a, b, i, t) + \alpha_{\text{F5rs}}(t) (\mathbf{AIP}(a, b, t) \mathbf{F5ePREC}(i, t))$$

$$\mathbf{W}_{\text{AIP} \rightarrow \text{F5TRI}}(a, b, i, t + 1) = \mathbf{W}_{\text{AIP} \rightarrow \text{F5TRI}}(a, b, i, t) + \alpha_{\text{F5rs}}(t) (\mathbf{AIP}(a, b, t) \mathbf{F5eTRI}(i, t))$$

$$\mathbf{W}_{\text{AIP} \rightarrow \text{F5POW}}(a, b, i, t + 1) = \mathbf{W}_{\text{AIP} \rightarrow \text{F5POW}}(a, b, i, t) + \alpha_{\text{F5rs}}(t) (\mathbf{AIP}(a, b, t) \mathbf{F5ePOW}(i, t))$$

$$\mathbf{W}_{\text{AIP} \rightarrow \text{F5SIDE}}(a, b, i, t + 1) = \mathbf{W}_{\text{AIP} \rightarrow \text{F5SIDE}}(a, b, i, t) + \alpha_{\text{F5rs}}(t) (\mathbf{AIP}(a, b, t) \mathbf{F5eSIDE}(i, t))$$

2.7.4 Wrist rotation

Infants starting at 7 months old begin to pre-orient their hands to match an object's affordances when reaching for that object (Witherington 2005). By 9 months old, infants are skilled at hand pre-orientation and adjustment and increase reach and grasp efficiency (Morrongiello and Rocca 1989). Neurons have been described in area F2 that become active in relation to specific orientations of visual stimuli and to corresponding hand/wrist movements (Raos et al. 2004). That same paper showed that 66% of grasp neurons in F2 were highly selective for grasp type and that 72% were highly selective for wrist orientation. In addition to reach target selection, the dorsal premotor cortex is implicated in wrist movements (Riehle and Requin 1989; Kurata 1993). Raos et al. (2006) show that F5 neurons combine selectivity for grip type and wrist orientation, and that 21 out of the 38 they tested for wrist orientation selectivity showed high selectivity for a particular orientation. The most plausible hypothesis that reconciles these findings is that the dorsal premotor cortex is involved in coding reach direction and the ventral premotor cortex is involved in coding grasps, and that interconnections between F2 and F5 (Marconi et al. 2001) allow the two regions to converge on a wrist orientation appropriate for the selected reach direction and grasp type.

We model the F2/F5 wrist rotation network as a preparation- and an execution-related population, similarly to the other premotor modules. The signal-related population contains a three-dimensional DNF, with each unit selective for a combination of the angles of the DOFs of the wrist within its joint angle limits (see Online Resource 1, Table 1). The F2/F5 module receives inputs from the AIP, F5, and F7 module and therefore selects the wrist orientation based on the affordance representation and selected grasp and reach offset. The input to the three-dimensional preparation-related DNF is given by:

$$\begin{aligned} \mathbf{IN}_{\text{WRs}}(t) &= \mathbf{AIP}(t) \mathbf{W}_{\text{AIP} \rightarrow \text{WR}}(t) + \mathbf{F7sDIR}(t) \\ &\times \mathbf{W}_{\text{F7} \rightarrow \text{WR}}(t) + \mathbf{F5sPREC}(t) \mathbf{W}_{\text{F5PREC} \rightarrow \text{WR}}(t) \\ &+ \mathbf{F5sTRI}(t) \mathbf{W}_{\text{F5TRI} \rightarrow \text{WR}}(t) + \mathbf{F5sPOW}(t) \\ &\times \mathbf{W}_{\text{F5POW} \rightarrow \text{WR}}(t) + \mathbf{F5sSIDE}(t) \mathbf{W}_{\text{F5SIDE} \rightarrow \text{WR}}(t) \\ &+ \epsilon_{\text{WR}} \end{aligned}$$

Note that the preparation-related F2/F5 DNF receives input from the *F7sDIR* population encoding the direction of the reach offset, but it does not get reach offset radius information from the *F7sRAD* population since the wrist rotation should not depend on the offset radius. The execution-related population also contains a three-dimensional DNF, with its input given by:

$$\mathbf{IN}_{\text{WRc}}(t) = \mathbf{WRs}(t) \mathbf{W}_{\text{WR} \rightarrow \text{WR}} + \mathbf{GP}(t) + \epsilon_{\text{WR}}$$

As with F5 and F7, the weights of the afferent connections of the preparation-related populations were updated using the REINFORCE learning rule (Sutton and Barto 1998):

$$\mathbf{W}_{\text{AIP} \rightarrow \text{WR}}(a, b, i, j, k, t + 1) = \mathbf{W}_{\text{AIP} \rightarrow \text{WR}}(a, b, i, j, k, t) + \alpha_{\text{WRrs}}(t) (\mathbf{AIP}(a, b, t) \mathbf{WRc}(i, j, k, t))$$

$$\mathbf{W}_{\text{F7} \rightarrow \text{WR}}(a, b, i, j, k, t + 1) = \mathbf{W}_{\text{AIP} \rightarrow \text{F7}}(a, b, i, j, k, t) + \alpha_{\text{WRrs}}(t) (\mathbf{F7eDIR}(a, b, t) \mathbf{WRc}(i, j, k, t))$$

$$\mathbf{W}_{\text{F5PREC} \rightarrow \text{WR}}(a, i, j, k, t + 1) = \mathbf{W}_{\text{F5PREC} \rightarrow \text{WR}}(a, i, j, k, t) + \alpha_{\text{WRrs}}(t) (\mathbf{F5ePREC}(a, t) \mathbf{WRc}(i, j, k, t))$$

$$\mathbf{W}_{\text{F5TRI} \rightarrow \text{WR}}(a, i, j, k, t + 1) = \mathbf{W}_{\text{F5TRI} \rightarrow \text{WR}}(a, i, j, k, t) + \alpha_{\text{WRrs}}(t) (\mathbf{F5eTRI}(a, t) \mathbf{WRc}(i, j, k, t))$$

$$\mathbf{W}_{\text{F5POW} \rightarrow \text{WR}}(a, i, j, k, t + 1) = \mathbf{W}_{\text{F5POW} \rightarrow \text{WR}}(a, i, j, k, t) + \alpha_{\text{WRrs}}(t) (\mathbf{F5ePOW}(a, t) \mathbf{WRc}(i, j, k, t))$$

$$\mathbf{W}_{\text{F5SIDE} \rightarrow \text{WR}}(a, i, j, k, t + 1) = \mathbf{W}_{\text{F5SIDE} \rightarrow \text{WR}}(a, i, j, k, t) + \alpha_{\text{WRrs}}(t) (\mathbf{F5eSIDE}(a, t) \mathbf{WRc}(i, j, k, t))$$

2.8 Training

Each training trial was run for 5 s with a 1-ms time step. All training simulations used the same protocol in which at 0.5 s a green object appeared in the model's field of view. Different objects were presented (cubes, rectangular prisms, cylinders, spheres, and flat plates) at random orientations and locations. The model input was obtained by getting the object's shape, color, size, orientation, and position from the physics simulator. At 1 s into the simulation the object turned red, triggering the release of inhibition from the execution-related premotor populations by setting the tonic inhibitory input *GP* to each of these populations to 0. In the last five time steps of each trial (corresponding to 5 ms of simulation time), *rs(t)* is set to $\text{DA}_{\text{success}}$ if the grasp is successful, and DA_{fail} if not. Although the bulk of learning in this model is

done simultaneously in all layers, we still found it necessary to use a somewhat staged learning approach to bootstrap the system. During AIP pretraining trials, no movements were attempted and $rs(t)$ was always equal to 0. During the wrist rotation pretraining trials, $rs(t)$ was set to $\frac{DA_{\text{success}}}{4}$ if palm contact was achieved at all.

2.8.1 AIP pretraining

We found that at the start of training, the neurons in AIP did not have sufficient activity rates to drive the premotor populations. This resulted in approximately 1000 trials in which no grasps were attempted, but the connection weights into AIP were slowly modified according to the SOM weight adjustment rule described above. Eventually large numbers of AIP units were significantly activated in widespread overlapping representations sufficient to activate premotor populations. This period may correspond to a period of visual experience before the development of skilled reaching in which no grasps are attempted and visual regions are shaped by unsupervised learning mechanisms.

Although we assume the existence of a pretrained reaching circuit, this period could correspond to a period of motor babbling in which internal models of the arm are learned (Bullock et al. 1993). This period of training may also correspond to the motor babbling period of infant development modeled by Kuperstein (1988).

2.8.2 Wrist rotation pretraining

The added realism of our simulator compared to that of ILGM comes at the price of a much lower probability of successfully grasping an object with random motor parameters. This makes the task of learning much more difficult. To surmount this problem, we used a period of pretraining in which any palm contact was rewarded and only the connection weights between F7 and F2/F5 were modified. At the end of pretraining, the system could at least orient the hand in the correct direction to make finger or palm contact with the object at various locations and orientations (Fig. 7), similar to the automatic wrist orienting mechanism used in some ILGM simulations. This period of training therefore corresponds to infant development from 7 to 9 months where infants learn to pre-orient their wrist in response to an object's affordances (Morrongiello and Rocca 1989; Witherington 2005).

2.8.3 Grasp training

After pretraining, the system was trained for 10,000 trials with different objects at random locations and orientations. The object type, position, and orientation were changed every 6 trials since each trial lasted 5 s and infants will repeatedly reach to a novel object for at least 30 s before habituation (Von

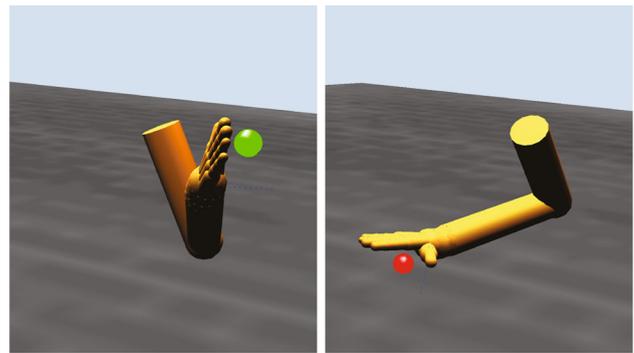


Fig. 7 Examples of two successful reaches in the wrist rotation pretraining stage. Note the rotation of the wrist which orients the palm toward the object

Hofsten and Spelke 1985). During these training trials, only stable grasps were positively reinforced and all modifiable connection weights (dashed arrows in Fig. 4) were subject to learning.

3 Results

After training, the model was able to generate stable grasps of each object type at various locations and orientations. The representation in AIP allowed the model to generalize across object properties enough to successfully grasp objects in novel configurations in a few attempts. Here, we demonstrate the ability of the model to generate successful grasps, analyze the learned representations in AIP, and generate predictions for future experiments.

3.1 AIP representation

Each object at each location and orientation in the training set elicited a slightly different activation pattern in the cIPS and V6A/MIP populations. However after pretraining AIP, objects with similar features elicited similar, overlapping, patterns of activation in this region (Fig. 8). This is an inherent property of SOMs and is what allows the model to successfully grasp novel objects in familiar locations and orientations (see Grasp training, below).

