An active inference model of hierarchical action understanding, learning and imitation

Riccardo Proietti¹, Giovanni Pezzulo^{2,*} and Alessia Tessari¹

¹*Department of Psychology, University of Bologna, Italy*

²*Institute of Cognitive Sciences and Technologies, National Research Council, Rome, Italy*

Corresponding author:

Giovanni Pezzulo

Institute of Cognitive Sciences and Technologies,

National Research Council,

00185 Rome, Italy

E-mail: giovanni.pezzulo@istc.cnr.it

Abstract

We advance a novel active inference model of the cognitive processing that underlies the acquisition of a hierarchical action repertoire and its use for observation, understanding and imitation. We illustrate the model in four simulations of a tennis learner who observes a teacher performing tennis shots, forms hierarchical representations of the observed actions, and imitates them. Our simulations show that the agent's oculomotor activity implements an active information sampling strategy that permits inferring the kinematics aspects of the observed movement, which lie at the lowest level of the action hierarchy. In turn, this low-level kinematic inference supports higher-level inferences about deeper aspects of the observed actions: proximal goals and intentions. Finally, the inferred action representations can steer imitative responses, but interfere with the execution of different actions. Our simulations show that hierarchical active inference provides a unified account of action observation, understanding, learning and imitation and helps explaining the neurobiological underpinnings of visuomotor cognition, including the multiple routes for action understanding in the dorsal and ventral streams and mirror mechanisms.

Keywords: action understanding; active inference; hierarchical model; action representation; visuomotor control

1. Introduction: action understanding as a hierarchical inference problem

Understanding actions performed by other people is vital for social cognition. An action can be defined as a sequence of kinematic bodily movements (e.g., movements of the left arm and fingers) elicited and monitored by a goal (e.g., grasping an object). Hence, we assume that action understanding amounts to inferring the actor's goal by observing her movement kinematics, such as the positions of her limbs, angles of joints, their relative positions, and respect to objects and the context (Jacob and Jeannerod, 2005). In turn, since an action may result from many movements, the inference of action goals from observed movements constitutes an inverse problem (Csibra and Gergely, 2007; Jacob and Jeannerod, 2005).

Several researchers proposed that to solve this inverse (one-to-many) problem, the brain adopts a probabilistic strategy, which can be formulated in Bayesian terms as:

$$p(goal | movement) = p(goal) * p(movement | goal) / p(movement)$$

In the above equation, the (*posterior*) probability of an action goal given the observed movement $p(goal | movement)$ is proportional to the (*prior*) probability of the goal $p(goal)$ before observing the movement and the probability $p(movement | goal)$ that a specific goal generates the observed movement (likelihood). These two (prior and likelihood) terms constitute a so-called *generative model* of how actions are generated, that the brain uses (technically, "inverts") for action recognition - or the inference about the action goal that may have produced an observed movement (Baker et al., 2009, 2006; Dindo et al., 2011; Kilner et al., 2007).

It has been proposed that the brain uses the same generative model for both action execution and recognition and this generative model is structured hierarchically (Byrne and Russon, 1998; Grafton and de C. Hamilton, 2007; Jeannerod, 1994; Kilner, 2011; Pacherie, 2008). In this study, we present a computational analysis of these hierarchical action representations, by considering four levels of representation of an exemplificative set of skilled movements, namely, tennis movements; see Figure 1. At the bottom of the hierarchy, the *kinematic level* encodes the kinematic features of movements (e.g., speed and acceleration) of actions. For example, in our tennis context, the kinematic level regards specific body parts involved in the execution of a tennis shot. One step higher in the hierarchy, we formalize a *postural level* where different combinations of body parts constitute postures that are the building blocks of tennis movements. The further hierarchical level encodes the *proximal goal* of the action, such as the aim to execute a particular tennis shot, by enacting a sequence of postures. The highest, *intention level* represents the ultimate reason for executing the action, which in this setting corresponds to a general class of shots: forehand, backhand, or smash. Notably, each class of shots has multiple realizations at the proximal goal and kinematic levels (e.g., a forehand can be done to the left or right), but the intention level is more generic and agnostic about particular realizations.

Hierarchical action representations: an example in the context of tennis

Figure 1. This schematic illustrates the hierarchical representation of an action repertoire in the example scenario that we adopt in this paper: playing tennis. The hierarchy comprises three levels of hidden variables (postures, proximal goals, and intentions) and an observation layer (kinematics). The higher the level of the hidden variable, the more the representation is associated with abstract and generalized action meanings. Here, kinematics aspects are exemplified with body parts which spatial relations constitute different body postures; proximal goals denote sequences of postures that realize tennis shots, in this case a left or right forehand; rather, intentions denote the most general class of shots, here forehand, that entails multiple realizations at the proximal goal and kinematic levels. Note that this schematic (and our computational model) omits the motor level, which is responsible for the specification of patterns of muscle activity to execute the movements specified at the kinematic level. This schematic also illustrates that in this hierarchical setting, action understanding amounts to inferring the higher levels of the hierarchy (e.g., goals and intentions) by only having access to low-level (kinematic) observations.

In sum, the hierarchical arrangement of the generative model illustrated in Figure 1 implies that the brain encodes a rich cognitive representation of actions, which links together more abstract features of action at higher levels (i.e., the intentions they realize and the contexts in which they can be successfully deployed) to functional or anatomo-kinematic rules at lower levels (Grafton, 2009). Importantly, in this hierarchical setting, action understanding amounts to inferring the higher levels of the hierarchy (e.g., goals and

intentions) by only having access to movement observations. Action understanding also paves the way to imitating the observed actions (Hurley and Chater, 2005).

The aim of this paper is to advance a unified computational account of the cognitive processing underlying action observation, understanding, learning and imitation abilities. This novel proposal is grounded in the framework of active inference, which provides a normative perspective on brain computations and behaviour (Parr et al., 2022). While our proposal is domain-general, we illustrate it in a generative model of a "tennis task", in which a naive tennis player (henceforth, the learner) infers and imitates the actions executed by an expert player (henceforth, the teacher). Please consider that in our study, the game of tennis is taken as an exemplificative scenario; the naive and expert tennis players do not recognize or reproduce the bodily movements of human tennis players, but only an abstract representation of these movements (i.e., combinations of abstract movement features, see below the model description). While the task structure consists of an abstract, toy-example, the model aims to provide a general and biologically plausible account of action understanding, learning, and imitation. Our simulations will show that the inferential dynamics supporting the recognition, learning and imitation of (tennis) actions, as well as the saccadic movements that promote (active) perception, emerge naturally from the active inference formulation and align well with empirical results and neurobiological models from previous literature.

In the next sections of this paper, we present and discuss four simulations of the tennis task using active inference. Each simulation is designed to illustrate one specific aspect of the framework that is relevant to explain how we understand and imitate actions. The first simulation illustrates the core mechanisms of active inference in play during action understanding, highlighting its inferential and active nature. This simulation shows the empirical predictions (e.g., beliefs updating, oculomotor dynamics and response selection and times) that can be derived by casting action understanding as a process of (active) *inference over a hierarchical generative model* of postures, goals, and intentions. The second simulation illustrates how during the observation of familiar actions, oculomotor dynamics are guided by the *imperative of reducing uncertainty* (about which action one is observing), which is automatically elicited in active inference in ambiguous contexts. This simulation also illustrates the way prior beliefs and the motor knowledge of the observed movements can influence the observation pattern. The third simulation illustrates how during the observation and learning of novel actions, oculomotor dynamics are guided by the *imperative to pursue novelty*, which arises automatically in active inference when the observed actions are unknown and the context affords learning. This simulation also illustrates that the learning process creates a more reliable (likelihood) mapping between lower and higher levels of representation and affords a deep (semantic) level of action understanding. Finally, the fourth simulation illustrates how in active inference, action observation automatically facilitates the execution of imitative actions but interferes with the execution of alternative actions. This is because both action understanding and execution rest on shared (ideomotor) codes across perception and action. Finally, we illustrate the neurobiological implications and the possible neural implementation of our proposed model.

2. Active inference simulations during action understanding and imitation

2.1 Simulation 1: hierarchical inference about actions in a tennis task

The tennis task involves two people: an experienced player (teacher), who demonstrates how to perform different kinds of tennis shots correctly; and a naive player (learner), who observes the teacher's actions and tries to understand and imitate them. The computational model describes formally how the learner performs action understanding and imitation, by observing how the teacher's body movements unfold over time.

As illustrated in Figure 2D, the learner's visual scene encodes the movements of the teacher. The visual elements of the scene are the teacher's *body parts*, and their relative position in a four-slot quadrant. Each combination of body parts corresponds to a specific *posture*. Each posture is univocally defined by the relative position of two body parts. Specifically, posture 1 occurs whenever the *head* is beside the *foot*; posture 2 occurs whenever the *head* is beside the *racket*; and posture 3 occurs whenever the *head* stands diagonally with respect to the *racket*. The body part *hand* covers the two remaining locations in each configuration of postures. The transitions between postures represent the teacher's movements to hit the ball, or her *proximal goals*. Here, the teacher has six possible proximal goals (*right or left forehand, right or left backhand, right or left smash*), each defined by a unique sequence of four postures. For example, the sequence of postures 3-1-3-2 corresponds to the *right backhand*. Finally, proximal goals can be further classified into three *intentions* (*forehand, backhand or smash*) that abstract away from the specifics of (left or right) actions.

Simulated action observation, understanding and imitation

Figure 2. Results of Simulation 1: action observation, belief updates during action understanding and imitative motor response. The Panels A-D correspond to the four hierarchical levels of the model shown in Figure 1. Here, we describe them in the reverse order, starting from the bottom (kinematic observation) level to the top (intention) level. The D panel shows the teacher's sequence of postures during a tennis shot (here, a right backhand). Each posture corresponds to a different combination of body parts. The blue line represents the simulated learner's eye movement trajectory, and the red dots her saccade locations. Note that the learner can only make saccades to one body part at a time but can integrate this information over time, to infer the current posture of the teacher, shown in panel C. This information is represented as probabilistic beliefs: the darker the colour of each cell, the larger the posterior belief about a specific posture. Please note that these (and the other) beliefs change over time, as the learner obtains more information about the teacher's postures. In panel B are shown the dynamics of the probabilistic beliefs associated with the sequences of postures, where each sequence corresponds to one of the proximal goals of the teacher. Panel A shows the probabilistic beliefs associated with the highest level of representation: the intentions. Finally, panel E shows the imitative response executed by the learner where she correctly executes the same sequence of postures that she has observed.

In the model, action understanding is cast as the inference of the three hidden variables (intentions, proximal goals, and postures) based on movement (body part kinematics) observations. In turn, this hierarchical inference can be conceptually divided into three kinds of inferences, see Figure 2. The first kind of inference only regards the most superficial features of action, i.e., the inference of postures from kinematic observations. In the model, the learner can only execute a saccade and observe a single body part of the teacher at a time (i.e., the positioning of head, hand or leg) but she can accumulate this information over time to infer the teacher's posture, defined here as a specific combination of body parts. Figure 2D shows the example of a learner performing saccades (red dots) to different body parts, which are informative about which of three possible postures she is observing. Figure 2C shows the dynamics of the inference about postures in the computational model. The model maintains beliefs (as probability distributions, where darker colours denote higher probabilities) about postures and updates them when it obtains novel information via further saccades. Figure 2B shows the belief dynamics associated to the second kind of inference, which regards a deeper aspect of the observed movements: the one that concern *proximal goals*. This second kind of inference uses the results of the first kind of inference (i.e., the inferred postures) as observations. The specific example illustrated in Figure 2B regards the inference that the teacher's proximal goal is executing a right backhand, based on the fact that the sequence of observed postures is 3-1-3-2. Figure 2A shows the third kind of inference, which regards the most general aspect of the observed movement: the intention. The specific example illustrated in Figure 2A regards the inference of the *backhand* intention. Finally, Figure 2E shows that the learner can imitate the observed action by executing the sequence of postures that she previously inferred (or if needed, another sequence that realizes the same intention). The possibility to imitate observed actions rests on the fact that the model infers (forms beliefs about) the observed sequence of actions or policy. In other words, imitation is simply the enactment of the policy that was inferred with the highest probability. See methods for an explanation of how beliefs about the observed action are translated into a motor response.