Murata et al. (2000) tested the response of AIP neurons to the sight of various types of objects and found that a high proportion of visually responsive AIP neurons were highly shape selective, responding strongly to one particular object shape and weakly to all others. They used multidimensional scaling (MDS) to look at how moderately object-selective neurons encode the similarity of objects. MDS is a method of reducing a high-dimensional input space into a lower-dimensional space while preserving topological relations between vectors. It was found that moderately object-selective visually

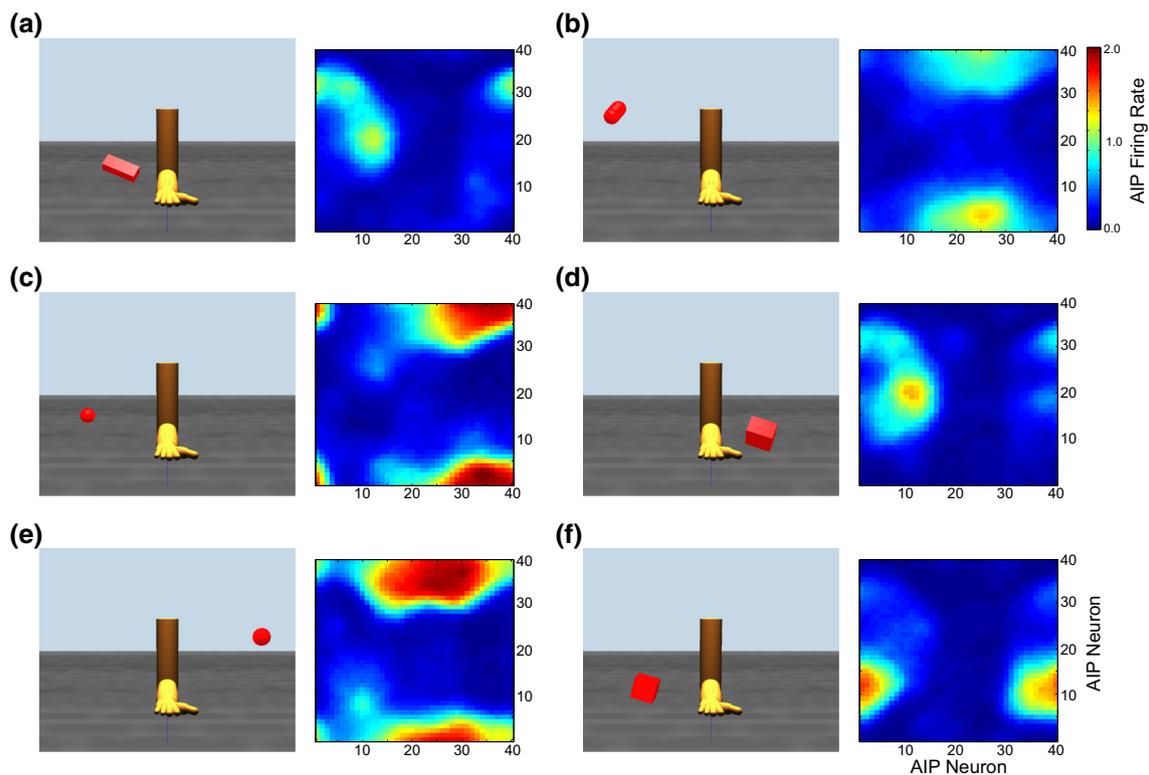


Fig. 8 AIP activation after pretraining for different objects in various locations and orientations. Each panel shows a third-person view (left) and AIP activation (right)

responsive AIP neurons respond to common combinations of geometric features shared by similar objects such as shape, size, and/or orientation. However, note that the objects used in this experiment were not natural, complex objects, but geometric primitives. We suggest that AIP codes combinations of features of object components and that given a complex object, AIP neurons selective for features of each of its components will be activated.

We did not include any explicit encoding of object shape in the inputs to AIP (although the CYL and RECT populations selectively respond to features of particularly shaped objects). However, we found that after training AIP contained a mixture of highly, moderately, and weakly shape-selective neurons (Fig. 9). To characterize a neuron's object preference, we used the same technique as Raos et al. (2006) where the object specificity of a neuron is defined as a preference index (PI) :

$$PI = \frac{n - \left(\frac{\sum r_i}{r_{\text{pref}}} \right)}{n - 1}$$

where n is the number of object shapes tested, r_i is the mean activity of the neuron for objects with shape i , and r_{pref} is the mean activity for the preferred object shape in the current epoch. This measure can range from 0 to 1, with 0 meaning

the neuron responds equally to all object shapes and 1 indicating activity for only one shape. We classified neurons as highly shape selective if they had a PI greater than .75, moderately selective if they had a PI between .25 and .75, and non-shape selective if they had a PI less than .25. The PI of each neuron was evaluated in blocks of 500 trials throughout the entire training period.

We found that at the start of training, all AIP neurons were moderately shape selective (Fig. 9b), responding to combinations of features such as size and orientation. After about 500 trials, a small amount of neurons became non-shape selective and simply responded to the presence of any object shape. By the end of the AIP pretraining period, nearly 80% of AIP neurons were highly shape selective, some reaching a maximum PI near 1.0, indicating they only responded to specific combinations of features that signaled a particular object shape. This selectivity was maintained throughout the entire training period, even after the model began to attempt grasps.

At the end of training, we computed the PI for each AIP neuron for object shape, shoulder-centered direction (φ_s , θ_s), orientation (o_x , o_y , o_z), and size (o_x , o_y , o_z). The PI for each AIP neuron for each object feature is shown in Fig. 10. A significant portion of AIP neurons were moderately selective for object shape, position, orientation, and size. Each AIP neuron responded to a combination of these object

Fig. 9 **a** Object shape specificity statistics for the AIP population during training (*solid* maximum PI, *dashed* mean PI, *dotted* minimum PI). **b** Numbers of highly (*solid*), moderately (*dashed*), and non- (*dotted*) shape-specific neurons throughout training

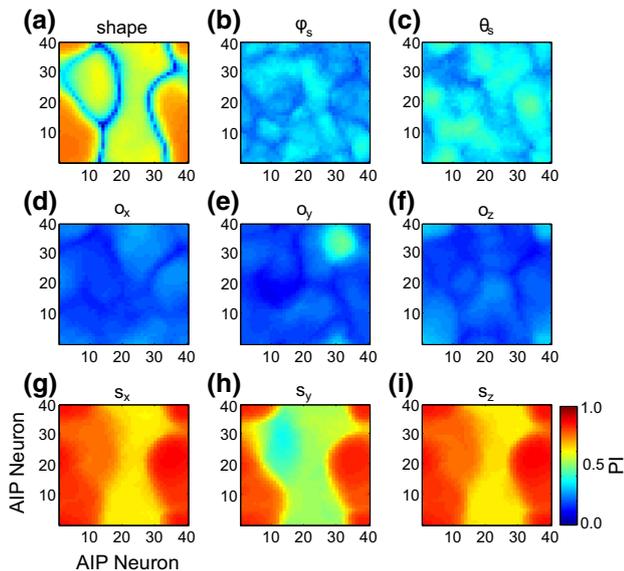
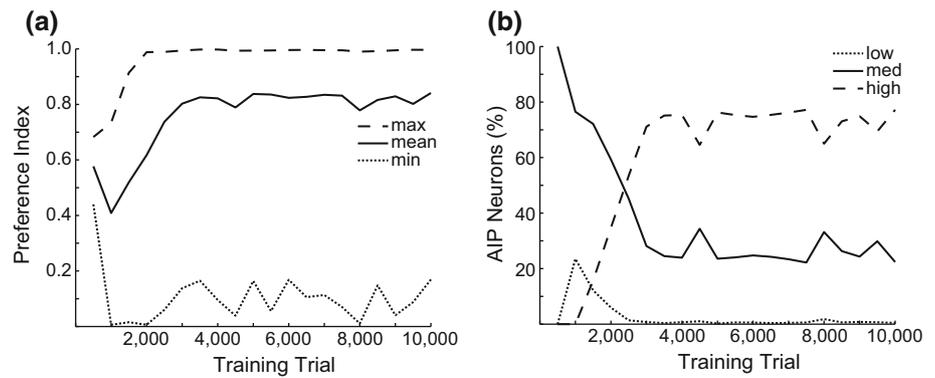


Fig. 10 Preference index (PI) for each AIP neuron for different object features after training. **a** Shape, **b** shoulder-centered direction ϕ_s , **c** shoulder-centered direction θ_s , **d** orientation, o_x , **e** orientation, o_y , **f** orientation o_z , **g** size, s_x , **h** size, s_y , **i** size, s_z

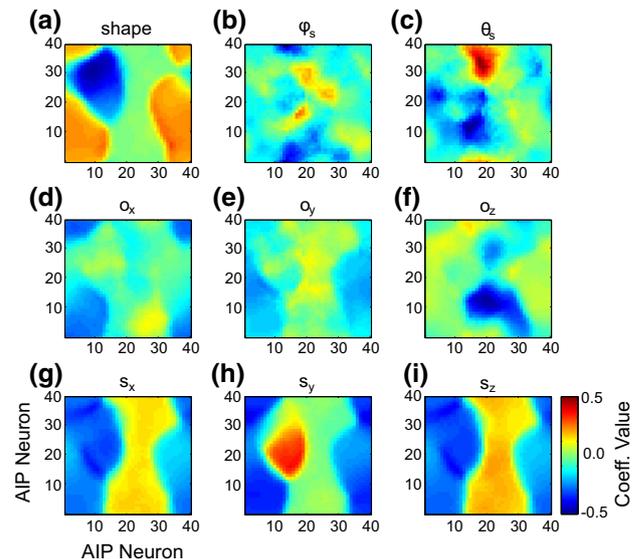


Fig. 11 Value of each fitted linear coefficient for each AIP neuron after training. **a** Shape, **b** shoulder-centered direction ϕ_s , **c** shoulder-centered direction θ_s , **d** orientation, o_x , **e** orientation, o_y , **f** orientation o_z , **g** size, s_x , **h** size, s_y , **i** size, s_z

features. We fitted the activity of each AIP neuron to a linear model using the z-scored value of each object feature as independent variables. The coefficients for each variable for all AIP neurons are shown in Fig. 11. While many neurons are highly selective for the object shape, neurons within those populations preferentially respond to a combination of different orientations, positions, and sizes.

We tested the ability of AIP representations to encode different object orientations and positions by presenting the model with a cylinder of a given size at various orientations and positions (Fig. 12). In all trials, the cylinder elicited a similar pattern of activation in AIP, but certain cells were modulated by the orientation and/or position of the object. The AIP module therefore represented not only the grasp affordance represented by the object (a power grasp in this case), but also the metric information needed to parameterize the grasp according to the object’s position and orientation.

3.2 Grasp training

For each object shape and configuration, the model replicated the results of the ILGM, which used a simplified AIP model that signaled the presence, location, or orientation of an object in different simulations. As in ILGM, the learned connection weights between the AIP module and the premotor populations encoded grasp parameters most likely to result in the performance of a stable grasp. The model was able to generate stable grasps of each object tested in various positions with different orientations (Figs. 13, 15, see Online Resources 2–6 for videos of sample grasps), successfully grasping the presented object in nearly 70% of the trials after 5000 training trials (Fig. 14). Halfway through the training period, we introduced a set of novel object trials, presenting the model with objects having the same shapes as objects it had seen before, but in new positions, orientations, and having different sizes. Performance initially decreased

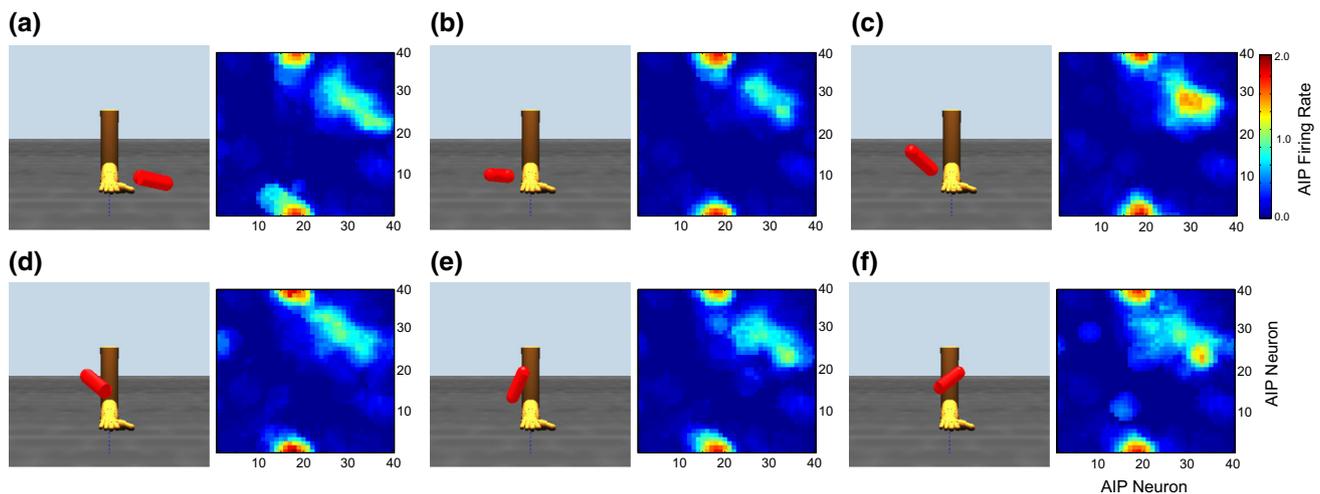


Fig. 12 Response of each AIP neuron after training to a cylinder of the same size presented at different orientations and positions

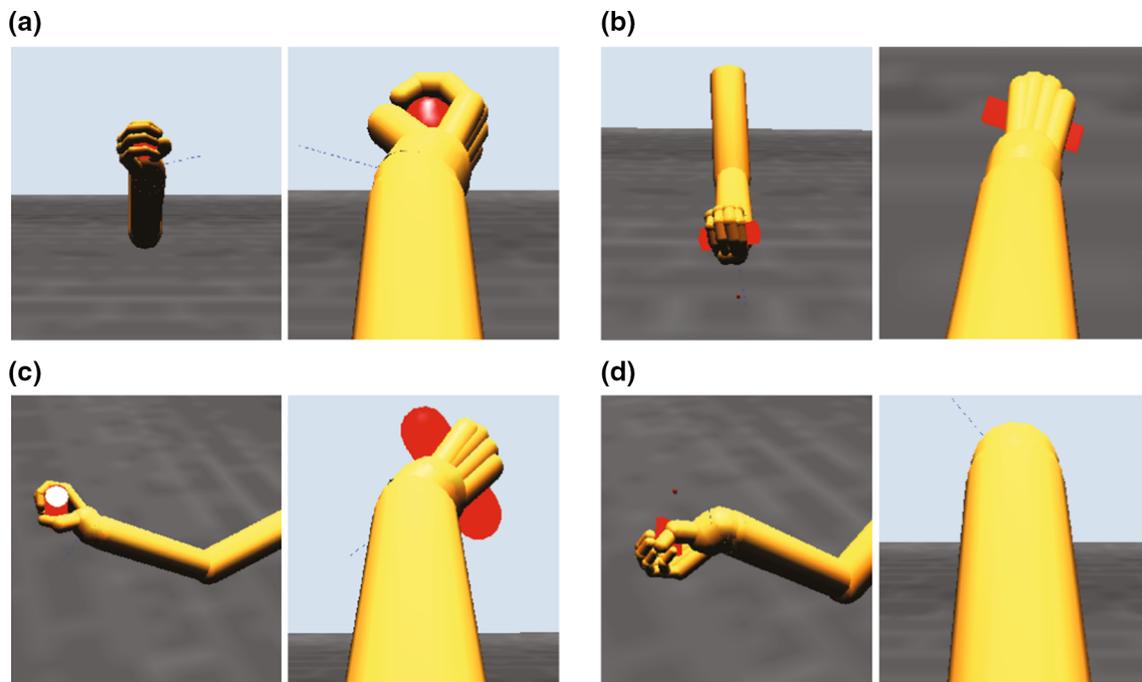


Fig. 13 Stable power grasps generated by the model of different objects with various positions and orientations. *Each panel* shows a third-person view on the *left* and the model's first-person view on the *right*. **a** Sphere, **b** rectangular prism, **c** cylinder, **d** plate

to approximately 50 % at 6000 trials, but then increased to nearly 80 % after 8000 trials.

ILGM showed that even without hand preshaping and virtual finger enclosure, precision pinches could result from trial-and-error reinforcement learning, although not often. We have replicated this result, but because the more realistic physics simulator we used is much less forgiving in evaluating grasp stability, we had to build the ability to pre-shape the hand for a precision pinch into the model. In spite of this limitation, we have shown that it is possible to learn appropriate affordance representations and motor parameters

for precision pinches without feedback-based control. We therefore believe that the development of precision pinching coincides with the development of feedback-based, skilled grasping (see Sect. 4). A sequence of frames from a precision pinch generated by the model during training is shown in Fig. 15 (see Online Resource 2 for the video). Hand preshaping and wrist rotation begins shortly after object presentation (Fig. 15c), and the enclosure phase is triggered once the grasp reaches its maximal aperture (Fig. 15d). The forefinger first makes contact with the object (Fig. 15f), and the thumb contacts the other side of the object's surface (Fig. 15g).

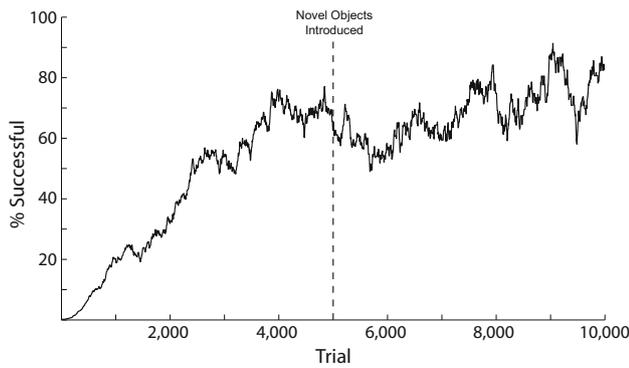


Fig. 14 Percentage of trials resulting in a stable grasp during training

The opposition axis between the thumb and forefinger was aligned close enough to stabilize the object, aided by friction, and the resulting grasp is judged as stable (Fig. 15h).

As demonstrated above, the learned AIP representation preserved those features represented in cIPS and V6A/MIP that are essential for programming the grasp, allowing the model to successfully grasp objects in various orientations and positions. Activity in cIPS, AIP, and premotor cortex is shown in Fig. 16 during two grasps of the same size cylinder at the same location (sometimes the object is displaced by the hand after grasping), but with different orientations. The cIPS AOS cylinder population encodes the three-dimensional orientation of the main axis of the cylinder as a two-dimensional population code. Based on this representation and those in the other cIPS populations and V6A/MIP, the AIP module forms a distributed representation of combinations of the object’s features that are important for grasping. AIP neurons selective for the object shape and/or position are active during each grasp trial, resulting in highly similar patterns of activity in AIP. However, there are some AIP neurons selective for the orientation of the object, and this causes the patterns of

AIP activity to be slightly different depending on the object’s orientation in each trial. These neurons bias motor parameter selection in the premotor populations, resulting in the selection of parameters appropriate for the current object’s orientation. The differences in F2–F5 activity encode the different wrist rotations that must be used to successfully grasp the cylinder at each orientation.

In this model, F5 neurons encode grasp types, since their activity is decoded in order to perform the grasp. In contrast, AIP neurons come to represent affordances—combinations of object features that signal the possibility for grasping. A comparison of AIP and F5 activity while grasping the same object with different types of grasps is shown in Fig. 17. Each row shows AIP and F5 activity while grasping the same object with three different types of grasps. The AIP representation is almost exactly the same during each grasp, but the F5 activation patterns are completely different. AIP neurons in this model therefore do not specify the type of grasp, but represent an affordance that can be acted on using several types of grasps.

Because of its learned affordance representation, ILGA can generalize learned grasp plans to novel objects. After training, the model was presented with a novel object at a similar orientation and position to those in the training set. Figure 18a shows the input to each premotor population during the first trial with the novel object, and Fig. 18b shows the resulting premotor population activity in that trial. The circles denote regions of grasp parameter space where the model was biased toward selecting grasp parameters that had proved successful with similar objects during training. These parameters resulted in successful grasps of the novel object in some proportion of trials, but after 35 trials, a better strategy was found, demonstrated by the shifts in the peaks of input activity to each premotor population in Fig. 18c.

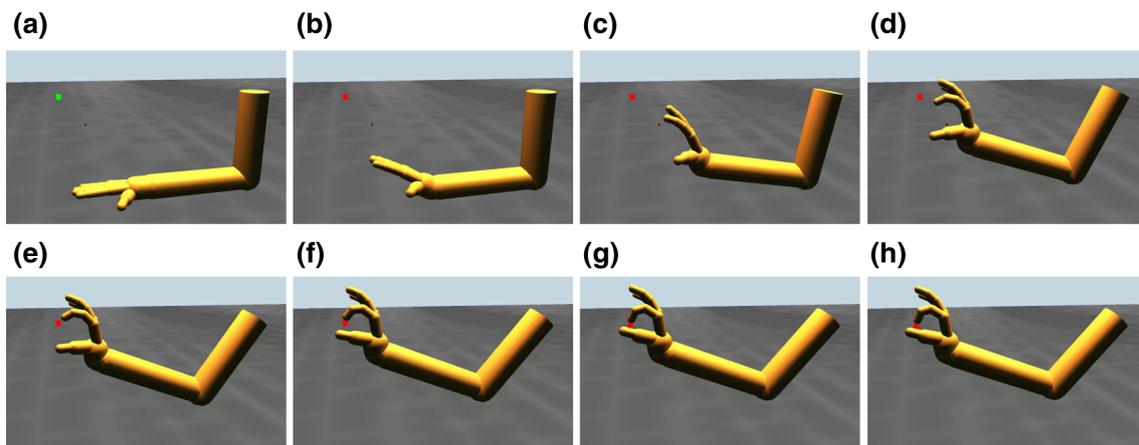


Fig. 15 A series of frames showing the progression of a precision pinch of a small cube generated by the model. After the go signal (b), hand preshaping and wrist rotation has begun (c). The enclose phase is triggered in d and the object is first contacted in f