2.2 Simulation 2: observation of familiar actions and the role of uncertainty and prior knowledge in active sensing

In the previous simulation, we explained that the learner employed saccades to gather information about the teacher's postures, proximal goals, and intentions, but we did not explain how she selected the next saccades. The second simulation aims to show that in active inference, action understanding is achieved by engaging sequences of eye movements that gather *salient* information (Donnarumma et al., 2017; Friston et al., 2015, 2012; Parr and Friston, 2017). Salience is determined by a prediction on where to perform the next fixation to help reducing uncertainty about the observed action. This is in keeping with theories of perception as an active process, which involves planning sequences of (oculomotor) actions that lead to optimal information foraging and proactively gathers information to reduce uncertainty (Blake and Yuille, 1992; O'Regan, 2011;

Tsotsos, 2011). This uncertainty-reduction imperative is one of the components of the (expected free energy) equation that guides the selection of policies in active inference; see the Methods for details. This imperative becomes more prominent when there is some uncertainty to resolve, such as when the learner does not know which action she is observing, but loses importance if the learner has prior knowledge about the action – unless this prior knowledge is misleading.

Consider a learner who is observing the teacher executing a *right smash* shot in two different conditions (see Figure 3). In one condition, illustrated in the left panels of Figure 3, the learner has no prior knowledge about the teacher's goal (i.e., she has a flat prior belief). In the other condition, illustrated in the right panels of Figure 3, the learner has been informed that the teacher will perform a *left forehand* shot (i.e., she has a misleading prior belief). In both cases, eye movements are engaged to solve the uncertain aspects of the action as they are directed towards the most salient location of the scene, which tests the learner's hypotheses. However, what is deemed to be the most salient location to gaze depends on the learner's hypotheses or prior beliefs (which are of course updated after gathering novel observations). As an effect, the patterns of eye movements and belief updating of left panels of Figure 3 and the right panels of Figure 3 are different. In both cases, the learner eventually recognizes the correct proximal goal (*right smash*); however, this process is slightly slower in the latter case, when she starts from a misleading prior. This can be observed by looking at the simulated reaction times (see the Methods), below the D panels.

The role of uncertainty and prior knowledge

Figure 3. Results of Simulation 2: the dynamics of action understanding of two learners, with no prior knowledge (left learner) or incorrect prior knowledge (right learner). The **D** panels show the different

oculomotor dynamics for the two cases, respectively. The A, B and C panels show the different belief dynamics for the two cases, with the same notation as in Figure 2. Note that at timestep 3, the right learner (with misleading priors) strongly believes to observe a left forehand and posture 2. However, after inferring she is observing posture 1, she changes her mind and correctly infers she is observing a right smash. On the contrary, at timestep 3, the left learner (with flat priors) assigns the same probability to the left forehand and the right smash and to the two possible incoming postures (1 and 2). After observing posture 1, she correctly infers she is observing a right smash, without significant surprise. Although the right one employed more time to correctly categorize the observed action, both learners correctly imitate the action, see panel E.

This simulation permits appreciating how active perception and the dynamics of belief updating - and hence whether recognition will be successful or unsuccessful, fast or slow - depend on the learner's prior beliefs. This is because the prior beliefs determine the salience of visual locations and which location to gaze next. In this perspective, salience is not associated with basic properties of the stimulus such as noise, but it is defined formally as an information gain - and is therefore high in locations that have a high potential to resolve uncertainty (Donnarumma et al., 2017; Friston et al., 2015, 2012). Specifically, salience can be formalized as a Bayesian surprise, which corresponds to the divergence between the prediction of a generative model and the actual outcome. Before a saccade, the learner can use her generative model to estimate Bayesian surprise and hence the salience of each location - and then select the highest salience location (Itti et al., 1998; Itti and Baldi, 2009; Itti and Koch, 2000). Using this method, a salient location corresponds to the one expected to reduce uncertainty, if a saccade were performed towards it (Parr and Friston, 2017). Empirical research shows that this notion of Bayesian surprise characterizes well which aspects of a visual scene capture human attention (Itti and Baldi, 2009; Mirza et al., 2016; Posner and Petersen, 1990).

Another important point to notice is that prior information can come from different sources: it can be provided exogenously (for example, via verbal or contextual cues) or be generated endogenously, on the basis of the learner's own motor knowledge. A consistent body of evidence has shown that eye movements are engaged coherently during both action execution and observation, suggesting that they come from the same generative model that encodes one's own motor knowledge (Flanagan and Johansson, 2003). More broadly, this suggests that the parameters employed for categorizing and encoding action representations are shared with those for action execution - which is a feature implemented in our model (see the Methods for a description of the model parametrization). This feature has been widely studied in relation to the mirror system (Aglioti et al., 2008; Bonini, 2016; Calvo-Merino et al., 2006, 2005; Cross et al., 2006; Maranesi et al., 2019, 2014; Pedullà et al., 2020; Pezzulo, 2013; Rizzolatti and Craighero, 2004). Later, we will elaborate further on this topic.

2.3 Simulation 3: observation of novel actions: the role of novelty in active sensing and learning

So far, we assumed that the learner already knows the teacher's actions that she observes. This simulation illustrates what happens when the learner firstly observes novel postures and then learns a never-observed proximal goal.

There is a fundamental difference between the (active) perception of known and unknown actions. As discussed above, the perception of known actions is directed towards *salient* locations that reduce uncertainty about what action a person is observing. In turn, to calculate saliency, the learner needs to be endowed with a generative model that permits predicting “what would I see if I observe, in a specific moment, that location of the teacher’s body?” By definition, during the observation of unknown actions, the learner's model is incomplete and hence she cannot make reliable action predictions that are necessary to assess what is salient. In this case then, her oculomotor behaviour is mainly driven by the *novelty* of sensory evidence rather than saliency or Bayesian surprise. Observing unknown actions induces a novelty-seeking, curious behaviour that is mainly driven by bottom-up processes such as attentional capture (e.g., orienting responses) (Sokolov, 1963) or perceptual curiosity, defined as “the interest in and giving attention to novel perceptual stimulation” (Berlyne, 1954; Collins et al., 2004), as opposed to the prominence of top-down, expectation-guided processes during the pursuit of salient information.

The patterns of oculomotor movements related to salience (Figure 2) and novelty (Figure 4) are significantly different. Visually sampling unknown actions generally requires more saccades because many locations provide novel information, and the learner cannot predict in advance which posture she will observe next; this is because she has no knowledge of how postures are determined in the unknown action. The novelty-seeking behaviour is crucial to resolve the learner's “semantic knowledge gap” regarding the acquisition of the unknown mapping between body parts and postures in the learner's generative model. In Bayesian terms, this unknown mapping corresponds to the likelihood $p(o|s)$ and it regards the relations between hidden states (s) at a higher hierarchical level – here, the level of postures – and the possible outcomes (o) at the level below – here, the body part kinematics. As illustrated in Figure 4, sampling novel locations permits the learner to fill in this knowledge gap, and to learn the likelihood mapping between body parts and their combinations (Baranes et al., 2015; Barto et al., 2013; Gottlieb and Oudeyer, 2018; Schwartenbeck et al., 2019).

Observation of novel actions

Figure 4. Results of Simulation 3: observation and imitation of a novel action. In panel **D** are shown the simulated oculomotor dynamics when the learner observes an unknown action, here corresponding to the sequence of postures 1-3-2-1. In this case, visual attention is driven by both surprise and novelty. The latter induces the sampling of several novel locations, resulting in a larger number of saccades compared to previous cases, in which the observed actions were known. In panel **C** and **B**, it can be noted how the probabilistic beliefs are more spread across the hidden states resulting in much larger amount of uncertainty. This leads the student to an “undecided” intention categorization in panel **A**.

Figure 5 illustrates schematically how the novel knowledge that the agent acquired through novelty sampling permits it to learn the likelihood mappings of the higher hierarchical levels of the generative model. Specifically, the likelihood mapping illustrated in Figure 5 encodes the probabilistic relations between the levels of the postures and the proximal goals. Figure 5A shows the real statistics of the environment (called *generative process* in active inference), which in our example corresponds to the real and objective mapping between proximal goals and postures. Each row of the panel shows the sequence of postures expected under a particular proximal goal: the first three columns represent the postures (1 to 3) at the quadrant location 1; the second three columns represent the postures at the quadrant location 2; the third three columns represent

the postures at the quadrant location 3; and the last four columns represent the postures at the quadrant location 4. For example, the fourth row shows that performing a *left backhand* elicits the sequence of postures 3-2-1-3 (these postures are marked in black, which correspond to the fact that they have a high probability). Figure 5B shows the learner's *generative model* before learning (to put it simply, her knowledge about the statistics of the environment). This generative model is analogous to the generative process, but the row corresponding to the *left backhand* is grey, which indicates that the learner has no knowledge about the sequence of posture elicited by this proximal goal (technically, she has a flat belief state). The (likelihood) mapping between the *left backhand* goal and the corresponding sequence of postures is the “semantic knowledge gap” that the learner has to fill by learning. Finally, Figure 5C shows that after some trials, the learner has (almost completely) filled in this gap as she learned a sufficiently reliable (likelihood) mapping between the *left backhand* goal and the corresponding sequence of postures.

Crucially, we assume that the quality of the likelihood mapping of the learner's generative model determines her semantic understanding of the observed actions. Specifically, a strong likelihood mapping between two levels of the action representation hierarchy affords to a deep, semantic understanding of the action meaning. Conversely, a weak (or flat) likelihood mapping cannot afford semantic understanding. This is because the likelihood mapping establishes a link between the different levels of action description in the hierarchy of Figure 1. Learning a likelihood mapping creates a link between two representation levels, therefore expanding the learner's action vocabulary with novel conceptual representations. Furthermore, and importantly, this learning process grounds the novel action representations at higher hierarchical levels into sensorimotor experiences at lower hierarchical levels – which is key for semantic understanding within theories of embodied and grounded cognition (Barsalou, 2008; Pezzulo et al., 2011; Rumiati et al., 2010; van Elk et al., 2014). The embodied approach argues that processing conceptual representations requires eliciting sensorimotor experiences, for example, by simulating them (Von Hofsten, 2007). Motor representations are then automatically elicited when we activate conceptual representations (Aziz-Zadeh and Damasio, 2008; Galetzka, 2017; Meteyard et al., 2012), as shown in the context of language (Borghi et al., 2019; Borghi and Binkofski, 2014; Borghi and Cimatti, 2010; Davis and Yee, 2021; Hauk et al., 2004). Note that here the term “sensorimotor experiences” should be intended broadly, to encompass multiple exteroceptive, proprioceptive and interoceptive modalities – all of which can become linked to action representation during the learning of likelihood mappings.

Learning action representations

Figure 5. Results of Simulation 3: how the generative model is updated when a novel action is learned. In this example, the part of the generative model that is updated is the (likelihood) mapping between postures and goals. For each panel (A, B, C), the rows represent the goals and the columns represent the three postures for each of the four quadrant locations. In each square, a darker tint represents a larger probabilistic contingency between the states (goals) and the observations (postures). The panel A shows the contingencies of the so-called generative process, which correspond to the true statistics of the environment. The generative process represents the real probability of observing a posture given a goal. In panel B is represented the likelihood of the learner's generative model before learning. In the generative process the goal left backhand corresponds to the sequence of postures 3-2-1-3. However, before learning, the learner's generative model does not include yet this mapping, but, instead, a "flat" mapping between the left backhand goal and all the possible postures (grey tint, as probabilities are spread). The panel C shows that, after learning from 12 simulation trials, the learner has acquired knowledge about these contingencies and

the likelihood mapping in the generative model is almost completely aligned with that of the generative process (tints are close to black and white).

In summary, conceptual-semantic learning is the process that allows a novel, meaningless action to be associated with semantic meaning by establishing links between action representations across hierarchical levels, such as between conceptual representations of the action vocabulary and sensorimotor processes. This perspective implies that a novel observed action has no semantic meaning until the connections between higher and lower levels of description of the action (or likelihood mappings) are established. Until semantic meaning (in the sense described above) is established, the learner will be able to discriminate between actions at the kinematic level, but not to understand their meaning or the underlying intention – as these actions are not present in her action vocabulary, nor they are linked to any existing intention. In this condition, their representation is supposed to be episodic in memory but not semantic (Tessari et al., 2006). This distinction between a shallow level of action understanding that only involves its kinematics and a deeper level that also involves proximal goals and intentions is key when the representation is used to execute an action such as during imitation, as we illustrate in the next simulation.