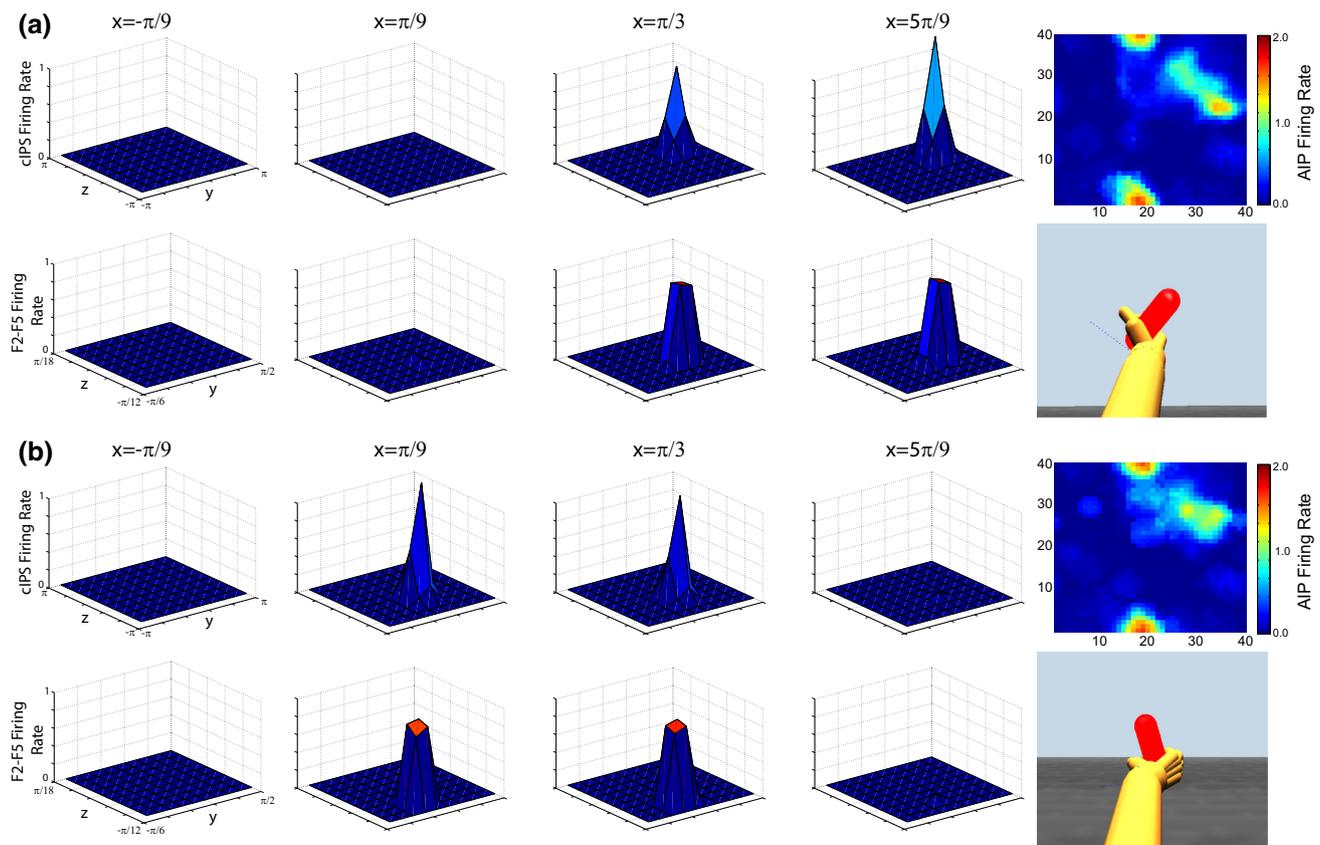


Fig. 16 Top row of each panel shows (from left to right) firing rates of slices of the cIPS AOS cylinder population at $x = -\pi/9, \pi/9, \pi/3, 5\pi/9$ and the AIP population. The bottom row of each panel shows (from left to right) firing rates of slices of the F2/F5 population at $x = -\pi/9, \pi/9, \pi/3, 5\pi/9$ and the resulting grasp. **a**

Network activity and resulting grasp in response to a cylinder rotated 60 degrees about the x axis. **b** Network activity and resulting grasp in response to a cylinder rotated 120 degrees about the x -axis. Activity in other populations was not significantly different during each grasp and is therefore not shown

4 Discussion

The main computational ingredients in ILGA and their contributions to the main results of our simulations are summarized in Table 1. We have shown not only that the current model explains the development of F5 canonical neurons controlling grasping, as did ILGM, but that it also gives an account of the development of visual neurons in area AIP. In our simulations, highly shape-specific neurons developed in this region as the result of unsupervised learning before grasps were even attempted. This specificity was maintained even as grasps were performed. The result was a mixture of AIP neurons that only fire for particular shapes and those that are activated by combinations of features representing grasp affordances. This allows the model to respond to novel objects in positions and orientations similar to ones it has already successfully grasped.

Parietal representations in both ILGA and its precursor, ILGM, use population codes to represent object features, but ILGA represents multiple object features at once and

combines them into a representation of a grasp affordance. ILGM's premotor module uses a probabilistic coding followed by a rewriting of activity as a population code, while ILGA uses a more realistic noisy WTA process. In ILGM the wrist rotation, object-centered reach offset and hand enclosure rate are selected by the premotor module, but ILGA comes closer to FARS in including the grasp type and maximum aperture in addition to wrist rotation and object-centered reach offset.

Oztop et al. (2006) presented a model of AIP related to ILGM and ILGA. Like ILGA, this model used a SOM to model visual-dominant AIP neurons. The basic result of this model was the demonstration of how AIP could extract higher-level information from simpler visual inputs and map them onto hand postures resulting in stable grasps. In this sense, the results of ILGA are similar, showing how AIP can combine simpler visual input into higher-level affordance representations and map them onto motor parameters that will result in stable grasps. However, Oztop et al. (2006) used backpropagation, a biologically implausible learning rule, to

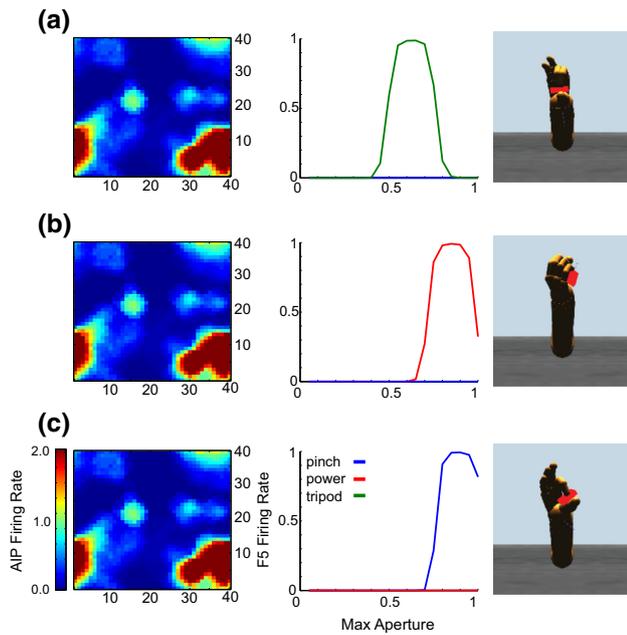


Fig. 17 Firing rates of neurons in AIP (left column) and F5 (middle column) while grasping a plate with a tripod grasp (a), power grasp (b), and precision pinch (c)

shape the connections between AIP and F5, while ILGA uses reinforcement learning. The inputs and outputs to Oztop et al.’s (2006) model are also less biologically plausible than in ILGA. In their model, the input to AIP was a depth map and the F5 representation was set of hand joint angles. This is inconsistent with the neurophysiological data showing representation of object geometric properties in cIPS (Sakata et al. 1998) and its projection to AIP (Nakamura et al. 2001), as

well as data showing that F5 neurons are tuned to a particular grasp rather than the specific posture of the hand at any point during the grasp (Umiltà et al. 2007).

4.1 Skilled grasping

Normally, grasping is controlled in a feedforward manner (Santello et al. 2002; Winges et al. 2003); however, several studies have found online feedback-based corrections of the hand shape relative to the object in the later part of the movement (Schettino et al. 2003; Nataraj et al. 2014) especially during object perturbations (Gritsenko et al. 2009) and as a compensatory strategy after stroke (Raghavan et al. 2010). This indicates that humans are capable of both types of grasp control and can switch between them based on task demands and availability of sensory information (Hoff and Arbib 1993; Gritsenko et al. 2009). Indeed, transient inactivation of human area AIP using TMS causes deficits in error correction during online control of grasping (Tunik et al. 2005). While the current model accounts for the development of visually dominant AIP neurons, it does not include motor and visuomotor neurons. It is thought that these neurons complete a feedback loop between AIP and area F5 that is used to perform feedback-based control of grasping and other manual actions. In order for this model of open loop control of infant grasping to include feedback control to guide the fingertips to the object’s surface, two extensions are required. The first is a region to represent patches of an object’s surface to serve as targets to bring the fingertips to, and the second is an inverse kinematics model of the hand and wrist to bring the desired virtual fingers to these targets. Introduction of these features to the model would both increase the granularity of

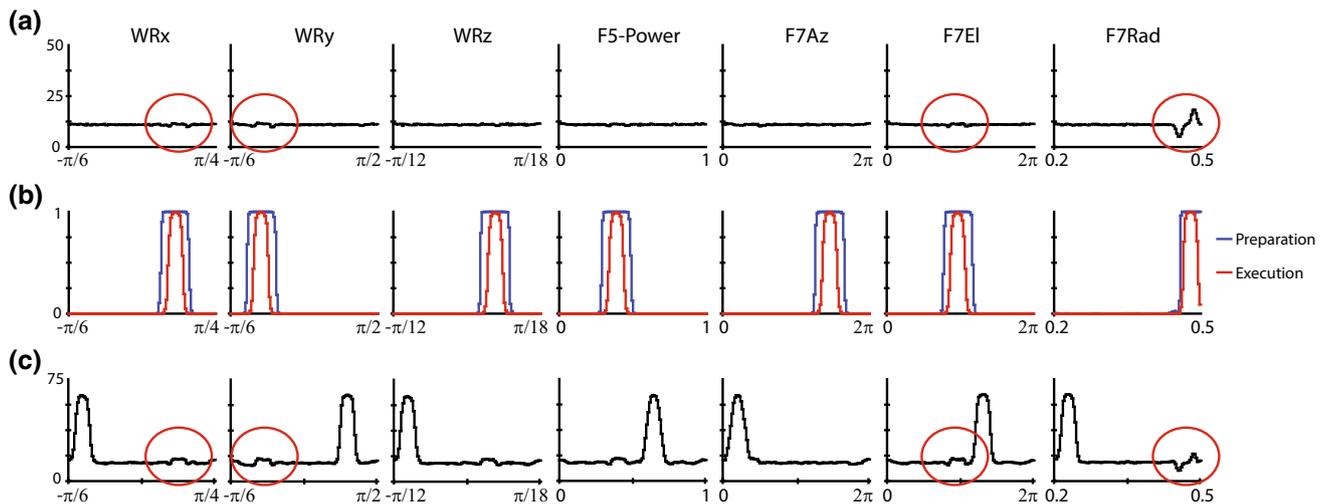


Fig. 18 An example of the generalization abilities of ILGA. **a** Input to each premotor population during the first trial with a novel object. The circles highlight biases based on previous successful grasps of similar

objects. **b** Premotor signal and execution subpopulation activity on the first trial with the novel object. **c** Input to each premotor population after 35 trials with the novel object

Table 1 Summary of ILGA components and results

Component	Result
Object feature population codes	Aids generalization between objects with slightly different features (i.e., objects of a similar size or orientation)
Self-organizing map (SOM)	Aids generalization between objects with similar combinations of features and novel objects with feature combinations similar to those seen before
Modulation of SOM learning rate by reinforcement	Causes SOM to preferentially represent feature combinations of successfully grasped objects
Dynamic neural fields (DNFs)	Allow selection of motor parameters based on input strength
Noise in DNF activity	Promotes exploration of parameter space
Learned connection weights between AIP and premotor DNFs	Bias the selection of reach and grasp parameters to successfully grasp object
Learned connection weights between premotor DNFs	Bias the selection of reach and grasp parameters that depend on values of other parameters (i.e., the appropriate wrist rotation depends on the chosen reach offset)
Fixed connection weights between V6a/MIP and premotor	Allow accurate reaches. Future versions could change with learning to simulate reach learning
Dynamic movement primitives (DMPs)	Implement reach trajectory planning in a way that is easily extended to handle more complex trajectories

object representation and improve fine control of the hand, expanding the range of possible grasp affordances that can be extracted and acted upon.

Since the orientation of the surface patches matters in programming a grasp, a candidate region for the object surface patch representation would be cIPS since surface-orientation-selective cells have been found there (Sakata et al. 1997). To make the target surface representation invariant to object location, it may be represented in an object-centered reference frame. Although such an organization has not been reported in cIPS, this may be due to a lack of experiments using eye tracking to control for gaze position with respect to the center of the object. A potential problem with this idea is that cIPS neurons have only been found to respond to visible surfaces, whereas fingertips may potentially contact a surface on the opposite side of an object during a grasp. The location of potential target surfaces on the opposite side of a visible object must therefore be inferred from the 2 1/2-dimensional sketch of the object provided by cIPS.

The current model uses an inverse kinematics model of the arm to convert target wrist locations into target joint angles for the shoulder and elbow. In order to perform more precise and dextrous grasp and manipulation tasks, the extended model must include an inverse kinematics model of the arm and hand that can convert target locations for any combination of virtual fingers into target joint angles for the shoulder, elbow, wrist, thumb, and finger joints. ILGA uses the pseudo-inverse of the Jacobian matrix for computing inverse reach kinematics, requiring a 3×4 matrix (4 controlled degrees of freedom, 3-dimensional wrist position). If the extended model used the same method, this would require multiple

Jacobian matrices, one for each combination of virtual fingers and mapping to real fingers. Assuming that only two fingers will contact the object, each Jacobian would therefore be a 6×22 matrix (22 controlled DOFs, and two 3-dimensional virtual finger positions). Since multiple combinations of virtual fingers, mappings to real fingers, and contact points are possible, this method may not be tractable.

A model like ILGA could provide the scaffolding for such a model by generating a range of stable grasps that can be used to learn the prediction of unseen target surfaces. Given a stable grasp, the haptic feedback from finger and hand contacts with the object can be used to calculate an object-centered representation of the location of these contact points using the location of the object (represented in V6A/MIP) and the posture of the hand and arm. These contact point locations could then be used as a training signal for a network that predicts the location of target surfaces given the representation of visible surfaces in cIPS. Such a model would then develop AIP cells that are responsive to the affordances characterized by these combinations of surface patches. Computing the positions of hand–object contact points given the arm/hand posture requires a *forward* kinematics model of the entire arm and hand including shoulder, elbow, wrist, thumb, and finger joints.

The extended model therefore requires a set of inverse/forward model pairs for the entire arm/hand. The models could be learned using multiple model-based reinforcement learning (MMRL, Doya et al. 2002) during the ILGA training period. MMRL is a method of using reinforcement learning and predictor error to learn multiple inverse/forward model pairs, and has been formulated in

both discrete and continuous time and state cases. This architecture is composed of multiple modules containing inverse/forward model pairs that compete to learn new tasks, with certain modules becoming specialized for different tasks with repeated training. These models could be learned offline while ILGA training progresses, until the prediction errors of the forward models become small enough to allow their associated inverse models to control the arm and hand. At this point, the forward model pairs would continue to compete for control of the arm and hand and become specialized for controlling different types of grasps.

4.2 Context-dependent grasps

At least two studies show that infants 1–2 years old selectively modify their actions based on future planned actions. [McCarty et al. \(1999\)](#) demonstrated that 9- and 14-month-old infants grasp a spoon with their preferred hand, regardless of its orientation and whether the next action was to bring the spoon to the mouth or another location. At 19 months, infants had learned to coordinate hand selection with the action goal and the spoon's orientation in order to facilitate the smooth execution of the next action. [Claxton et al. \(2003\)](#) measured the arm kinematics of 10-month-old infants reaching for a ball and then either throwing it or fitting it down a tube. They found that the reach to the ball was faster if they then intend to throw it. Both of these studies suggest that infants pre-plan segments of compound actions at some level. This point brings to light a shortcoming in both the ILGM and the ILGA models. Both of these models use some evaluation of the stability of the grasp as the metric for reinforcement. A more realistic model would use success in a subsequent action with the object such as throwing or placing as the grasp reinforcement criterion. This would require some representation of the task or goal with the ability to bias selection of grasp targets and affordance representations to perform grasps appropriate for the planned action.

In the original FARS model, working memory and task-specific associations in prefrontal cortex bias grasp execution by modulating the grasp selection in F5 such that those grasps appropriate for the current task are selected. Using synthetic brain imaging, a method to compare global model activity with PET or fMRI data, [Arbib et al. \(2002\)](#) showed that a projection from PFC to AIP rather than F5 better explained human PET data. A projection from PFC to AIP had not been reported at that time, but this prediction was later validated both anatomically ([Borra et al. 2007](#)) and neurophysiologically ([Baumann et al. 2009](#)). This could be included in ILGA by the addition of a prefrontal cortex module encoding the task context with projections to the AIP module that are also modifiable through reinforcement. Thus, task representations in prefrontal cortex would presumably become associated with the affordances and actions that lead to reward. Such

a simulation could shed insight into the interplay between cognitive and motor development by examining the operation of the model without or without a pretrained motor system or prefrontal cortex. It could be that associative signals from an already trained prefrontal cortex could interfere with the normal development of the parieto-premotor connection weights.

The issue of context-dependent grasps also requires a more sophisticated mechanism for representing multiple affordances. We suggest that the dorsal visual stream analyzes the affordances of the components of complex objects, but in general, objects are not made of spheres and rectangular solids. Rather, they may have diverse shapes, yet the developing brain learns to embed an opposition space in them that affords a stable grasp. There is additionally a binding problem here if there are multiple affordances, which may require encoding a “focus of attention” to give an approximate localization of the affordance with learning refining this to the localization of the grasp axis to guide the grasp-plan-specific guidance of reaching.

4.3 Predictions

Using biologically plausible learning rules and inputs, we have shown that ILGA can learn to represent affordances for grasping and to select motor parameters appropriate to act on them. This model makes several testable predictions concerning (a) the encoding of object features in area AIP, (b) shifts in AIP activation during learning, and (c) the existence of an object-centered spatial representation for reach-to-grasp movements.

ILGA reproduces experimental data showing that many AIP neurons are moderately object shape selective, showing responses to multiple objects ([Murata et al. 2000](#)). Since ILGA is a developmental model, it allows us to go beyond available experimental data and predict that AIP neurons are initially selective for various object features and become shape selective early in development, before grasping has developed. In ILGA, the global reinforcement signal elicited by successful grasps modulates the rate of unsupervised learning occurring in AIP. This causes neurons in AIP to preferentially encode features of objects that can be successfully grasped, resulting in a representation of grasp affordances rather than strictly geometric features. To our knowledge, no studies have looked for shifts in AIP activation during the course of grasp learning. This model predicts that if grasps of certain objects are disrupted (through local muscimol injection or physical perturbation), AIP cell selectivity will shift, with more cortical representation eventually given to the features of other objects that are successfully grasped. For example, the selectivity of any AIP neurons that prefer round, elongated objects should shift over many trials if grasps of

cylinders are repeatedly disrupted by spatial perturbing the target object.

While we are not aware of any data showing the existence of an object-centered reference frame for reaching, we found it necessary to use such a representation in order to plan the direction of the hand's approach to the object. One possible reason that such a representation has not been found is that most experiments use either a pure reaching, wrist rotation, or naturalistic grasping task. Once ILGA has been trained, the reach offset direction is highly correlated with the wrist rotation so that the hand will approach the object with the correct orientation for grasping. Therefore, selectivity for object-centered offset directions cannot be experimentally demonstrated without trials in which the offset direction is held constant while the wrist rotation is varied. This is similar to the situation in the interpretation of motor cortex activity, where it has been shown that intrinsic and extrinsic and kinematic and kinetic variables are highly correlated during a commonly used experimental reaching task (Chan and Moran 2006). In order to determine whether a region encodes the object-centered reach offset independent of wrist rotation, a reaching task must be used in which the subject must reach to a target object from different directions with varying wrist orientations. On the basis of its object-centered representation for saccades, and arm-related activity in its ventral portion, we predict an object-centered spatial representation in ventral F7.

4.4 Related models

Most related grasping models focus on learning inverse kinematics for the hand, selecting contact points on the object's surface, and developing feedback-based control of the hand. Most models plan the grasp in terms of kinematics, but at least one model stresses control of grasp forces. While many models use trial-and-error learning, they are not developmental models like ILGM and ILGA in the sense that they begin learning at a phase corresponding to grasping development in infants 9 months and older. Some models are based on neurophysiological data, but many are built purely using machine learning and robotics techniques.

Two models that stress learning inverse kinematics transformations for the hand are those of Molina-Vilaplana et al. (2007), Rezzoug and Gorce (2003) and Gorce and Rezzoug (2004). Molina-Vilaplana et al.'s (2007) model first learns the inverse kinematics functions of the fingers and then learns to associate object properties with grasp postures (a function they relate to AIP/F5 functionality). The model first learns inverse kinematics for the thumb, index, and middle fingers so that it knows the relationships between finger motor commands and their sensory consequences (proprioceptive and visual), and then learns to associate object features with grasp postures with a local network called GRASP. The input to

GRASP is a 7-dimensional vector encoding object shape (cube, sphere, or cylinder), object dimensions, and whether to grasp with two or three fingers. Similarly, Rezzoug and Gorce (Rezzoug and Gorce 2003; Gorce and Rezzoug 2004) have a model with a modular architecture that first learns inverse kinematics for the fingers using backpropagation and then learns hand configuration for grasping using reinforcement. The Fingers Configuration Neural Network (FCNN) learns finger postures given the desired position of the fingertip.

Several models plan grasps in terms of finger contact points on the object. Molina-Vilaplana et al. (2007) use heuristics to select contact points on the object surface for the fingers. Rezzoug and Gorce's (Rezzoug and Gorce 2003; Gorce and Rezzoug 2004) model uses a Hand Configuration Neural Network (HCNN) that learns contact locations for each finger. Kamon et al. (1996) split the problem of grasp learning into two problems: choosing grasping points and predicting the quality of the grasp. Each grasp type has a certain set of location and quality parameters used to select grasp locations and predict quality, which are supplied beforehand as task-specific knowledge. Grasp quality is predicted by the angles between the fingers and the surface normals at the contact points, and the distance between the opposition axis and the center of mass. They suggest that the grasp specification and evaluation modules run in an alternating manner until a suitable grasp is selected. Faldella et al. (1993) present an interesting model where a neural mechanism for matching object geometry to hand shape interacts with a symbolic rule-based expert system. The symbolic system performs some geometric analyses such as identifying curvature type, selects candidate hand contact positions, identifies symmetric situations (to reduce neural module input), and ranks the selected grasp according to task constraints. Multilayer perceptrons trained using backpropagation were used to determine potential grasps based on geometric information.