2.4 Simulation 4: facilitation and interference between observed and executed actions

There is increasing evidence that the neuronal underpinnings of action execution and observation are largely shared, at least in humans and other primates (but also possibly in other animals). This was most iconically demonstrated by the discovery of mirror neurons: a set of neurons (originally found in the pre-motor area F5 of the macaque, di Pellegrino et al., 1992) that discharge both during the observation of an action directed to an object and its execution. Mirror neurons are part of a much wider action-observation network (AON) constituted by three bilateral cortical areas reciprocally connected: the ventral premotor cortex, the inferior parietal lobule, and the superior temporal sulcus (Rizzolatti and Craighero, 2004). Further studies have suggested that the engagement of the mirror system depends on the observer's motor knowledge (Aglioti et al., 2008; Calvo-Merino et al., 2006, 2005; Cross et al., 2006; Pedullà et al., 2020). To explain the above findings, various researchers have proposed that action observation and understanding are based on motor resonance and the automatic activation of the motor system in the observer's brain (Rizzolatti et al., 2001; Rizzolatti and Sinigaglia, 2010); on a “motor simulation” of the observed action in one's own nervous system (Aglioti et al., 2008; Calvo-Merino et al., 2006, 2005; Cross et al., 2006; Gergely and Csibra, 2003; Grafton, 2009; Jeannerod, 1994; Pedullà et al., 2020); or on shared (ideomotor) neural codes for perception and action (Hommel et al., 2001; James, 1890).

Another convergent line of research has shown that action observation and execution reciprocally and continuously influence each other: observing an action can produce an automatic visuomotor priming (Brass et al., 2001) and the automatic imitation of the same action (Heyes, 2011; Tessari et al., 2002) but interferes

with the simultaneous execution of different actions (Kilner et al., 2003). Observing an action compatible with the performed movement (e.g. lifting the index finger while observing an index finger move upwards) facilitates reaction times, whereas observing a movement incompatible with the performed one slows down reaction times and accuracy (e.g. lifting the index finger while observing a finger press; (Brass et al., 2001) – possibly, because both actions need to be simultaneously encoded in the mirror system (Catmur et al., 2009).

In this simulation, we illustrate how the model accounts for automatic facilitation and interference effects, given that its representations of (beliefs about) actions are shared across perception and action in the model hierarchy. We compare two cases: in the first case (left panels of Figure 6), the learner is instructed to observe an action (left smash) and then to produce an imitative response. In the second case (right panels of Figure 6), the learner is instructed to observe the left smash action but to produce an incongruent action (a left forearm).

Motor priming and interference

Figure 6. Results of Simulation 4: the dynamics of motor priming (left panels) and interference (right panel) for two learners. The left learner observes a left smash and imitates the same action (motor priming). The right learner observes a right forearm but executes a left smash (motor interference). The discrepancy can be observed in the less precise “target postures” in the E panel to the right, where the grey tint indicates candidate target postures that corresponds to the observed sequence 1-3-2-3.

Figure 6 shows that observing an action facilitates the execution of congruent, imitative responses (left panels) but interferes with the execution of incongruent actions (right panels) as shown empirically. This

becomes evident by considering that during execution of the same (left smash) action, the beliefs about policies (see the Methods) are more uncertain in the E panel on the right compared to the E panel on the left.

2.5 Summary and discussion of the simulations

The simulations described in this section help illustrating the key mechanisms of active inference that are relevant to model how we understand and imitate actions. Here, we summarize them by highlighting three main points. First, action observation and understanding can be cast as inferential processes that engage active information gathering strategies (see Simulation 1). Understanding an action requires a hierarchical generative model to produce a cascade of perceptual hypotheses and predictions (about what action one is currently observing, the sequence of motor acts and the ensuing observations) and active sampling strategies to test the predictions by selecting the most salient observations. The fact that an action repertoire is hierarchically arranged is a common assumption across theoretical and neurobiological studies (Grafton and de C. Hamilton, 2007; Jacob and Jeannerod, 2005; Pacherie, 2008). Here, we move from theoretical descriptions of hierarchies to a fully implemented computational model. Our proposal emphasizes that the hierarchical action representation constitutes a generative model that supports inferential dynamics and the active selection of relevant information. This perspective therefore aligns action understanding to a broader view of predictive coding and inference in the brain (Friston, 2008, 2005; Friston et al., 2017; Parr et al., 2022; Pezzulo et al., 2021, 2018).

Second, there is a fundamental difference between oculomotor dynamics during the observation of familiar actions (Simulation 1 and 2) versus novel actions (Simulation 3). In the former case, when the observed action is already part of the learner's action repertoire (and generative model), oculomotor strategies are guided by the imperative of reducing uncertainty about what action one is observing. In the latter case, when the action is not part of the learner's action repertoire (and generative model), oculomotor strategies are guided by novelty-seeking. Importantly, the two imperatives of uncertainty-reduction and novelty-seeking emerge naturally from the active inference framework (and in particular from the expected free energy equation, see the Methods). In both cases, the observation policy maximizes information gain; however, the type of uncertainty is different during the observation of familiar and novel actions. In the former case, the model needs to resolve uncertainty about the hidden states causing the observations (this is called *hidden state exploration* and is driven by salience) whereas in the latter case the model needs to establish how hidden states generate outcomes driven by uncertainty about the model parameters (this is called *parameter exploration* and is driven by novelty). Notably, these two imperatives prescribe two different behaviours (Schwartenbeck et al., 2019). Hidden state exploration predicts that the learner actively seeks *salient* observations that allow for unambiguously infer the hidden states that generates the outcomes; whereas parameter exploration predicts that the learner actively seeks for *novel* combinations of hidden states and outcomes, because this determines the learning of how outcomes are generated. It is only via this novelty-seeking behaviour that the learner can acquire knowledge about the relationships between different levels of

action representation and between hidden states and observations (for novel actions), hence filling the gaps in her generative model. This is important insofar we assume that the meaning in action semantics emerges from pursuing novelty across the multiple hierarchical levels of action representation. The learner can still track novel actions even without having an integrated action representation, but this would not correspond in our account to a deep semantic representation.

Third, in the proposed model, hierarchical action representations and internal models are shared across action execution and perception. The sharing can prime the execution of actions that are congruent with an observed action and interfere with incongruent actions. This is important insofar it highlights the importance of engaging one's action repertoire during the observation of actions executed by others, as highlighted by a large body of neurobiological evidence - as we will elaborate below.

3. Neurobiological underpinnings of the proposed model

So far, we described the functioning of the hierarchical active inference model but we disregarded its relations with neurobiological findings about action execution and understanding. Here, we briefly discuss the ways the hierarchical active inference model links to neurobiological findings about the functioning of dorsal and ventral visual streams and the action observation network.

3.1 Mapping the model's computational strategies into the multiple-routes model for action understanding

Simulations 1, 2 and 3 suggest that the visuomotor system recapitulates and solves at the neurocognitive level the trade-off between engaging explorative saccades and goal-directed saccades (see the Methods for the rationale). Furthermore, the visuomotor system solves the trade-off between salience and novelty driven exploratory behaviour, particular relevant for modulating learning. To recap, both novelty and salience (as Bayesian surprise) elicit attentional capture and arousal that direct appropriate responses and enhance learning. Even though they are usually confounded they are distinct concepts: while novelty is associated with the acquisition of representations (as we showed), surprise relates to the improvement of predictions. Indeed, novelty entails a divergence from memory while surprise entails a divergence from expectations (Barto et al., 2013). Thus, the fundamental prediction of our formalization is that during action observation, if the context allows for learning, observation will be driven by both novelty and surprise. On the opposite side, when observation is restricted on inference about movement kinematics it will be driven only by surprise. We suggest that surprise and novelty might recruit partially distinct neuronal circuits and computational strategies. A strategy named *parameter exploration* (Schwartenbeck et al., 2019) is driven by visual novelty and is engaged to obtain meaning about actions, which translation is associated with skilful motor control; a strategy named *hidden states exploration* is driven by salience (Bayesian surprise) and is engaged to recognize the kinematic aspect of actions and the motor translation is associated with short-term, online, sensorimotor, and flexible responses .

This description fits well with multiple-routes models for action understanding, tool-use and imitation. A two-routes cognitive model (initially proposed by (Rothi et al., 1991), and later developed by others (e.g., Buxbaum & Randerath, 2018; Cubelli et al., 2000; Rumiati & Tessari, 2002; Tessari et al., 2007); based on the study of apraxic patients, proposes that the visuomotor system, during the translation of the visual input (gestures or objects) into a motor program (the corresponding action), can follow a *semantic route* recalling conceptual and semantic representations of action, or a *direct route* that bypasses semantic information transforming the visual input into a motor act. Here, we propose a parallelism: within the semantic route, actions are observed according to salience and novelty driven visual attention and then imitated according to procedural motor control (see the Methods for an illustration of how motor control is implemented in the model). On the other hand, within the direct route actions are observed according only to salience driven attention and imitated according to goal-directed motor control. Therefore, we suggest that the two routes may encode two distinct cognitive and behavioural strategies that correspond to the salience driven hidden state's exploration over movements with goal-directed control for the direct route and both novelty and salience driven parameters exploration with procedural control for the semantic route. Furthermore, the semantic encoding is hierarchical in terms of cognitive representation, which is also reflected in the neuroanatomy as it will be explained soon. Our prediction is supported by studies in the context of imitation where, when the context is not stable enough, the cognitive strategy for action recognition and imitation is empowering the direct route with surprise driven exploration (*hidden state exploration*) even when meaningful actions are displayed (see (Tessari et al., 2022, 2021, 2007, 2006; Tessari and Rumiati, 2004). The alternative strategy (*parameter exploration*) is engaged when there are enough stable regularities in the world to be learned. Therefore, surprise driven hidden state exploration will be useful in conditions where behaviour must be adapted to trial-by-trial changes.

Similarly, our account of the two strategies can be associated with the two-action system model for encoding object affordances in the context of tool use: one pathway entails affordances based on objects structure and is specialized for visuomotor interactions with objects based on currently observed visual information, such as shape, size, and location that are constantly updated and processed online (i.e., the *Structure system*). The other pathway entails their functional manipulation and relies upon long-term, conceptual representations and extracts the features of the action that remain constant across occurrences (i.e., the *Function system*; Binkofski & Buxbaum, 2013; Jax & Buxbaum, 2010). Once again, at the computational level, the Structure system resembles the idea of surprise driven hidden state exploration and goal-directed control. The Function system resembles what we have proposed for novelty driven parameter exploration, as it permits extracting the constant features of the action and learning the contingencies between kinematic and higher-level, semantic aspects.

The cognitive processes underlying the two routes have been extensively investigated in several neuropsychological studies with brain-damaged patients with a focus on the neural and cognitive correlates of imitation of familiar and novel actions (Achilles et al., 2019, 2016; Bartolo et al., 2001; Cubelli et al., 2000; Goldenberg and Hagmann, 1997; Mengotti et al., 2013; Peigneux et al., 2004; Rumiati et al., 2005;

Tessari et al., 2007) and support a network in the left hemisphere: lesions of the ventro-dorsal stream (from medial superior temporal area, MT/MST, to the inferior parietal lobule, and then to the ventral premotor cortex) produce impairments to more conceptual aspects of action representation, such as skilled use and pantomime of objects (e.g., (Martin et al., 2016a, 2016b; Tessari et al., 2021, 2007). On the contrary, the direct route and the processing of new movements have been associated with the dorso-dorsal stream (from V3a to V6 to V6a, to the superior parietal lobule, and then to the dorsal pre-motor areas (Binkofski and Buxbaum, 2013; Hoeren et al., 2014; Martin et al., 2016a; Mengotti et al., 2015; Tessari et al., 2021, 2007). At last, the processing of known gestures and the semantic route have been related to regions belonging to both the ventral (from V2 and V4 to the posterior inferotemporal, the central and the anterior inferotemporal areas) and the ventro-dorsal streams (Hoeren et al., 2014; Martin et al., 2016a, 2016b; Rijntjes et al., 2012; Rumiaty et al., 2005; Tessari et al., 2021; Weiller et al., 2011, 2009), suggesting that the ventral stream might decode the meaning of a movement (and intransitive gestures particularly), and the ventro-dorsal stream the tool-related, meaningful gestures. The dorso-ventral structures also allow to infer the possible functions from structure (Binkofski and Buxbaum, 2013; Goldenberg and Hagmann, 1998; Heilman et al., 1997) and can discover the alternative functions of familiar tools (Heilman et al., 1997; Roy and Square, 1985). This ability recalls the parameter exploration strategy we have argued to emerge during learning.

In sum, we propose that there is a correspondence between the salience driven low-level hidden state exploration plus the goal-directed control (at the computational-cognitive level of description), the direct route (at the algorithmic level) and the dorso-dorsal route (at the anatomical level). On the other hand, we propose a correspondence between salience and novelty driven parameter exploration plus procedural control (at the computational-cognitive level), the semantic route (at the algorithmic level) and the ventral and ventro-dorsal streams (at the anatomical level). Finally, higher semantic aspects, such as abstract and symbolic representations, can be associated with the ventral stream. See Figure 7 for a graphical illustration of this proposed taxonomy.