For Grupen and Coelho (2000), grasping is primarily a force domain task, emphasizing *force* closure around an object over *form* (of the hand) closure. Their model used closed loop control, utilizing tactile feedback to reposition contact forces based on models of interaction between the contacts and object surface. A Markov decision process (MDP) framework is used to select a sequence of controllers to maximize object stability without knowing the object's identity, geometry, or pose. This model is complementary to those that focus on visual-based kinematic grasp planning in that it learns to use haptic feedback in order to reposition contact forces to stabilize the object. Including such a mechanism in a model like ILGA would increase the number of successful grasps early in training and may provide a way to adjust grasp plans that lead to initially unstable grasps that can be subsequently stabilized with corrective hand movements.

The Two Route, Prefrontal Instruction, Competition of Affordances, Language Simulation (TRoPICALS) model is a neural network model similar to FARS which includes the influence of the dorsal and ventral stream in affordance selection and grasp planning (Caligiore et al. 2010). Like ILGA, TRoPICALS utilizes SOMs, winner-take-all processes, and a staged learning approach to learn to represent object shapes and motor parameters appropriate for grasping them. While ILGA does not include the ventral visual stream, TRoPICALS includes the influence of language and ventral stream recognition of object identity on affordance selection and is able to reproduce the pattern of reaction times shown by human subjects when presented with conflicting grasp instructions and affordances. However, ILGA uses a more complex simulation environment, captures the emergence of affordance representations in AIP, and addresses the linkage between reaching and grasping in ways that TRoPICALS does not.

Each of these models addresses some aspect of grasping that is neglected in this model. In ILGA, only the inverse kinematics for the arm is used to plan reaching movements, while grasping movements are controlled by rotating the wrist and preshaping then enclosing the fingers. Grasping is controlled in an open loop manner without a notion of target contact points on the object surface, and kinetics is handled by PD controllers that are not concerned with balancing force application along an opposition axis. However, ILGA is more biologically plausible in that it uses neural representations based on neurophysiological data, modules with connectivity constrained by anatomical data, and biologically inspired learning rules. ILGA could be extended to a more complete model by including aspects of these other models such as inverse kinematics for the hand and feedback-based grasp control (described in Skilled grasping, above), using visual feedback for the control of the hand, and utilizing haptic feedback for corrective movements.

Integration of ILGA with models of the primate mirror system could provide the system with a feedback signal for skilled grasping. The neurons in the AIP module of ILGA correspond to visual-dominant neurons in monkey area AIP and project to the F5 module, but do not receive reciprocal connections from F5. In reality, area AIP also contains neurons classified as visuomotor and motor dominant and receives feedback from F5 (Sakata et al. 1995). Motor-dominant cells respond during grasping in the light and the dark, while visual-dominant cells respond only during object fixation and grasping in the light and visuomotor cells respond during object fixation and grasping in the dark but fire most strongly during grasping in the light. Sakata et al. (1995) offer a conceptual model of feedback-based grasping in which F5 canonical neurons provide AIP motor-dominant neurons with a copy of the grasp motor command, which then pass this signal to AIP visuomotor neurons which combine

this information with information from AIP visual-dominant neurons and project back to F5. In this way, if the ongoing grasp does not match the encoded affordance, the grasp plan in F5 is modified or aborted. AIP visual-dominant neurons are classified into object-type neurons that fire during object fixation, and non-object-type neurons that fire during grasping in the light but not object fixation and may respond to the sight of the hand during the grasp. Non-object-type neurons are seldom mentioned in discussions of AIP, but make up half of visual-dominant neurons in the region (Sakata et al. 1995). Interestingly, their existence fits with the hypothesis outlined by Oztop and Arbib (2002) that F5 mirror neurons evolved to provide visual feedback of the shape of the hand relative to the object's affordances. We suggest that non-object-type AIP neurons obtain their properties by projections from F5 mirror neurons, and that these projections are used for visual feedback-based control of grasping. It has been shown that reversible inactivation of F5 mirror neurons by muscimol injection in the cortical convexity of the arcuate sulcus results in slower, clumsy grasps (Fogassi et al. 2001), consistent with the idea of F5 mirror neurons playing a role in providing visual feedback. Our conceptual model predicts that muscimol injection in the cortical convexity of F5 will abolish the response of non-object-type AIP neurons to the sight of the grasp.

In addition to reversible inactivation of mirror neurons, Fogassi et al. (2001) tested the effects of muscimol injection in the bank of the arcuate sulcus where most F5 canonical neurons are located. In this case, the hand preshape was impaired but monkeys were still able to grasp the object by contacting it and then making appropriate corrective movements using tactile feedback. This seems similar to the process modeled by Grupen and Coelho (2000) in which haptic feedback is used to reposition contact forces. The fact that corrective movements can still be made after F5 inactivation suggests that they are not based on F5 activity and may be implemented by the direct projection included in the FARS model from the primary somatosensory area S1 to the primary motor cortex.

4.4.1 Future directions

ILGA combines grasp learning ability with a mature reaching system that is smooth and accurate. However, in infants, reaching and grasping develop simultaneously. By allowing the development of simultaneously maturing reach and grasp systems to influence each other, future versions of the model may be able to address data showing how the reach and grasp movements are coordinated based on the object to be grasped (Supuk et al. 2005) or the distal goal of the movement (i.e., grasping to lift versus to place, Ansuini et al. 2006).

We have used as realistic simulation of the arm, hand, and object as possible, within the limitations of our simu-

lation engine, in order to facilitate translation of ILGA to a robotic platform. However, several issues would need to be addressed for this transition. To simplify the task, we did not include gravity, so that objects had mass, but no weight. This required the model to preshape the hand and enclose it around the object to stabilize it, but did not require the arm to counteract the weight of the object. Additionally, we did not simulate objects made of different materials. All of our objects had the same friction coefficient and therefore a robotic implementation of ILGA would require the computation of the appropriate contact forces for objects of different types.

ILGA represents reach and grasp rotation parameters using Euler angles. While this representation is more compact than rotation matrices, for example, it suffers from several problems such as composition difficulty (Pastor and Righetti 2011) and singularities which can cause gimbal lock (Feix and Romero 2013). Quaternions use one more parameter than Euclidean angles, but do not suffer from these problems while still having lower dimensionality than rotation matrices (Feix and Romero 2013). It is not known what representation the brain uses to represent joint rotation, but future versions of ILGA could explore the use of quaternion representations to avoid the issues caused by Euler angles and reduce dimensionality.

5 Conclusion

ILGA is the only developmental model of grasping to date that simultaneously learns to extract affordances from object features and select motor parameters to successfully grasp them. We have shown that the model develops distributed representations in area AIP similar to those reported in the experimental literature and can use these representations to generalize grasp plans to objects of varying sizes and at different orientations and positions. Finally, we presented several neurophysiologically testable predictions made by the model and discussed ways in which it could be extended to handle context-dependent grasping of complex objects and skilled manipulation.

Acknowledgments This material is based in part on work supported by the National Science Foundation under Grant No. BCS-1343544 “INSPIRE Track 1: Action, Vision and Language, and their Brain Mechanisms in Evolutionary Relationship” (Michael A. Arbib, Principal Investigator).

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Amari S, Arbib M (1977) Competition and cooperation in neural nets. In: Metzler J (ed) *Systems neuroscience*. Academic Press, New York, pp 119–166
- Ansuini C, Santello M, Massaccesi S, Castiello U (2006) Effects of end-goal on hand shaping. *J Neurophysiol* 95(4):2456–2465. doi:10.1152/jn.01107.2005
- Arbib M, Iberall T, Lyons D (1985) Coordinated control programs for movements of the hand. *Exp Brain Res* 111–129
- Arbib MA, Bonaiuto JB, Jacobs S, Frey SH (2009) Tool use and the distalization of the end-effector. *Psychol Res* 73:441–462. doi:10.1007/s00426-009-0242-2
- Arbib MA, Fagg AH, Grafton ST (2002) Synthetic PET imaging for grasping: from primate neurophysiology to human behavior. In: Soomer FT, Wichert A (eds) *Exploratory analysis and data modeling in functional neuroimaging*. MIT Press, Cambridge, pp 231–250
- Bar-Gad I, Morris G, Bergman H (2003) Information processing, dimensionality reduction and reinforcement learning in the basal ganglia. *Prog Neurobiol* 71:439–473. doi:10.1016/j.pneurobio.2003.12.001
- Barto AG (1995) Adaptive critics and the basal ganglia. In: Houk JC, Davis JL, Beiser DG (eds) *Models of information processing in the basal ganglia*. MIT Press, Cambridge, pp 215–232
- Battaglia-Mayer A, Caminiti R, Lacquaniti F, Zago M (2003) Multiple levels of representation of reaching in the parieto-frontal network. *Cereb Cortex* 13:1009–1022
- Battaglia PP, Muzur A, Galletti C et al (2002) Effects of lesions to area V6A in monkeys. *Exp Brain Res* 144:419–422. doi:10.1007/s00221-002-1099-4
- Baumann MA, Fluet MC, Scherberger H (2009) Context-specific grasp movement representation in the macaque anterior intraparietal area. *J Neurosci* 29:6436–6448. doi:10.1523/JNEUROSCI.5479-08.2009
- Berthier NE (1996) Learning to reach: a mathematical model. *Dev Psychol* 32:811–823
- Berthier NE, Clifton RK, Gullapalli V, McCall DD, Robin DJ (1996) Visual information and object size in the control of reaching. *J Mot Behav* 28(3):187–197. doi:10.1080/00222895.1996.9941744
- Berthier NE, Rosenstein MT, Barto AG (2005) Approximate optimal control as a model for motor learning. *Psychol Rev* 112:329–346. doi:10.1037/0033-295x.112.2.329
- Bhat A, Heathcock J, Galloway JC (2005) Toy-oriented changes in hand and joint kinematics during the emergence of purposeful reaching. *Infant Behav Dev* 28:445–465
- Bhat AN, Galloway JC (2006) Toy-oriented changes during early arm movements: hand kinematics. *Infant Behav Dev* 29:358–372. doi:10.1016/j.infbeh.2006.01.005
- Borra E, Belmalih A, Calzavara R et al (2007) Cortical connections of the macaque anterior intraparietal (AIP) area—Borra et al. 10.1093/cercor/bhm146—Cerebral Cortex. *Cereb Cortex*. doi:10.1093/cercor/bhm146
- Bullock D, Grossberg S, Guenther FH (1993) A self-organizing neural model of motor equivalent reaching and tool use by a multijoint arm. *J Cogn Neurosci* 5:408–435
- Caligiore D, Borghi AM, Parisi D, Baldassarre G (2010) TRoPICALS: a computational embodied neuroscience model of compatibility effects. *Psychol Rev* 117:1188–1228. doi:10.1037/a0020887
- Caligiore D, Parisi D, Baldassarre G (2014) Integrating reinforcement learning, equilibrium points, and minimum variance to understand the development of reaching: a computational model. *Psychol Rev* 121(3):389–421. doi:10.1037/a0037016