Figure 7. A putative neural implementation of the proposed hierarchical model (this is a modified version of the neuro-cognitive model proposed by Tessari, Proietti & Rumiati, 2022, on Cognitive Neuropsychology). Visuomotor cognition is distributed onto multiple routes in the dorsal and ventral pathways. According to our proposal, the dorso-dorsal stream is mainly dedicated to lower-level aspects of action representations, the dorso-ventral pathway affords more general aspects of action semantics and the ventral stream is dedicated to the most abstract features. Please see the main text for explanation.

3.2 The action observation network

In simulation 4, we investigated the interactions between action observation and execution. It has been proposed that these two processes reciprocally and continuously influence each other, since motor and visual knowledge share a common representation (Gazzola and Keysers, 2009). In our model, understanding and execution are influenced by a common memory parameter, as shown in simulation 2 (and in details in the Methods), where prior experience (including motor knowledge) can influence eye movements at the observation level.

At a neurophysiological level, the mirror neuron system (a cortical circuit connecting ventral premotor cortex, inferior parietal lobule and superior temporal sulcus) was associated with understanding the meaning of the observed (object-oriented) action by extracting and representing its meaning and goal (Fogassi et al., 2005; Rizzolatti et al., 2001; Rizzolatti and Craighero, 2004). This has led to the *direct-matching hypothesis* (Iacoboni et al., 1999) that argues that action understanding occurs when the visual representation of the

observed action is mapped onto the observer’s own motor representation of the same action. More recently, a predictive coding account of the mirroring phenomenon has been advanced: it has been proposed that the mirror neuron system follows a predictive coding scheme in relation to action observation where the most likely cause of the observed movement is inferred by minimizing prediction error in the cortical hierarchy of the action-observation network (AON; (Friston et al., 2011; Kilner et al., 2007)). Generally speaking, on the predictive coding account, what the visuomotor system does during action observation is resolving the inverse (inference) problem created by the one-to-many mapping of action-to-goals. For this, the visuomotor system represents the best explanation of the observed action as a generative model ($p(\text{goal} \mid \text{movement})$), where the AON predicts the sensory consequences of what would be the most likely set of movements to be executed to achieve that goal. The predictive coding scheme supports the continuous comparison between the predicted sensory information and the actual sensory input, and ultimately the inference of the action goal (when prediction error is minimized). A crucial assumption is that mirror neurons discharge during the observation of an action because they are part of the same generative model that is predicting the sensory input related to that action (Friston et al., 2011; Kilner et al., 2007). Hence, the functional role associated with the mirror neuron system is predictive action monitoring. Motor resonance is not achieved by a direct matching mechanism but by emulative action inference or motor, embodied, simulation (Csibra and Gergely, 2007; Gallese and Goldman, 1998; Jeannerod, 2001, 2006; Pezzulo et al., 2013a). Please note that the predictive coding formalization is very similar to the Bayesian belief updating method adopted in our model, with two main differences. First, our model is formulated in discrete time, whereas predictive coding is formulated in continuous time – which implies not just a formal difference but also a different form for the “neuronal message passing” in the two schemes (see Parr et al., 2022 for details). Second, and importantly, here we use active inference, which extends predictive coding to also cover action dynamics. In the model proposed here, the generative model (putatively corresponding to the AON) supports oculomotor control in addition to movement prediction. In other words, our model would imply the AON and the mirror neuron system in the active perception of observed movements (Donnarumma et al., 2017) – a hypothesis that remains to be tested in future studies.

4. General discussion

We have showed how action observation, understanding and execution can be represented in a hierarchical active inference model to illustrate a plausible way of how the visuomotor system may compute visual information and translate it into a motor program.

Our model highlights that perception can be treated as an inferential process based on a hierarchical generative model; and that it is based on active information sampling strategies. Indeed, in our proposed model, building a perceptual representation requires an active process of selecting salient data to test the expectations of an internal generative model – and how this saliency can be influenced by prior information. Furthermore, our model highlights a fundamental difference between oculomotor strategies that are guided

by uncertainty reduction and prior knowledge during the observation of familiar actions, versus novelty-seeking during the observation of novel actions. Crucially, the pursuit of novelty generates exploratory behaviour that permits acquiring knowledge about the relationships between the several levels of action representation. We argue that this novelty-based learning process is key to acquire meaningful action semantics. Another relevant aspect of our model is that it reuses the same internal models and codes across action prediction, understanding and response preparation. The sharing of common models and codes produces mirror responses, priming or interference effects, depending on the congruency between observed and executed actions. Finally, we discussed the putative neuroanatomical correlates of key mechanisms of the model. We highlighted that the different strategies of visual exploration and motor control used by the model could map to different routes of visuomotor cognition within the dorsal and ventral pathways.

Interestingly, most of the appealing features of our simulations stem directly from active inference, which provides a general framework to model the functioning of control hierarchies, in which the different hierarchical levels are not modularized, but rather influence each other reciprocally (Pezzulo et al., 2018, 2015). While the idea that action perception, understanding and imitation exploit hierarchical representations is not novel (Grafton and de C. Hamilton, 2007; Pacherie, 2008), the computational implementation offered here could be used to produce more specific, quantitative predictions. One set of predictions regards the neuronal dynamics that we would expect to observe during action observation and imitation, if they correspond to (active) inferential processes, as assumed here. The belief dynamics plotted in Figures 2-6 can be easily mapped into (simulated) neuronal population dynamics, using the methods illustrated by Friston et al., 2017. More broadly, the model proposed here would suggest that during action observation, we should observe a hierarchy of action predictions, as it has been recently shown during natural language comprehension (Heilbron et al., 2022).

Another set of predictions regards the unique assumptions that our model makes about the structure of the generative model – and the ways messages are propagated from one level to the others during inference – which distinguish it from previous hierarchical and deep active inference models used to address (for example) reading dynamics (Friston et al., 2018). For example, a novelty of our proposal is the fact that the Action Execution level (that is responsible for imitative responses in our simulations) is simultaneously influenced by the Action Understanding and the Action Observation levels, via different pathways; see the methods and Figure 8 for a clarification of the mutual dependencies between levels. This organization of the model provides a formal perspective that permits understanding several experimental findings and formulating novel empirical predictions, which could be tested experimentally in future studies. For example, at the behavioural level, this model predicts that during the semantic learning of an action, humans will vary their pattern of observation according to the amount of stimulus novelty and salience that drives the sampling of new saccade location. This highlights the different contribution of novelty and surprise to attentional capture (Barto et al., 2013). Furthermore, we extended the notion that cues and expectations can alter motor behaviour by showing that these alterations (e.g., Jacquet et al., 2012) can be level-of-representation-specific, as in the case of affordances (Sciulli et al., in preparation). This may suggest further

investigations of mirroring or motor resonance, which seems to be stronger at lower levels of representations (e.g., object oriented actions) and to disappear as the action becomes more symbolic and abstracted (Csibra and Gergely, 2007). The systematic testing of the model predictions in experimental conditions is an objective that we intend to pursue in future research.

Finally, the present model can support neuropsychological studies of patients, following the computational phenotyping approach (Friston et al., 2014; Krakauer and Shadmehr, 2007; Schwartenbeck and Friston, 2016). This approach permits simulating a “digital version” of the patient engaged in a particular task, by providing the actions and observations that the patient experiences in the task as inputs to the model, and 'inverting' the model so that it recovers the parameters (e.g., prior beliefs) that better fit the patient data. This approach permits understanding, for example, what specific parameters of the model better permit to account for specific neuropsychological disorders, see Proietti et al., (in preparation) for a development in this direction. Finally, while the goal of the present papers was to propose a general formal (active inference) scheme for visuomotor cognition, the proposed model could be further elaborated by including more realistic and detailed motor control models that permit simulating sophisticated behavioural tasks, beyond the simple tennis example used here by associating the kinematic level with a generative model for continuous variables, as shown in Parr et al., 2021 and Parr & Friston, 2018.

5. Methods

In this section, we provide a summary description of the active inference framework and introduce the hierarchical generative model that we used to realize the simulations of this paper.

5.1 Active inference under the free energy principle

Active inference is a recent formal framework to explain cognition and behaviour that is gaining prominence across multiple domains of cognitive science and neuroscience (Parr et al., 2022; Pezzulo et al., 2015). Active inference derives from the *free energy principle*, which describes in statistical and information-theoretical terms how a biological agent restricts itself within a limited set of sensory states to avoid the natural tendency to reach thermodynamic equilibrium and, as a consequence, death (Friston, 2010). A key assumption of active inference is that to avoid surprising and potentially disruptive sensory states, the agent must encode and update a probabilistic generative model, or a probabilistic mapping between sensory states and the present and future observations that they generate. The generative model serves to implement approximate (variational) free energy minimization that - in active inference - implements both perception and action planning. Specifically, in addition to the generative model, variational inference requires an auxiliary distribution (aka recognition density q) that encodes an approximate posterior over the agent's beliefs; see (Beal, 2003) used to minimize approximate surprisal in the present (variational free energy) and in the future (expected free energy).

5.2 Variational free energy (F)

Variational free energy is minimized to perform perception and perceptual learning, by estimating posterior beliefs about hidden states. The equation below shows how the *variational free energy* (F) for a given policy π (sequence of actions) is calculated. Variational free energy can be decomposed into two terms. The former is a *complexity* term that scores a (Kullback Leibler or KL) divergence between the posterior beliefs about states of the auxiliary distribution ($q(s|\pi)$), and the posterior beliefs about states of the generative model ($p(s|\pi)$). The latter is an *accuracy*, which scores (logarithm of the) the probability of observations given model beliefs about states ($\ln p(o|s)$). These two terms jointly ensure that the agent engages in a continuous *perception-action* cycle; namely, it updates its (posterior) beliefs about the states represented on the model to better fit its observations; and selects courses of actions that actively sample the observations predicted by the model; please see Friston et al., 2017 for a derivation of variational free energy and technical details.

$$\underbrace{F(\pi)}_{\text{Variational FE}} = \underbrace{D_{KL} [q(s|\pi) || p(s|\pi)]}_{\text{Complexity}} - \underbrace{E_{q(s|\pi)}[\ln p(o|s)]}_{\text{Accuracy}}$$

5.3 Expected free energy (G)

Active inference can also be used to model *planning*, which corresponds to selecting a course of actions that minimizes the agent's free energy in the long run. Planning consists in inferring the variables upon which the states that can be controlled are conditioned on. This requires calculating and minimizing another formal quantity, the *expected free energy* associated to each policy π . Once the expected free energy of each policy is calculated, it is translated into a prior value of policies $p(\pi)$ through a softmax function. Using this prior value for action selection ensures that the selected policies are those expected to minimize more free energy in the long run.

As shown in the equation below,

$$\underbrace{G(\pi)}_{\text{Expected FE}} = \underbrace{D_{KL} [q(o|\pi) || p(o)]}_{\text{Risk}} + \underbrace{E_{q(s|\pi)}[H[p(o|s)]]}_{\text{Expected ambiguity}}$$

expected free energy (G) can be decomposed in two terms. The former is *risk*, which scores the (KL) divergence between the probability of outcomes given a policy ($q(o|\pi)$) and the distribution of preferred outcomes in the agent's generative model ($p(o)$). The latter is an expected *ambiguity*, which scores the expected value of the entropy (H) of the model's likelihood function ($p(o|s)$). Taken together, these two terms ensure that the agent plans adaptively and balances exploitation (or utility maximization) and exploration (or

information gain). Minimizing risk leads to the exploitation of preferred observations and determines goal-directed behaviour. This is because by minimizing the first term, the agent comes closer to its preferences or goals. Minimizing ambiguity generates instead exploratory, information-seeking behaviour that aims to minimize the uncertainty about the states of the world.

In sum, expected free energy drives an agent to maximize utilitarian behaviour and explore all options in a way that confers an exploratory aspect upon behaviour. Expected free energy represents a single objective function shared by information-seeking behaviour and the achievement of prior preferences. It is worth noting that in most practical settings, when there is uncertainty about what course of actions will be more effective, the agent needs to minimize ambiguity first (to reduce the uncertainty) before being able to minimize risk (to achieve goals). This solution to the problem of balancing exploration and exploitation emerges automatically in active inference (Friston et al., 2015).

5.4 Policy selection

Once expected free energy has been calculated, the policy is selected according to the following equation, which represents a (prior) distribution over the policies:

$$\pi_0 = \sigma(\ln \mathbf{E} - \gamma \mathbf{G})$$

This distribution comprises two components. The former is a learned prior over the policies (an $\mathbf{E}$ vector), which reflects the number of times a policy has been previously selected and can be considered a habitual component of policy selection. The latter is the expected free energy $\mathbf{G}$ that is calculated anew at each simulation and can be considered a deliberative component of policy selection. The balance between these two terms is determined by a precision term (γ), which encodes the confidence of beliefs about $\mathbf{G}$. Finally, the equation includes σ , which is a normalized exponential (softmax) function that normalizes its values to ensure that it is a probability distribution.

5.5 The hierarchical generative model used in the simulations

Active inference is a generic framework to model cognitive phenomena but realizing each set of simulations requires designing a specific generative model. To realize the simulations reported in this article, we used the hierarchical generative model shown in Figure 8. The hierarchical arrangement of the generative model shown in Figure 8 reflects – and formalizes – the hierarchical organization of action representations shown in Figure 1; see (Friston et al., 2018; Mirza et al., 2016) for related models in active inference and (Botvinick, 2008; Friston, 2008; Kiebel et al., 2009, 2008; Lee and Mumford, 2003) for a broader view of hierarchical models in the brain.

Figure 8: Schematic of the hierarchical generative model used in this article. (A) The Action Understanding level. At this level, there are two hidden state factors and two outcome modalities. The first factor is the proximal goal performed by the teacher (the sequence of postures) while the second factor is the categorization of the intention. The two outcome modalities are exteroception and feedback. (B) The Action Observation level. At this level, there are two hidden state factors and two outcome modalities. The hidden factors are the three possible postures of the teacher and the four possible eye positions of the agent. The two outcome modalities are exteroception (what) and proprioception (where). (C) The Action Execution level. At this level, there are one hidden factor and one outcome modality. The hidden factor is the posture

executed. The outcome modality is proprioception. The notation used in the generative model follows standard conventions to describe generative models under active inference (Friston et al., 2017). Nodes denote probability distributions and edges denote probabilistic relations between them. The capitalized letters **A-E** denote the matrices of the model: the **A** (likelihood), **B** (transition function), **C** (prior over observations), **D** (prior about hidden states), **E** (prior about policies) matrices. The **G** denotes expected free energy. The lowercase letters **a**, **d** and **e** denote parameters of the respective matrices. The lowercase letters *s*, *o* and π denote hidden states, observations and policies, respectively. The symbol γ is a precision parameter used for policy selection. See the main text and Friston et al., 2018 for technical details.

In the generative model shown in Figure 8, *hidden state factors* correspond to different states of the world, such as the location (where) and category (what) of an element (e.g., the teacher body parts). *Outcome modalities* correspond to the possible observations. Multiple outcome modalities are involved in accounting for parallel sources of sensory input, such as visual inputs and proprioceptive sensations.

The generative model is organized as a *partially observable Markov decision process* (POMDP) that entails two general features. The first is the partial observability, which means that the true states of the world are hidden and can only be inferred by observations that may have a degree of uncertainty. The second feature is the Markov property, according to which beliefs about future states, on which the agent organizes its behaviour, depend only on the states at the current time and not on the past ones (Friston et al., 2017).

Our generative model comprises three levels that are reciprocally connected. The Action Understanding level models the inference of proximal goals from postures and of intentions from proximal goals (hence formalizing the functioning of levels 3 and 4 in Figure 1). The Action Observation level models the recognition of movement kinematics and postures (hence formalizing the functioning of levels 1 and 2 in Figure 1) as well as action observation dynamics, or the active sampling of information. Finally, the Action Execution level models motor responses.

Action recognition exploits the reciprocal connections between the Action Understanding and the Action Observation levels. The hidden states at the Action Understanding level are propagated top-down to influence states transitions at the Action Observation level, by setting their prior beliefs (Friston et al., 2018). In other words, the Action Understanding level generates the hypotheses that are tested by engaging saccades at the Action Observation level. On the other hand, the posterior beliefs about postures at the Action Observation level are propagated bottom-up and become observable outcomes at the Action Understanding level (Friston et al., 2018). This provides bidirectional interactions between variables at different levels of the hierarchy, because the inference of each latent variable is simultaneously influenced by messages passed from different directions.

Please note also that as usual in hierarchical models, state transitions at different levels proceed at different time rates. Specifically, the Action Understanding level proceeds at a slower time rate because to generate an outcome at each time step, it must “wait” for all the outcomes from the Action Observation level. In our

simulations, to categorize a sequence of postures and their goal, the agent employs five timesteps and each of them requires at most four timesteps of posture categorization at the Action Observation level. In addition, the simulated reaction times are obtained by running the *toc* function in Matlab to evaluate the time employed for the categorization of each posture and the associated proximal goal.

The execution of motor actions exploits the reciprocal connections between all the three levels. The translation from observed movements into goals occurs at the Action Understanding level, whereas the translation from goals into procedural motor responses uses the connections between the Action Understanding and Action Execution levels. Finally, goal-directed motor responses occur thanks to the connections between the Action Observation and Action Execution levels. Please see below for a discussion of procedural versus goal-directed pathways in the model. Note that the Action Execution level proceeds at its own pace and it is not temporally linked to the other levels.

The generative model showed in Figure 8 generates outcomes as follows: a policy is selected at the Action Understanding level using a softmax function of expected free energy. The state-transition probabilities (**B** matrices) prescribed by the selected policy determine the sequence of hidden states. These hidden states generate outcomes at this level (through the **A** matrixes) and the initial hidden states at the Action Observation level (through to the **D** vectors). Please note that during inference, new observation can lead to a revision of hidden states and of the expected free energy (via the **C** matrixes), hence influencing policy selection. This means that the agent starts acting driven by a given policy but can select new policies along the way based on new observations.

Computationally, the agent solves the task by engaging in approximate Bayesian inference and (variational) belief updating (Beal, 2003). Variational Bayes is based on introducing an arbitrary distribution called *variational density* or *recognition density* $q(s)$ and in rendering it maximally similar to the true posterior probability of the generative model $p(s|o)$. This is done by assuming a specific form for q (e.g., Gaussian, known as the Laplace assumption, see (Friston et al., 2007) and then optimizing it by the iterative updating of its sufficient statistics through a gradient descent algorithm. Interestingly, active inference assumes that this form of variational inference can be directly mapped into a set of hypotheses about how messages are transmitted at the neural level and used to predict neuronal responses during hidden states estimation and action selection (Friston et al., 2017; Friston and Kiebel, 2009). Under the hypothesis that the brain encodes a generative model of its sensations, and that different groups of neurons encode the beliefs specified in the generative model (e.g., beliefs about intentions, proximal goals and postures), the belief dynamics illustrated in Figures 2-6 can be equated to the dynamics of neural activity of these groups of neurons, which evolve in a way that minimize free energy.

5.6 Action recognition dynamics are driven by salience and/or novelty, in context-dependent ways

Action recognition dynamics are guided by different drives, depending on the context. In the context of familiar, meaningful actions (Simulations 1, 2 and 4) observation policies are driven by the need to reduce

uncertainty about hidden states. This uncertainty is represented by Bayesian surprise as it scores the divergence between the predicted and the actual outcomes. Generally speaking, the model tends to predict familiar sequence of postures that it has seen most often in the past (see below) and that are as encoded into prior beliefs over the initial states $\mathbf{D}$ at the Action Understanding level. Hence, during the recognition of a familiar, meaningful action, high-level prior expectations generate hypotheses about movement kinematics that are verified by engaging epistemic actions – eye movements that sample information to reduce uncertainty. Eye movements are directed towards *salient* locations, which are the locations that the model expects to reduce uncertainty the most, if a saccade were performed towards them (Donnarumma et al., 2017; Friston et al., 2012; Parr and Friston, 2017). In this perspective, salience is defined as information gain, or the potential resolution of uncertainty about hidden states (Friston et al., 2015).

The observation of novel actions (Simulation 3) is substantially different from the observation of familiar actions, as the model has the opportunity to learn novel contingencies between proximal goals and postures. In our model, these contingencies are encoded in the ($\mathbf{a}$ parameters of the) $\mathbf{A}$ matrix that links hidden states (proximal goals) and outcomes (postures) at the Action Understanding level. Crucially, when the model is allowed to learn ($\mathbf{a}$) parameters, the expected free energy equation has to be expanded to include a new (novelty) term:

$$\underbrace{G(\pi)}_{\text{Expected FE}} = \underbrace{D_{KL} [q(o|\pi) || p(o)]}_{\text{Risk}} + \underbrace{E_{q(s|\pi)}[H[p(o|s)]]}_{\text{Expected ambiguity}} - \underbrace{E_{p(o|s)q(s|\pi)}[D_{KL}[q(\mathbf{A}|o, s) || q(\mathbf{A})]]}_{\text{Novelty}}$$

The novelty term in this expanded equation scores how much the beliefs in the $\mathbf{a}$ parameters (mapping states to outcomes) are expected to change after a new observation. In turn, including this new term in the expected free energy ensures that the agent will tend to select policies that pursue novelties. This entails a forms of active learning, during which the active inference agent will preferentially look at (novel) locations that it expects to change the $\mathbf{a}$ parameters the most (Schwartenbeck et al., 2019). In sum, the expanded free energy equation highlights the importance of engaging in two forms of exploratory behaviour, which are guided by salience and novelty, respectively. While pursuing salience is useful to reduce the uncertainty about hidden states (or ambiguity), pursuing novelty is useful to reduce uncertainty about model parameters (or ignorance).

5.7 Two varieties of model learning: model expansion and statistical sequence learning

The model engages in two varieties of learning. The former kind of learning is a model expansion. As discussed above, in Simulation 3, we allowed the model to expand its repertoire of semantic action representations by learning novel contingencies between proximal goals and postures in the $\mathbf{A}$ matrix. This has been done by setting a column that encodes one of the proximal goals (*left backhand*) in the $\mathbf{A}$ matrix as

a uniform distribution. This way, the likelihood mapping for that state entails equal probability for all observations, carrying no information. We allowed learning over 12 consecutive trials, but we prevented the learner from trying to recognize the *left backhand* during the first 11 trials learning. By doing so, we ensured that the learning of the *left backhand* was done in an unsupervised manner, without feedback from correct or incorrect recognition, which would have been misleading in early learning phases; see (Smith et al., 2020) for a similar approach.

The latter kind of learning is a process of statistical sequence learning. When the process of action recognition is repeated for multiple trials, the agent accumulates information over time and acquires familiarity with the observed sequences of postures. Formally, the agent sequentially updates its prior beliefs about sequence of postures, by accumulating Dirichlet distributions in the $\mathbf{d}$ parameter (Beal, 2003; Blei et al., 2003; Friston et al., 2016). This statistical learning process ensures that the agent can assign a higher prior probability to the sequences of postures that were observed more often in the past.

5.8 Motor control and its two components: procedural and goal-directed

In the model, the selection and control of imitative responses is determined by two components. The first component is the posterior belief about the proximal goals, as encoded by the prior beliefs over the policies ($\mathbf{E}$) and the associated $\mathbf{e}$ parameter at the Action Execution level. The $\mathbf{e}$ parameter constitutes a form of procedural motor knowledge that guides action selection in a habitual manner - which means that the agent expects to execute the (next action in the) sequence it has observed more often in the past. The second component that determines action selection is the posterior belief about the (recognized) postures, which – via a form of visuo-spatial working memory (Cai et al., 2018; Logie, 1986) – influences the agent’s prior preferences ($\mathbf{C}$) and hence in turn the expected free energy ($\mathbf{G}$). In short, the agent expects to execute the same action that it has inferred by observing the teacher’s movements. We consider the former a *procedural* memory (Engelkamp and Zimmer, 1989, 1985, 1984) as it uses cached ($\mathbf{E}$) values that are learned over time and the latter a *goal-directed* form of control as it uses expected free energy ($\mathbf{G}$) computations that take the agent’s prior preferences ($\mathbf{C}$) into account.

Which specific action is executed depends on the weight given to the procedural ($\mathbf{E}$) and goal-directed ($\mathbf{G}$) components. The relative influence of the two components is determined by the parameter γ (gamma) that correspond to the precision of beliefs about the expected free energy. This process of precision weighting corresponds to attentional modulation at the cognitive level and to synaptic gain at the neurophysiological level (Clark, 2013; Feldman and Friston, 2010; Parr and Friston, 2017). In our simulations, the modulation of the γ parameter is implicated in the imitative responses for familiar and novel actions. When the learner observes a familiar action, the imitative response is mostly driven by the procedural memory ($\mathbf{E}$) of the action sequence, while goal-directed control generated by $\mathbf{G}$ can be engaged with a lower level of precision, as a form of online monitoring of the movements. Notably, by engaging procedural memories, the model processes an entire action sequence as a whole rather than evaluating its component actions one by one - a

mechanism that has been sometimes called "chunking" (Botvinick et al., 2009; Dezfouli et al., 2014; Dezfouli and Balleine, 2013; Miller et al., 1960). In active inference, this sequential knowledge is encoded in a (**V**) matrix that relies on the prior beliefs about the policies (**E**). On the other hand, when the learner observes a novel action, she needs to rely on expected free energy (**G**) computations to perform a step-by-step online encoding and monitoring of the movement. In this case, the weight of the procedural component (**E**) is attenuated, reflecting the lack of prior experience. Hence, while in the former case imitative responses are more driven by procedural knowledge, in the latter case they are more driven in a goal-directed manner. This trade-off between goal-directed and procedural (or habitual) control strategies has been often related to a difference between model-based and model-free methods of reinforcement learning (Daw et al., 2005; Lee et al., 2014; Maisto et al., 2019; Pezzulo et al., 2013b).