- Caminiti R, Ferraina S, Johnson PB (1996) The sources of visual information to the primate frontal lobe: a novel role for the superior parietal lobule. *Cereb Cortex* 6:319–328
- Caminiti R, Johnson PB, Galli C et al (1991) Making arm movements within different parts of space: the premotor and motor cortical representation of a coordinate system for reaching to visual targets. *J Neurosci* 11:1182–1197
- Cangelosi A, Schlesinger M (2013) *Developmental robotics: from babies to robots*. MIT Press/Bradford Books, Cambridge
- Chan SS, Moran DW (2006) Computational model of a primate arm: from hand position to joint angles, joint torques and muscle forces. *J Neural Eng* 3:327–337. doi:[10.1088/1741-2560/3/4/010](https://doi.org/10.1088/1741-2560/3/4/010)
- Chentanez N, Barto AG, Singh SP (2004) Intrinsically motivated reinforcement learning. In: *Advances in neural information processing systems* pp. 1281–1288
- Cisek P, Kalaska JF (2002) Simultaneous encoding of multiple potential reach directions in dorsal premotor cortex. *J Neurophysiol* 87:1149–1154
- Claxton LJ, Keen R, McCarty ME (2003) Evidence of motor planning in infant reaching behavior. *Psychol Sci* 14:354–356
- Clifton RK, Muir DW, Ashmead DH, Clarkson MG (1993) Is visually guided reaching in early infancy a myth? *Child Dev* 64:1099–1110
- Crammond DJ, Kalaska JF (2000) Prior information in motor and premotor cortex: activity during the delay period and effect on pre-movement activity. *J Neurophysiol* 84:986–1005
- Doya K, Samejima K, Katagiri K, Kawato M (2002) Multiple model-based reinforcement learning. *Neural Comput* 14:1347–1369. doi:[10.1162/089976602753712972](https://doi.org/10.1162/089976602753712972)
- Dum RP, Strick PL (2005) Frontal lobe inputs to the digit representations of the motor areas on the lateral surface of the hemisphere. *J Neurosci* 25:1375–1386
- Durbin R, Mitchison G (1990) A dimension reduction framework for understanding cortical maps. *Nature* 343:644–647. doi:[10.1038/343644a0](https://doi.org/10.1038/343644a0)
- Erlhagen W, Schoner G (2002) Dynamic field theory of movement preparation. *Psychol Rev* 109:545–572
- Fagg AH, Arbib MA (1998) Modeling parietal-premotor interactions in primate control of grasping. *Neural Netw* 11:1277–1303. doi:[10.1016/S0893-6080\(98\)00047-1](https://doi.org/10.1016/S0893-6080(98)00047-1)
- Faldella E, Fringuelli B, Zanichelli F (1993) A hybrid system for knowledge-based synthesis of robot grasps. In: *IEEE/RSJ international conference on intelligent robots and systems*
- Fattori P, Kutz DF, Breveglieri R et al (2005) Spatial tuning of reaching activity in the medial parieto-occipital cortex (area V6A) of macaque monkey. *Eur J Neurosci* 22:956–972
- Feix T, Romero J (2013) A metric for comparing the anthropomorphic motion capability of artificial hands. *IEEE Trans Robot* 29(1):82–93
- Fogassi L, Gallese V, Buccino G et al (2001) Cortical mechanism for the visual guidance of hand grasping movements in the monkey: A reversible inactivation study. *Brain* 124:571–586
- Fogassi L, Raos V, Franchi G et al (1999) Visual responses in the dorsal premotor area F2 of the macaque monkey. *Exp Brain Res* 128:194–199
- Freund HJ (1990) Premotor area and preparation of movement. *Rev Neurol* 146:543–547
- Fu QG, Suarez JJ, Ebner TJ (1993) Neuronal specification of direction and distance during reaching movements in the superior precentral premotor area and primary motor cortex of monkeys. *J Neurophysiol* 70:2097–2116
- Fujii HN, Mushiake, Tanji J (1996) Rostrocaudal differentiation of dorsal premotor cortex with physiological criteria. *Soc Neurosci Abstr*
- Fujii N, Mushiake H, Tanji J (2002) Distribution of eye- and arm-movement-related neuronal activity in the SEF and in the SMA and Pre-SMA of monkeys. *J Neurophysiol* 87:2158–2166
- Gallese V, Murata A, Kaseda M et al (1994) Deficit of hand preshaping after muscimol injection in monkey parietal cortex. *Neuroreport* 5:1525–1529
- Galletti C, Battaglini PP, Fattori P (1993) Parietal neurons encoding spatial locations in craniotopic coordinates. *Exp Brain Res* 96:221–229
- Galletti C, Fattori P, Kutz DF, Battaglini PP (1997) Arm movement-related neurons in the visual area V6A of the macaque superior parietal lobule. *Eur J Neurosci* 9:410–413
- Galletti C, Fattori P, Kutz DF, Gamberini M (1999) Brain location and visual topography of cortical area V6A in the macaque monkey. *Eur J Neurosci* 11:575–582
- Galletti C, Kutz DF, Gamberini M et al (2003) Role of the medial parieto-occipital cortex in the control of reaching and grasping movements. *Exp Brain Res* 153:158–170. doi:[10.1007/s00221-003-1589-z](https://doi.org/10.1007/s00221-003-1589-z)
- Gibson J (1966) *The senses considered as perceptual systems*. Houghton Mifflin, Oxford
- Goodale MA, Milner AD (1992) Separate visual pathways for perception and action. *Trends Neurosci* 15:20–25. doi:[10.1016/0166-2236\(92\)90344-8](https://doi.org/10.1016/0166-2236(92)90344-8)
- Gorce P, Rezzoug N (2004) A method to learn hand grasping posture from noisy sensing information. *Robotica* 22:309–318
- Gordon J, Ghilardi MF, Ghez C (1994) Accuracy of planar reaching movements. I. Independence of direction and extent variability. *Exp Brain Res* 99:97–111
- Gritsenko V, Yakovenko S, Kalaska JF (2009) Integration of predictive feedforward and sensory feedback signals for online control of visually guided movement. *J Neurophysiol* 102(2):914–930
- Gruppen R, Coelho JAJ (2000) Structure and growth: a model of development for grasping with robot hands. In: *IEEE/RSJ international conference on intelligent robots and systems*. pp 1987–1992
- Gurney K, Prescott TJ, Redgrave P (2001) A computational model of action selection in the basal ganglia. I. A new functional anatomy. *Biol Cybern* 84:401–410
- Hoff B, Arbib MA (1993) Models of trajectory formation and temporal interaction of reach and grasp. *J Mot Behav* 25(3):175–192
- Ijspeert AJ, Nakanishi J, Schaal S (2002) Movement imitation with nonlinear dynamical systems in humanoid robots. In: *IEEE international conference on robotics and automation*. New York, pp 1398–1403
- Jeannerod M, Arbib MA, Rizzolatti G, Sakata H (1995) Grasping objects: the cortical mechanisms of visuomotor transformation. *Trends Neurosci* 18:314–320
- Jeannerod M, Decety J, Michel F (1994) Impairment of grasping movements following a bilateral posterior parietal lesion. *Neuropsychologia* 32:369–380. doi:[10.1016/0028-3932\(94\)90084-1](https://doi.org/10.1016/0028-3932(94)90084-1)
- Johnson PB, Ferraina S, Bianchi L, Caminiti R (1996) Cortical networks for visual reaching: physiological and anatomical organization of frontal and parietal lobe arm regions. *Cereb Cortex* 6:102–119
- Johnson PB, Ferraina S, Caminiti R (1993) Cortical networks for visual reaching. *Exp Brain Res* 97:361–365
- Takei S, Hoffman DS, Strick PL (1999) Muscle and movement representations in the primary motor cortex. *Science* 285(5436):2136–2139
- Takei S, Hoffman DS, Strick PL (2001) Direction of action is represented in the ventral premotor cortex. *Nat Neurosci* 4:1020–1025. doi:[10.1038/nn726](https://doi.org/10.1038/nn726)
- Takei S, Hoffman DS, Strick PL (2003) Sensorimotor transformations in cortical motor areas. *Neurosci Res* 46:1–10
- Kalaska JF (2009) From intention to action: motor cortex and the control of reaching movements. *Adv Exp Med Biol* 629:139–178. doi:[10.1007/978-0-387-77064-2_8](https://doi.org/10.1007/978-0-387-77064-2_8)
- Kamon I, Flash T, Edelman S (1996) Learning to grasp using visual information. In: *IEEE international conference on robotics and automation*. citeseer, Minneapolis, MN, pp 2470–2476