Interestingly, the two (procedural and goal-directed) components of action can trigger different actions in the case of non-imitative actions, as in Simulation 4. If the agent were instructed to execute a non-imitative action, the goal-directed (**G**) component would correctly infer the to-be-executed action. However, as discussed above, the agent's procedural component (**E**) is implicitly tuned in a way to execute imitative actions. Hence, it would trigger an imitative response that interferes with the correct action; see the right panel of Figure 6.

5.9 The contributions of mirror mechanisms to action observation, understanding and selection

In our model, mirror mechanisms contribute to action observation, understanding and selection in two main ways. First, the agent's motor knowledge about action sequences encoded in the **e** parameter influences the **d** parameter that encodes prior expectations about the next observed actions, which guides in a cascade action observation processes (Donnarumma et al., 2017; Friston et al., 2011). In this way, the agent implicitly expects to observe the action sequence that she knows the most. Second, as discussed above, posterior beliefs about the (inferred) postures at the Action Recognition level become prior preferences (**C**) at the Action Execution level. This means that when a posture is recognized, the agent automatically prepares an imitative response that might facilitate the execution of congruent actions but impair the execution of incongruent actions, as often observed experimentally (Brass and Heyes, 2005).

Software note

The computational model was implemented using standard routines (**spm_MDP_VB_X.m**) that are freely available as Matlab code in the latest version of SPM academic software: <http://www.fil.ion.ucl.ac.uk/spm/>.

Acknowledgments

This research received funding from the European Union's Horizon 2020 Framework Programme for Research and Innovation under the Specific Grant Agreement No. 945539 (Human Brain Project SGA3) to GP and the European Research Council under the Grant Agreement No. 820213 (ThinkAhead) to GP.

References

- Achilles, E.I.S., Ballweg, C.S., Niessen, E., Kusch, M., Ant, J.M., Fink, G.R., Weiss, P.H., 2019. Neural correlates of differential finger gesture imitation deficits in left hemisphere stroke. *Neuroimage Clin* 23, 101915. <https://doi.org/10.1016/j.nicl.2019.101915>
- Achilles, E.I.S., Fink, G.R., Fischer, M.H., Dovern, A., Held, A., Timpert, D.C., Schroeter, C., Schuetz, K., Kloetzsch, C., Weiss, P.H., 2016. Effect of meaning on apraxic finger imitation deficits. *Neuropsychologia* 82, 74–83. <https://doi.org/10.1016/j.neuropsychologia.2015.12.022>
- Aglioti, S.M., Cesari, P., Romani, M., Urgesi, C., 2008. Action anticipation and motor resonance in elite basketball players. *Nat Neurosci* 11, 1109–1116. <https://doi.org/10.1038/nn.2182>
- Aziz-Zadeh, L., Damasio, A., 2008. Embodied semantics for actions: Findings from functional brain imaging. *Journal of Physiology-Paris, Links and Interactions Between Language and Motor Systems in the Brain* 102, 35–39. <https://doi.org/10.1016/j.jphysparis.2008.03.012>
- Baker, C.L., Saxe, R., Tenenbaum, J.B., 2009. Action understanding as inverse planning. *Cognition, Reinforcement learning and higher cognition* 113, 329–349. <https://doi.org/10.1016/j.cognition.2009.07.005>
- Baker, C.L., Tenenbaum, J.B., Saxe, R.R., 2006. Bayesian models of human action understanding, in: *Advances in Neural Information Processing Systems* 18. MIT Press, pp. 99–106.
- Baranes, A., Oudeyer, P.-Y., Gottlieb, J., 2015. Eye movements reveal epistemic curiosity in human observers. *Vision Research* 117, 81–90. <https://doi.org/10.1016/j.visres.2015.10.009>
- Barsalou, L.W., 2008. Grounded Cognition. *Annual Review of Psychology* 59, 617–645. <https://doi.org/10.1146/annurev.psych.59.103006.093639>
- Barto, A., Mirolli, M., Baldassarre, G., 2013. Novelty or surprise? *Front Psychol* 4, 907. <https://doi.org/10.3389/fpsyg.2013.00907>
- Bartolo, A., Cubelli, R., Della Sala, S., Drei, S., Marchetti, C., 2001. Double dissociation between meaningful and meaningless gesture reproduction in apraxia. *Cortex* 37, 696–699. [https://doi.org/10.1016/s0010-9452\(08\)70617-8](https://doi.org/10.1016/s0010-9452(08)70617-8)
- Beal, M., 2003. Variational algorithms for approximate Bayesian inference /. PhD thesis.
- Berlyne, D.E., 1954. A Theory of Human Curiosity. *British Journal of Psychology. General Section* 45, 180–191. <https://doi.org/10.1111/j.2044-8295.1954.tb01243.x>
- Binkofski, F., Buxbaum, L.J., 2013. Two action systems in the human brain. *Brain Lang* 127, 222–229. <https://doi.org/10.1016/j.bandl.2012.07.007>
- Blake, A., Yuille, A.L. (Eds.), 1992. *Active Vision, Artificial Intelligence Series*. MIT Press, Cambridge, MA, USA.

- Blei, D.M., Ng, A.Y., Jordan, M.I., 2003. Latent dirichlet allocation. *J. Mach. Learn. Res.* 3, 993–1022.
- Bonini, L., 2016. The Extended Mirror Neuron Network Anatomy, Origin, and Functions. *Neuroscientist* 1073858415626400. <https://doi.org/10.1177/1073858415626400>
- Borghi, A.M., Barca, L., Binkofski, F., Castelfranchi, C., Pezzulo, G., Tummolini, L., 2019. Words as social tools: Language, sociality and inner grounding in abstract concepts. *Physics of Life Reviews* 29, 120–153. <https://doi.org/10.1016/j.plrev.2018.12.001>
- Borghi, A.M., Binkofski, F., 2014. The WAT Proposal and the Role of Language, in: Borghi, A.M., Binkofski, F. (Eds.), *Words as Social Tools: An Embodied View on Abstract Concepts*, SpringerBriefs in Psychology. Springer, New York, NY, pp. 19–37. https://doi.org/10.1007/978-1-4614-9539-0_2
- Borghi, A.M., Cimatti, F., 2010. Embodied cognition and beyond: acting and sensing the body. *Neuropsychologia* 48, 763–773. <https://doi.org/10.1016/j.neuropsychologia.2009.10.029>
- Botvinick, M.M., 2008. Hierarchical models of behavior and prefrontal function. *Trends Cogn Sci* 12, 201–208. <https://doi.org/10.1016/j.tics.2008.02.009>
- Botvinick, M.M., Niv, Y., Barto, A.G., 2009. Hierarchically organized behavior and its neural foundations: a reinforcement learning perspective. *Cognition* 113, 262–280. <https://doi.org/10.1016/j.cognition.2008.08.011>
- Brass, M., Bekkering, H., Prinz, W., 2001. Movement observation affects movement execution in a simple response task. *Acta Psychol (Amst)* 106, 3–22. [https://doi.org/10.1016/s0001-6918\(00\)00024-x](https://doi.org/10.1016/s0001-6918(00)00024-x)
- Brass, M., Heyes, C., 2005. Imitation: is cognitive neuroscience solving the correspondence problem? *Trends Cogn Sci* 9, 489–495. <https://doi.org/10.1016/j.tics.2005.08.007>
- Buxbaum, L.J., Randerath, J., 2018. Limb apraxia and the left parietal lobe. *Handb Clin Neurol* 151, 349–363. <https://doi.org/10.1016/B978-0-444-63622-5.00017-6>
- Byrne, R.W., Russon, A.E., 1998. Learning by imitation: A hierarchical approach. *Behav Brain Sci* 21, 667–684. <https://doi.org/10.1017/S0140525X98001745>
- Cai, Y., Urgolites, Z., Wood, J., Chen, C., Li, S., Chen, A., Xue, G., 2018. Distinct neural substrates for visual short-term memory of actions. *Hum Brain Mapp* 39, 4119–4133. <https://doi.org/10.1002/hbm.24236>
- Calvo-Merino, B., Glaser, D.E., Grèzes, J., Passingham, R.E., Haggard, P., 2005. Action observation and acquired motor skills: an FMRI study with expert dancers. *Cereb Cortex* 15, 1243–1249. <https://doi.org/10.1093/cercor/bhi007>
- Calvo-Merino, B., Grèzes, J., Glaser, D.E., Passingham, R.E., Haggard, P., 2006. Seeing or doing? Influence of visual and motor familiarity in action observation. *Curr Biol* 16, 1905–1910. <https://doi.org/10.1016/j.cub.2006.07.065>
- Catmur, C., Walsh, V., Heyes, C., 2009. Associative sequence learning: the role of experience in the

development of imitation and the mirror system. *Philos Trans R Soc Lond B Biol Sci* 364, 2369–2380. <https://doi.org/10.1098/rstb.2009.0048>

Clark, A., 2013. The many faces of precision (Replies to commentaries on “Whatever next? Neural prediction, situated agents, and the future of cognitive science”). *Front. Psychol.* 4. <https://doi.org/10.3389/fpsyg.2013.00270>

Collins, R.P., Litman, J.A., Spielberger, C.D., 2004. The measurement of perceptual curiosity. *Personality and Individual Differences* 36, 1127–1141. [https://doi.org/10.1016/S0191-8869\(03\)00205-8](https://doi.org/10.1016/S0191-8869(03)00205-8)

Cross, E.S., Hamilton, A.F. de C., Grafton, S.T., 2006. Building a motor simulation de novo: observation of dance by dancers. *Neuroimage* 31, 1257–1267. <https://doi.org/10.1016/j.neuroimage.2006.01.033>

Csibra, G., Gergely, G., 2007. ‘Obsessed with goals’: Functions and mechanisms of teleological interpretation of actions in humans. *Acta Psychologica, Becoming an Intentional Agent: Early Development of Action Interpretation and Action Control* 124, 60–78. <https://doi.org/10.1016/j.actpsy.2006.09.007>

Cubelli, R., Marchetti, C., Boscolo, G., Della Sala, S., 2000. Cognition in action: testing a model of limb apraxia. *Brain Cogn* 44, 144–165. <https://doi.org/10.1006/brcg.2000.1226>

Davis, C.P., Yee, E., 2021. Building semantic memory from embodied and distributional language experience. *WIREs Cognitive Science* 12, e1555. <https://doi.org/10.1002/wcs.1555>

Daw, N.D., Niv, Y., Dayan, P., 2005. Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nat Neurosci* 8, 1704–1711. <https://doi.org/10.1038/nn1560>

Dezfouli, A., Balleine, B.W., 2013. Actions, Action Sequences and Habits: Evidence That Goal-Directed and Habitual Action Control Are Hierarchically Organized. *PLOS Computational Biology* 9, e1003364. <https://doi.org/10.1371/journal.pcbi.1003364>

Dezfouli, A., Lingawi, N.W., Balleine, B.W., 2014. Habits as action sequences: hierarchical action control and changes in outcome value. *Philos Trans R Soc Lond B Biol Sci* 369, 20130482. <https://doi.org/10.1098/rstb.2013.0482>

di Pellegrino, G., Fadiga, L., Fogassi, L., Gallese, V., Rizzolatti, G., 1992. Understanding motor events: a neurophysiological study. *Exp Brain Res* 91, 176–180. <https://doi.org/10.1007/BF00230027>

Dindo, H., Zambuto, D., Pezzulo, G., 2011. Motor simulation via coupled internal models using sequential Monte Carlo, in: *Proceedings of IJCAI 2011*. pp. 2113–2119.