- Kropotov JD, Etlinger SC (1999) Selection of actions in the basal ganglia-thalamocortical circuits: review and model. *Int J Psychophysiol* 31:197–217
- Kuhtz-Buschbeck JP, Stolze H, Johnk K et al (1998) Development of prehension movements in children: a kinematic study. *Exp Brain Res* 122:424–432
- Kuperstein M (1988) Neural model of adaptive hand-eye coordination for single postures. *Science* 239:1308–1311. doi:[10.1126/science.3344437](https://doi.org/10.1126/science.3344437)
- Kurata K (1993) Premotor cortex of monkeys: set- and movement-related activity reflecting amplitude and direction of wrist movements. *J Neurophysiol* 69:187–200
- Kurata K (1994) Site of origin of projections from the thalamus to dorsal versus ventral aspects of the premotor cortex of monkeys. *Neurosci Res* 21:71–76
- Kurata K, Hoffman DS (1994) Differential effects of muscimol microinjection into dorsal and ventral aspects of the premotor cortex of monkeys. *J Neurophysiol* 71:1151–1164
- Kusunoki M, Tanaka Y, Ohtsuka H et al (1993) Selectivity of the parietal visual neurons in the axis orientation of objects in space. *Soc Neurosci Abstr* 19:770
- Lasky RE (1977) The effect of visual feedback of the hand on the reaching and retrieval behavior of young infants. *Child Dev* 48:112–117
- Lockman JJ, Ashmead DH, Bushnell EW (1984) The development of anticipatory hand orientation during infancy. *J Exp Child Psychol* 37:176–186
- Luppino G, Hamed SB, Gamberini M et al (2005) Occipital (V6) and parietal (V6A) areas in the anterior wall of the parieto-occipital sulcus of the macaque: a cytoarchitectonic study. *Eur J Neurosci* 21:3056–3076
- Ma Umilta, Brochier T, Spinks RL, Lemon RN (2007) Simultaneous recording of macaque premotor and primary motor cortex neuronal populations reveals different functional contributions to visuomotor grasp. *J Neurophysiol* 98:488–501. doi:[10.1152/jn.01094.2006](https://doi.org/10.1152/jn.01094.2006)
- Marconi B, Genovesio A, Battaglia-Mayer A et al (2001) Eye-hand coordination during reaching. I. Anatomical relationships between parietal and frontal cortex. *Cereb Cortex* 11:513–527
- Matelli M, Govoni P, Galletti C et al (1998) Superior area 6 afferents from the superior parietal lobule in the macaque monkey. *J Comp Neurol* 402:327–352
- Matelli M, Luppino G, Rizzolatti G (1985) Patterns of cytochrome oxidase activity in the frontal agranular cortex of the macaque monkey. *Behav Brain Res* 18:125–136
- McCarty ME, Clifton RK, Collard RR (1999) Problem solving in infancy: the emergence of an action plan. *Dev Psychol* 35:1091–1101
- McIntyre J, Stratta F, Lacquaniti F (1998) Short-term memory for reaching to visual targets: psychophysical evidence for body-centered reference frames. *J Neurosci* 18:8423–8435
- Messier J, Kalaska JF (2000) Covariation of primate dorsal premotor cell activity with direction and amplitude during a memorized-delay reaching task. *J Neurophysiol* 84:152–165
- Molina-Vilaplana J, Feliu-Batlle J, Lopez-Coronado J (2007) A modular neural network architecture for step-wise learning of grasping tasks. *Neural Netw* 20:631–645
- Morrongiello BA, Rocca PT (1989) Visual feedback and anticipatory hand orientation during infants' reaching. *Percept Mot Ski* 69:787–802
- Mollazadeh M, Aggarwal V, Thakor NV, Schieber MH (2014) Principal components of hand kinematics and neurophysiological signals in motor cortex during reach to grasp movements. *J Neurophysiol* 112(8):1857–1870
- Murata A, Gallese V, Luppino G et al (2000) Selectivity for the shape, size, and orientation of objects for grasping in neurons of monkey parietal area AIP. *J Neurophysiol* 83:2580–2601
- Nakamura H, Kuroda T, Wakita M et al (2001) From three-dimensional space vision to prehensile hand movements: the lateral intraparietal area links the area V3A and the anterior intraparietal area in macaques. *J Neurosci* 21:8174–8187
- Nataraj R, Pasluosta C, Li ZM (2014) Online kinematic regulation by visual feedback for grasp versus transport during reach-to-pinch. *Hum Mov Sci* 36:134–153
- Newell KM, McDonald PV, Baillargeon R (1993) Body scale and infant grip configurations. *Dev Psychobiol* 26:195–205. doi:[10.1002/dev.420260403](https://doi.org/10.1002/dev.420260403)
- Olson CR, Gettner SN (1995) Object-centered direction selectivity in the macaque supplementary eye field. *Science* 269:985–988
- Oztop E, Arbib MA (2002) Schema design and implementation of the grasp-related mirror neuron system. *Biol Cybern* 87:116–140
- Oztop E, Bradley NS, Arbib MA (2004) Infant grasp learning: a computational model. *Exp Brain Res* 158:480–503. doi:[10.1007/s00221-004-1914-1](https://doi.org/10.1007/s00221-004-1914-1)
- Oztop E, Imamizu H, Cheng G, Kawato M (2006) A computational model of anterior intraparietal (AIP) neurons. *Neurocomputing* 69:1354–1361. doi:[10.1016/j.neucom.2005.12.106](https://doi.org/10.1016/j.neucom.2005.12.106)
- Pastor P, Righetti L (2011) Online movement adaptation based on previous sensor experiences. In: 2011 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). San Francisco, CA
- Pesaran B, Nelson MJ, Andersen RA (2006) Dorsal premotor neurons encode the relative position of the hand, eye, and goal during reach planning. *Neuron* 51:125–134
- Raghavan P, Santello M, Gordon AM, Krakauer JW (2010) Compensatory motor control after stroke: an alternative joint strategy for object-dependent shaping of hand posture. *J Neurophysiol* 103(6):3034–3043
- Raos V, Umiltà MA, Gallese V, Fogassi L (2004) Functional properties of grasping-related neurons in the dorsal premotor area F2 of the macaque monkey. *J Neurophysiol* 92:1990–2002. doi:[10.1152/jn.00154.2004](https://doi.org/10.1152/jn.00154.2004)
- Raos V, Umiltà M-A, Murata A et al (2006) Functional properties of grasping-related neurons in the ventral premotor area F5 of the macaque monkey. *J Neurophysiol* 95:709–729. doi:[10.1152/jn.00463.2005](https://doi.org/10.1152/jn.00463.2005)
- Rezzoug N, Gorce P (2003) A biocybernetic method to learn hand grasping posture. *Kybernetes* 32:478–490
- Riehle A, Requin J (1989) Monkey primary motor and premotor cortex: single-cell activity related to prior information about direction and extent of an intended movement. *J Neurophysiol* 61:534–549
- Rizzolatti G, Camarda R (1988) Functional organization of inferior area 6 in the macaque monkey. *Exp Brain Res* 6:491–507
- Rizzolatti G, Luppino G (2001) The cortical motor system. *Neuron* 31:889–901
- Rizzolatti G, Luppino G, Matelli M (1998) The organization of the cortical motor system: new concepts. *Electroencephalogr Clin Neurophysiol* 106:283–296
- Rizzolatti G, Matelli M (2003) Two different streams form the dorsal visual system: anatomy and functions. *Exp Brain Res* 153:146–157. doi:[10.1007/s00221-003-1588-0](https://doi.org/10.1007/s00221-003-1588-0)
- Sakata H, Taira M, Kusunoki M et al (1998) Neural coding of 3D features of objects for hand action in the parietal cortex of the monkey. *Philos Trans R Soc Lond B Biol Sci* 353:1363–1373
- Sakata H, Taira M, Kusunoki M et al (1997) The TINS Lecture. The parietal association cortex in depth perception and visual control of hand action. *Trends Neurosci* 20:350–357
- Sakata H, Taira M, Kusunoki M et al (1999) Neural representation of three-dimensional features of manipulation objects with stereopsis. *Exp Brain Res* 128:160–169
- Sakata H, Taira M, Murata A, Mine S (1995) Neural mechanisms of visual guidance of hand action in the parietal cortex of the monkey. *Cereb Cortex* 5:429–438

- Sakata H, Tsutsui K, Taira M (2005) Toward an understanding of the neural processing for 3D shape perception. *Neuropsychologia* 43:151–161. doi:[10.1016/j.neuropsychologia.2004.11.003](https://doi.org/10.1016/j.neuropsychologia.2004.11.003)
- Santello M, Flanders M, Soechting JF (2002) Patterns of hand motion during grasping and the influence of sensory guidance. *J Neurosci* 22(4):1426–1435
- Schettino LF, Adamovich SV, Poizner H (2003) Effects of object shape and visual feedback on hand configuration during grasping. *Exp Brain Res* 151:158–166. doi:[10.1007/s00221-003-1435-3](https://doi.org/10.1007/s00221-003-1435-3)
- Shikata E, Tanaka Y, Nakamura H et al (1996) Selectivity of the parietal visual neurones in 3D orientation of surface of stereoscopic stimuli. *Neuroreport* 7:2389–2394
- Shipp S, Blanton M, Zeki S (1998) A visuo-somatomotor pathway through superior parietal cortex in the macaque monkey: cortical connections of areas V6 and V6A. *Eur J Neurosci* 10:3171–3193
- Shipp S, Zeki S (1995) Segregation and convergence of specialised pathways in macaque monkey visual cortex. *J Anat* 187(Pt 3):547–562
- Singh SP, Sutton RS (1996) Reinforcement learning with replacing eligibility traces. *Mach Learn* 22:123–158
- Smeets JB, Brenner E (2001) Independent movements of the digits in grasping. *Exp Brain Res* 139:92–100
- Soechting JF, Flanders M (1989) Sensorimotor representations for pointing to targets in three-dimensional space. *J Neurophysiol* 62:582–594
- Supuk T, Kodek T, Bajd T (2005) Estimation of hand preshaping during human grasping. *Med Eng Phys* 27(9):790–797. doi:[10.1016/j.medengphy.2005.03.008](https://doi.org/10.1016/j.medengphy.2005.03.008)
- Sutton RS, Barto AG (1998) Reinforcement learning: an introduction. MIT Press, Cambridge
- Taira M, Mine S, Georgopoulos AP et al (1990) Parietal cortex neurons of the monkey related to the visual guidance of hand movement. *Exp Brain Res* 83:29–36
- Tanne J, Boussaoud D, Boyer-Zeller N, Rouiller EM (1995) Direct visual pathways for reaching movements in the macaque monkey. *Neuroreport* 7:267–272
- Thelen E, Corbetta D, Kamm K, Spencer JP, Schneider K, Zernicke RF (1993) The transition to reaching: mapping intention and intrinsic dynamics. *Child Dev* 64(4):1058–1098
- Todorov E (2000) Direct cortical control of muscle activation in voluntary arm movements: a model. *Nat Neurosci* 3:391–398. doi:[10.1038/73964](https://doi.org/10.1038/73964)
- Tsutsui K, Jiang M, Yara K et al (2001) Integration of perspective and disparity cues in surface-orientation-selective neurons of area CIP. *J Neurophysiol* 86:2856–2867
- Tsutsui K, Taira M, Sakata H (2005) Neural mechanisms of three-dimensional vision. *Neurosci Res* 51:221–229. doi:[10.1016/j.neures.2004.11.006](https://doi.org/10.1016/j.neures.2004.11.006)
- Tunik E, Frey SH, Grafton ST (2005) Virtual lesions of the anterior intraparietal area disrupt goal-dependent on-line adjustments of grasp. *Nat Neurosci* 8:505–511
- Von Hofsten C (1982) Eye-hand coordination in the newborn. *Dev Psychol* 18:450–461
- Von Hofsten C, Ronnqvist L (1988) Preparation for grasping an object: a developmental study. *J Exp Psychol Hum Percept Perform* 14:610–621
- Von Hofsten C, Spelke ES (1985) Object perception and object-directed reaching in infancy. *J Exp Psychol Gen* 114:198–212
- Weinrich M, Wise SP (1982) The premotor cortex of the monkey. *J Neurosci* 2:1329–1345
- Weitzenfeld A, Arbib MA, Alexander A (2002) The neural simulation language: a framework for brain modeling. MIT Press, Cambridge
- Winges SA, Santello M (2005) From single motor unit activity to multiple grip forces: mini-review of multi-digit grasping. *Integr Comp Biol* 45:679–682
- Winges SA, Weber DJ, Santello M (2003) The role of vision on hand preshaping during reach to grasp. *Exp Brain Res* 152(4):489–498
- Wise SP, Boussaoud D, Johnson PB, Caminiti R (1997) Premotor and parietal cortex: corticocortical connectivity and combinatorial computations. *Annu Rev Neurosci* 20:25–42. doi:[10.1146/annurev.neuro.20.1.25](https://doi.org/10.1146/annurev.neuro.20.1.25)
- Witherington DC (2005) The development of prospective grasping control between 5 and 7 months: a longitudinal study. *Infancy* 7:143–161
- Wu S, Amari S, Nakahara H (2002) Population coding and decoding in a neural field: a computational study. *Neural Comput* 14:999–1026. doi:[10.1162/089976602753633367](https://doi.org/10.1162/089976602753633367)
- Xiao J, Padoa-Schioppa C, Bizzi E (2006) Neuronal correlates of movement dynamics in the dorsal and ventral premotor area in the monkey. *Exp Brain Res* 168:106–119. doi:[10.1007/s00221-005-0074-2](https://doi.org/10.1007/s00221-005-0074-2)
- Yin H (2008) On multidimensional scaling and the embedding of self-organising maps. *Neural Netw* 21:160–169. doi:[10.1016/j.neunet.2007.12.027](https://doi.org/10.1016/j.neunet.2007.12.027)
- Zatsiorsky VM, Latash ML (2004) Prehension synergies. *Exerc Sport Sci Rev* 32:75–80