Donnarumma, F., Costantini, M., Ambrosini, E., Friston, K., Pezzulo, G., 2017. Action perception as hypothesis testing. *Cortex* 89, 45–60. <https://doi.org/10.1016/j.cortex.2017.01.016>

Engelkamp, J., Zimmer, H.D., 1989. Memory for action events: A new field of research. *Psychological Research* 51, 153–157. <https://doi.org/10.1007/BF00309142>

- Engelkamp, J., Zimmer, H.D., 1985. Motor programs and their relation to semantic memory. *German Journal of Psychology* 9, 239–254.
- Engelkamp, J., Zimmer, H.D., 1984. Motor programme information as a separable memory unit. *Psychological Research* 46, 283–299. <https://doi.org/10.1007/BF00308889>
- Feldman, H., Friston, K., 2010. Attention, Uncertainty, and Free-Energy. *Frontiers in Human Neuroscience* 4, 215. <https://doi.org/10.3389/fnhum.2010.00215>
- Flanagan, J.R., Johansson, R.S., 2003. Action plans used in action observation. *Nature* 424, 769–771. <https://doi.org/10.1038/nature01861>
- Fogassi, L., Ferrari, P.F., Gesierich, B., Rozzi, S., Chersi, F., Rizzolatti, G., 2005. Parietal Lobe: From Action Organization to Intention Understanding. *Science* 308, 662–667. <https://doi.org/10.1126/science.1106138>
- Friston, K., 2010. The free-energy principle: a unified brain theory? *Nat Rev Neurosci* 11, 127–138. <https://doi.org/10.1038/nrn2787>
- Friston, K., 2008. Hierarchical Models in the Brain. *PLOS Computational Biology* 4, e1000211. <https://doi.org/10.1371/journal.pcbi.1000211>
- Friston, K., 2005. A theory of cortical responses. *Philos Trans R Soc Lond B Biol Sci* 360, 815–836. <https://doi.org/10.1098/rstb.2005.1622>
- Friston, K., Adams, R., Perrinet, L., Breakspear, M., 2012. Perceptions as Hypotheses: Saccades as Experiments. *Frontiers in Psychology* 3, 151. <https://doi.org/10.3389/fpsyg.2012.00151>
- Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., O'Doherty, J.P., Pezzulo, G., 2016. Active inference and learning. *Neuroscience & Biobehavioral Reviews* 68, 862–879. <https://doi.org/10.1016/j.neubiorev.2016.06.022>
- Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., Pezzulo, G., 2017. Active Inference: A Process Theory. *Neural Computation* 29, 1–49. https://doi.org/10.1162/NECO_a_00912
- Friston, K., Kiebel, S., 2009. Predictive coding under the free-energy principle. *Philos Trans R Soc Lond B Biol Sci* 364, 1211–1221. <https://doi.org/10.1098/rstb.2008.0300>
- Friston, K., Mattout, J., Kilner, J., 2011. Action understanding and active inference. *Biol Cybern* 104, 137–160. <https://doi.org/10.1007/s00422-011-0424-z>
- Friston, K., Mattout, J., Trujillo-Barreto, N., Ashburner, J., Penny, W., 2007. Variational free energy and the Laplace approximation. *Neuroimage* 34, 220–234. <https://doi.org/10.1016/j.neuroimage.2006.08.035>
- Friston, K., Rigoli, F., Ognibene, D., Mathys, C., Fitzgerald, T., Pezzulo, G., 2015. Active inference and epistemic value. *Cogn Neurosci* 6, 187–214. <https://doi.org/10.1080/17588928.2015.1020053>

- Friston, K., Rosch, R., Parr, T., Price, C., Bowman, H., 2018. Deep temporal models and active inference. *Neuroscience & Biobehavioral Reviews* 90, 486–501. <https://doi.org/10.1016/j.neubiorev.2018.04.004>
- Friston, K.J., Stephan, K.E., Montague, R., Dolan, R.J., 2014. Computational psychiatry: the brain as a phantastic organ. *The Lancet Psychiatry* 1, 148–158. [https://doi.org/10.1016/S2215-0366\(14\)70275-5](https://doi.org/10.1016/S2215-0366(14)70275-5)
- Galetzka, C., 2017. The Story So Far: How Embodied Cognition Advances Our Understanding of Meaning-Making. *Front. Psychol.* 8. <https://doi.org/10.3389/fpsyg.2017.01315>
- Gallese, V., Goldman, A., 1998. Mirror neurons and the simulation theory of mind-reading. *Trends in Cognitive Sciences* 2, 493–501. [https://doi.org/10.1016/S1364-6613\(98\)01262-5](https://doi.org/10.1016/S1364-6613(98)01262-5)
- Gazzola, V., Keysers, C., 2009. The Observation and Execution of Actions Share Motor and Somatosensory Voxels in all Tested Subjects: Single-Subject Analyses of Unsmoothed fMRI Data. *Cerebral Cortex* 19, 1239–1255. <https://doi.org/10.1093/cercor/bhn181>
- Gergely, G., Csibra, G., 2003. Teleological reasoning in infancy: The naïve theory of rational action. *Trends in Cognitive Sciences* 7, 287–292. [https://doi.org/10.1016/S1364-6613\(03\)00128-1](https://doi.org/10.1016/S1364-6613(03)00128-1)
- Goldenberg, G., Hagmann, S., 1998. Tool use and mechanical problem solving in apraxia. *Neuropsychologia* 36, 581–589. [https://doi.org/10.1016/S0028-3932\(97\)00165-6](https://doi.org/10.1016/S0028-3932(97)00165-6)
- Goldenberg, G., Hagmann, S., 1997. The meaning of meaningless gestures: A study of visuo-imitative apraxia. *Neuropsychologia* 35, 333–341. [https://doi.org/10.1016/S0028-3932\(96\)00085-1](https://doi.org/10.1016/S0028-3932(96)00085-1)
- Gottlieb, J., Oudeyer, P.-Y., 2018. Towards a neuroscience of active sampling and curiosity. *Nat Rev Neurosci* 19, 758–770. <https://doi.org/10.1038/s41583-018-0078-0>
- Grafton, S.T., 2009. Embodied cognition and the simulation of action to understand others. *Ann N Y Acad Sci* 1156, 97–117. <https://doi.org/10.1111/j.1749-6632.2009.04425.x>
- Grafton, S.T., de C. Hamilton, A.F., 2007. Evidence for a distributed hierarchy of action representation in the brain. *Human Movement Science, European Workshop on Movement Science 2007* 26, 590–616. <https://doi.org/10.1016/j.humov.2007.05.009>
- Hauk, O., Johnsrude, I., Pulvermüller, F., 2004. Somatotopic Representation of Action Words in Human Motor and Premotor Cortex. *Neuron* 41, 301–307. [https://doi.org/10.1016/S0896-6273\(03\)00838-9](https://doi.org/10.1016/S0896-6273(03)00838-9)
- Heilbron, M., Armeni, K., Schoffelen, J.-M., Hagoort, P., de Lange, F.P., 2022. A hierarchy of linguistic predictions during natural language comprehension. *Proceedings of the National Academy of Sciences* 119, e2201968119. <https://doi.org/10.1073/pnas.2201968119>
- Heilman, K.M., Maher, L.M., Greenwald, M.L., Rothi, L.J.G., 1997. Conceptual apraxia from lateralized lesions. *Neurology* 49, 457–464. <https://doi.org/10.1212/WNL.49.2.457>
- Heyes, C., 2011. Automatic imitation. *Psychol Bull* 137, 463–483. <https://doi.org/10.1037/a0022288>

- Hoeren, M., Kümmerer, D., Bormann, T., Beume, L., Ludwig, V.M., Vry, M.-S., Mader, I., Rijntjes, M., Kaller, C.P., Weiller, C., 2014. Neural bases of imitation and pantomime in acute stroke patients: distinct streams for praxis. *Brain* 137, 2796–2810. <https://doi.org/10.1093/brain/awu203>
- Hommel, B., Müsseler, J., Aschersleben, G., Prinz, W., 2001. The Theory of Event Coding (TEC): A framework for perception and action planning. *Behavioral and Brain Sciences* 24, 849–878. <https://doi.org/10.1017/S0140525X01000103>
- Hurley, S., Chater, N. (Eds.), 2005. *Perspectives on Imitation: From Neuroscience to Social Science - Two Volume Set*, Social Neuroscience. A Bradford Book, Cambridge, MA, USA.
- Iacoboni, M., Woods, R.P., Brass, M., Bekkering, H., Mazziotta, J.C., Rizzolatti, G., 1999. Cortical mechanisms of human imitation. *Science* 286, 2526–2528. <https://doi.org/10.1126/science.286.5449.2526>
- Itti, L., Baldi, P., 2009. Bayesian surprise attracts human attention. *Vision Research* 49, 1295–1306. <https://doi.org/10.1016/j.visres.2008.09.007>
- Itti, L., Koch, C., 2000. A saliency-based search mechanism for overt and covert shifts of visual attention. *Vision Research* 40, 1489–1506. [https://doi.org/10.1016/S0042-6989\(99\)00163-7](https://doi.org/10.1016/S0042-6989(99)00163-7)
- Itti, L., Koch, C., Niebur, E., 1998. A model of saliency-based visual attention for rapid scene analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 20, 1254–1259. <https://doi.org/10.1109/34.730558>
- Jacob, P., Jeannerod, M., 2005. The motor theory of social cognition: a critique. *Trends in Cognitive Sciences* 9, 21–25. <https://doi.org/10.1016/j.tics.2004.11.003>
- Jacquet, P.O., Chambon, V., Borghi, A.M., Tessari, A., 2012. Object affordances tune observers' prior expectations about tool-use behaviors. *PLoS ONE* 7. <https://doi.org/10.1371/journal.pone.0039629>
- James, W., 1890. *The principles of psychology*, Vol I., *The principles of psychology*, Vol I. Henry Holt and Co, New York, NY, US. <https://doi.org/10.1037/10538-000>
- Jax, S.A., Buxbaum, L.J., 2010. Response interference between functional and structural actions linked to the same familiar object. *Cognition* 115, 350–355. <https://doi.org/10.1016/j.cognition.2010.01.004>
- Jeannerod, M., 2006. *Motor cognition: What actions tell the self*, *Motor cognition: What actions tell the self*. Oxford University Press, New York, NY, US. <https://doi.org/10.1093/acprof:oso/9780198569657.001.0001>
- Jeannerod, M., 2001. Neural simulation of action: a unifying mechanism for motor cognition. *Neuroimage* 14, S103-109. <https://doi.org/10.1006/nimg.2001.0832>
- Jeannerod, M., 1994. The representing brain: Neural correlates of motor intention and imagery. *Behavioral and Brain Sciences* 17, 187–245. <https://doi.org/10.1017/S0140525X00034026>
- Kiebel, S., Daunizeau, J., Friston, K., 2009. Perception and hierarchical dynamics. *Frontiers in*

- Neuroinformatics 3, 20. <https://doi.org/10.3389/neuro.11.020.2009>
- Kiebel, S., Daunizeau, J., Friston, K., 2008. A Hierarchy of Time-Scales and the Brain. *PLOS Computational Biology* 4, e1000209. <https://doi.org/10.1371/journal.pcbi.1000209>
- Kilner, J.M., 2011. More than one pathway to action understanding. *Trends in Cognitive Sciences* 15, 352–357. <https://doi.org/10.1016/j.tics.2011.06.005>
- Kilner, J.M., Friston, K., Frith, C.D., 2007. Predictive coding: an account of the mirror neuron system. *Cogn Process* 8, 159–166. <https://doi.org/10.1007/s10339-007-0170-2>
- Kilner, J.M., Paulignan, Y., Blakemore, S.J., 2003. An interference effect of observed biological movement on action. *Curr Biol* 13, 522–525. [https://doi.org/10.1016/s0960-9822\(03\)00165-9](https://doi.org/10.1016/s0960-9822(03)00165-9)
- Krakauer, J.W., Shadmehr, R., 2007. Towards a computational neuropsychology of action, in: Cisek, P., Drew, T., Kalaska, J.F. (Eds.), *Progress in Brain Research, Computational Neuroscience: Theoretical Insights into Brain Function*. Elsevier, pp. 383–394. [https://doi.org/10.1016/S0079-6123\(06\)65024-3](https://doi.org/10.1016/S0079-6123(06)65024-3)
- Lee, S.W., Shimojo, S., O’Doherty, J.P., 2014. Neural computations underlying arbitration between model-based and model-free learning. *Neuron* 81, 687–699. <https://doi.org/10.1016/j.neuron.2013.11.028>
- Lee, T.S., Mumford, D., 2003. Hierarchical Bayesian inference in the visual cortex. *J Opt Soc Am A Opt Image Sci Vis* 20, 1434–1448. <https://doi.org/10.1364/josaa.20.001434>
- Logie, R.H., 1986. Visuo-spatial processing in working memory. *The Quarterly Journal of Experimental Psychology A: Human Experimental Psychology* 38A, 229–247. <https://doi.org/10.1080/14640748608401596>
- Maisto, D., Friston, K., Pezzulo, G., 2019. Caching mechanisms for habit formation in Active Inference. *Neurocomputing* 359, 298–314. <https://doi.org/10.1016/j.neucom.2019.05.083>
- Maranesi, M., Bruni, S., Livi, A., Donnarumma, F., Pezzulo, G., Bonini, L., 2019. Differential neural dynamics underling pragmatic and semantic affordance processing in macaque ventral premotor cortex. *Sci Rep* 9, 11700. <https://doi.org/10.1038/s41598-019-48216-y>
- Maranesi, M., Livi, A., Fogassi, L., Rizzolatti, G., Bonini, L., 2014. Mirror Neuron Activation Prior to Action Observation in a Predictable Context. *J. Neurosci.* 34, 14827–14832. <https://doi.org/10.1523/JNEUROSCI.2705-14.2014>
- Martin, M., Dressing, A., Bormann, T., Schmidt, C., Kümmerer, D., Beume, L., Saur, D., Mader, I., Rijntjes, M., Kaller, C., Weiller, C., 2016a. Componential Network for the Recognition of Tool-Associated Actions: Evidence from Voxel-based Lesion-Symptom Mapping in Acute Stroke Patients. *Cerebral Cortex* 27. <https://doi.org/10.1093/cercor/bhw226>
- Martin, M., Nitschke, K., Beume, L., Dressing, A., Bühler, L.E., Ludwig, V.M., Mader, I., Rijntjes, M.,

- Kaller, C.P., Weiller, C., 2016b. Brain activity underlying tool-related and imitative skills after major left hemisphere stroke. *Brain* 139, 1497–1516. <https://doi.org/10.1093/brain/aww035>
- Mengotti, P., Corradi-Dell’Acqua, C., Negri, G.A.L., Ukmar, M., Pesavento, V., Rumiati, R.I., 2013. Selective imitation impairments differentially interact with language processing. *Brain* 136, 2602–2618. <https://doi.org/10.1093/brain/awt194>
- Mengotti, P., Ripamonti, E., Pesavento, V., Rumiati, R.I., 2015. Anatomical and spatial matching in imitation: Evidence from left and right brain-damaged patients. *Neuropsychologia*, Special Issue: Sensory Motor Integration 79, 256–271. <https://doi.org/10.1016/j.neuropsychologia.2015.06.038>
- Meteyard, L., Cuadrado, S.R., Bahrami, B., Vigliocco, G., 2012. Coming of age: A review of embodiment and the neuroscience of semantics. *Cortex, Language and the Motor System* 48, 788–804. <https://doi.org/10.1016/j.cortex.2010.11.002>
- Miller, G.A., Galanter, E., Pribram, K.H., 1960. The integration of plans, in: *Plans and the Structure of Behavior*. Henry Holt and Co, New York, NY, US, pp. 95–102. <https://doi.org/10.1037/10039-007>
- Mirza, M.B., Adams, R.A., Mathys, C.D., Friston, K., 2016. Scene Construction, Visual Foraging, and Active Inference. *Frontiers in Computational Neuroscience* 10.
- O’Regan, J.K., 2011. *Why Red Doesn’t Sound Like a Bell: Understanding the feel of consciousness*. Oxford University Press.
- Pacherie, E., 2008. The phenomenology of action: A conceptual framework. *Cognition* 107, 179. <https://doi.org/10.1016/j.cognition.2007.09.003>
- Parr, T., Friston, K.J., 2018. The Discrete and Continuous Brain: From Decisions to Movement—And Back Again. *Neural Comput* 30, 2319–2347. https://doi.org/10.1162/neco_a_01102
- Parr, T., Friston, K.J., 2017. Working memory, attention, and salience in active inference. *Sci Rep* 7, 14678. <https://doi.org/10.1038/s41598-017-15249-0>
- Parr, T., Limanowski, J., Rawji, V., Friston, K., 2021. The computational neurology of movement under active inference. *Brain* 144, 1799–1818. <https://doi.org/10.1093/brain/awab085>
- Parr, T., Pezzulo, G., Friston, K., 2022. *Active Inference: The Free Energy Principle in Mind, Brain, and Behavior*. MIT Press, Cambridge, MA, USA.
- Pedullà, L., Gervasoni, E., Bisio, A., Biggio, M., Ruggeri, P., Avanzino, L., Bove, M., 2020. The last chance to pass the ball: investigating the role of temporal expectation and motor resonance in processing temporal errors in motor actions. *Social Cognitive and Affective Neuroscience* 15, 123–134. <https://doi.org/10.1093/scan/nsaa021>
- Peigneux, P., Van der Linden, M., Garraux, G., Laureys, S., Degueldre, C., Aerts, J., Fiore, G., Moonen, G.,

- Luxen, A., Salmon, E., 2004. Imaging A Cognitive Model of Apraxia: The Substrate of Gesture-Specific Cognitive Processes. *Human brain mapping* 21, 119–42. <https://doi.org/10.1002/hbm.10161>
- Pezzulo, G., 2013. Studying mirror mechanisms within generative and predictive architectures for joint action. *Cortex* 49, 2968–2969. <https://doi.org/10.1016/j.cortex.2013.06.008>
- Pezzulo, G., Barsalou, L.W., Cangelosi, A., Fischer, M.H., McRae, K., Spivey, M.J., 2013a. Computational Grounded Cognition: a new alliance between grounded cognition and computational modeling. *Front Psychol* 3, 612. <https://doi.org/10.3389/fpsyg.2012.00612>
- Pezzulo, G., Barsalou, L.W., Cangelosi, A., Fischer, M.H., Spivey, M., McRae, K., 2011. The Mechanics of Embodiment: A Dialog on Embodiment and Computational Modeling. *Front. Psychol.* 2. <https://doi.org/10.3389/fpsyg.2011.00005>
- Pezzulo, G., Rigoli, F., Chersi, F., 2013b. The Mixed Instrumental Controller: Using Value of Information to Combine Habitual Choice and Mental Simulation. *Frontiers in Psychology* 4, 92. <https://doi.org/10.3389/fpsyg.2013.00092>
- Pezzulo, G., Rigoli, F., Friston, K., 2018. Hierarchical Active Inference: A Theory of Motivated Control. *Trends in Cognitive Sciences* 22, 294–306. <https://doi.org/10.1016/j.tics.2018.01.009>
- Pezzulo, G., Rigoli, F., Friston, K., 2015. Active Inference, homeostatic regulation and adaptive behavioural control. *Progress in Neurobiology* 134, 17–35. <https://doi.org/10.1016/j.pneurobio.2015.09.001>
- Pezzulo, G., Zorzi, M., Corbetta, M., 2021. The secret life of predictive brains: what’s spontaneous activity for? *Trends Cogn Sci* 25, 730–743. <https://doi.org/10.1016/j.tics.2021.05.007>
- Posner, M.I., Petersen, S.E., 1990. The Attention System of the Human Brain. *Annual Review of Neuroscience* 13, 25–42. <https://doi.org/10.1146/annurev.ne.13.030190.000325>
- Rijntjes, M., Weiller, C., Bormann, T., Musso, M., 2012. The dual loop model: its relation to language and other modalities. *Frontiers in Evolutionary Neuroscience* 4, 9. <https://doi.org/10.3389/fnevo.2012.00009>
- Rizzolatti, G., Craighero, L., 2004. The mirror-neuron system. *Annu Rev Neurosci* 27, 169–192. <https://doi.org/10.1146/annurev.neuro.27.070203.144230>
- Rizzolatti, G., Fogassi, L., Gallese, V., 2001. Neurophysiological mechanisms underlying the understanding and imitation of action. *Nat Rev Neurosci* 2, 661–670. <https://doi.org/10.1038/35090060>
- Rizzolatti, G., Sinigaglia, C., 2010. The functional role of the parieto-frontal mirror circuit: interpretations and misinterpretations. *Nat Rev Neurosci* 11, 264–274. <https://doi.org/10.1038/nrn2805>
- Rothi, L.J., Ochipa, C., Heilman, K.M., 1991. A cognitive neuropsychological model of limb praxis. *Cognitive Neuropsychology* 8, 443–458. <https://doi.org/10.1080/02643299108253382>
- Roy, E.A., Square, P.A., 1985. Common Considerations In The Study of Limb, Verbal And Oral Apraxia, in:

- Roy, E.A. (Ed.), *Advances in Psychology*. North-Holland, pp. 111–161. [https://doi.org/10.1016/S0166-4115\(08\)61139-5](https://doi.org/10.1016/S0166-4115(08)61139-5)
- Rumiati, R.I., Papeo, L., Corradi-Dell’Acqua, C., 2010. Higher-level motor processes. *Annals of the New York Academy of Sciences* 1191, 219–241. <https://doi.org/10.1111/j.1749-6632.2010.05442.x>
- Rumiati, R.I., Tessari, A., 2002. Imitation of novel and well-known actions: the role of short-term memory. *Exp Brain Res* 142, 425–433. <https://doi.org/10.1007/s00221-001-0956-x>
- Rumiati, R.I., Weiss, P.H., Tessari, A., Assmus, A., Zilles, K., Herzog, H., Fink, G.R., 2005. Common and differential neural mechanisms supporting imitation of meaningful and meaningless actions. *J Cogn Neurosci* 17, 1420–1431. <https://doi.org/10.1162/0898929054985374>
- Schwartenbeck, P., Friston, K., 2016. Computational Phenotyping in Psychiatry: A Worked Example. *eNeuro* 3. <https://doi.org/10.1523/ENEURO.0049-16.2016>
- Schwartenbeck, P., Passecker, J., Hauser, T.U., FitzGerald, T.H., Kronbichler, M., Friston, K., 2019. Computational mechanisms of curiosity and goal-directed exploration. *eLife* 8, e41703. <https://doi.org/10.7554/eLife.41703>
- Smith, R., Schwartenbeck, P., Parr, T., Friston, K., 2020. An Active Inference Approach to Modeling Structure Learning: Concept Learning as an Example Case. *Frontiers in Computational Neuroscience* 14, 41. <https://doi.org/10.3389/fncom.2020.00041>
- Sokolov, E.N., 1963. Higher Nervous Functions: The Orienting Reflex. *Annual Review of Physiology* 25, 545–580. <https://doi.org/10.1146/annurev.ph.25.030163.002553>
- Tessari, A., Bosanac, D., Rumiati, R.I., 2006. Effect of learning on imitation of new actions: implications for a memory model. *Exp Brain Res* 173, 507–513. <https://doi.org/10.1007/s00221-006-0395-9>
- Tessari, A., Canessa, N., Ukmar, M., Rumiati, R.I., 2007. Neuropsychological evidence for a strategic control of multiple routes in imitation. *Brain* 130, 1111–1126. <https://doi.org/10.1093/brain/awm003>
- Tessari, A., Mengotti, P., Faccioli, L., Tuoizzi, G., Boscarato, S., Taricco, M., Rumiati, R.I., 2021. Effect of body-part specificity and meaning in gesture imitation in left hemisphere stroke patients. *Neuropsychologia* 151, 107720. <https://doi.org/10.1016/j.neuropsychologia.2020.107720>
- Tessari, A., Proietti, R., Rumiati, R.I., 2022. Bottom-up and top-down modulation of route selection in imitation. *null* 1–16. <https://doi.org/10.1080/02643294.2022.2043264>
- Tessari, A., Rumiati, R.I., 2004. The Strategic Control of Multiple Routes in Imitation of Actions. *Journal of Experimental Psychology: Human Perception and Performance* 30, 1107–1116. <https://doi.org/10.1037/0096-1523.30.6.1107>
- Tessari, A., Rumiati, R.I., Haggard, P., 2002. Imitation without awareness. *Neuroreport* 13, 2531–2535.

<https://doi.org/10.1097/00001756-200212200-00030>

Tsotsos, J.K., 2011. *A Computational Perspective on Visual Attention*. MIT Press.

van Elk, M., van Schie, H., Bekkering, H., 2014. Action semantics: A unifying conceptual framework for the selective use of multimodal and modality-specific object knowledge. *Physics of Life Reviews* 11, 220–250. <https://doi.org/10.1016/j.plrev.2013.11.005>

Von Hofsten, C., 2007. Action in development. *Developmental Science* 10, 54–60. <https://doi.org/10.1111/j.1467-7687.2007.00564.x>

Weiller, C., Bormann, T., Saur, D., Musso, M., Rijntjes, M., 2011. How the ventral pathway got lost – And what its recovery might mean. *Brain and Language* 118, 29–39. <https://doi.org/10.1016/j.bandl.2011.01.005>

Weiller, C., Musso, M., Rijntjes, M., Saur, D., 2009. Please don't underestimate the ventral pathway in language. *Trends Cogn Sci* 13, 369–370; 370–371. <https://doi.org/10.1016/j.tics.2009.06.007>