Alma Mater Studiorum - Università di Bologna

DOTTORATO DI RICERCA IN

PSYCHOLOGY

Ciclo 35

Settore Concorsuale: 11/E1 - PSICOLOGIA GENERALE, PSICOBIOLOGIA E PSICOMETRIA

Settore Scientifico Disciplinare: M-PSI/01 - PSICOLOGIA GENERALE

HIERARCHICAL ACTIVE INFERENCE FOR COGNITIVE ARCHITECTURES

Presentata da: Riccardo Proietti

Coordinatore Dottorato	**Supervisore**
Elisabetta Crocetti	Alessia Tessari

Esame finale anno 2023

Hierarchical active inference for cognitive architectures

Abstract

This thesis explores the methods based on the free energy principle and active inference for modelling cognition. Active inference is an emerging framework for designing intelligent agents where psychological processes are cast as a form of approximate Bayesian inference. Here, I appeal to it to test the design of a set of cognitive architectures, via simulation. These architectures are defined in terms of probabilistic hierarchical generative models where an active inference agent executes a task under the assumption that all cognitive processes aspire to the same objective: the minimization of an approximation of sensory surprisal, named variational free energy. Chapter 1 introduces the free energy principle and its assumptions about self-organizing systems. Chapter 2 describes how from the mechanics of self-organization can emerge a minimal form of cognition able to achieve autopoiesis. These mechanics can be formalized in terms of active Bayesian inference where a cognitive agent behaves adaptively through action and perception under a probabilistic generative model. In chapter 3 I present the method of how I formalize these models in a hierarchical setting and discrete state spaces. The architectures proposed allow providing a more biologically plausible account of more complex cognitive processing that entails deep temporal features. I then present three simulation studies using these kinds of architectures that aim to show different aspects of cognition, their associated behavior and the underlying neural dynamics. In chapter 4, the first study proposes an architecture that represents the visuomotor system for the encoding of actions during action observation, understanding and imitation. In chapter 5, the generative model is extended and is lesioned to simulate brain damage and neuropsychological patterns observed in apraxic patients. In chapter 6, the third study proposes an architecture for cognitive control and the modulation of attention for action selection. At last, I argue how active inference can provide a formal account of information processing in the brain and how the adaptive capabilities of the simulated agents are a mere consequence of the architecture of the generative models. Cognitive processing, then, becomes an emergent property of the minimization of variational free energy.

Table of Contents

Abstract 1

Introduction 4

Chapter 1: The architecture of life: self-organization under the free energy principle 6

  1.1 The Markov blanket 6

  1.2 From non-equilibrium steady state dynamics to Bayesian mechanics 8

Chapter 2: Self-evidencing architectures: generative models for perception and action.. 11

  2.1 Generative models 11

  2.2 Perception as variational inference 12

  2.3 Active inference 13

Chapter 3: Hierarchical architectures: active inference in deep temporal models.. 15

  3.1 Active inference in discrete state spaces: partially observable Markov decision processes.................. 15

  3.2 Hierarchical active inference: deep temporal models.. 18

Chapter 4: 1st Architecture: an active inference model of hierarchical action understanding, learning and imitation.. 20

  4.1 Introduction: action understanding as a hierarchical inference problem.. 20

  4.2 Active inference simulations during action understanding and imitation 23

  4.3 Neurobiological underpinnings of the proposed model.. 36

  4.4 Methods 40

  4.5 General discussion 46

Chapter 5: 2nd Architecture: limb apraxia and active inference in the visuomotor pathways 48

  5.1 Introduction 48

  5.2 Limb apraxia 49

5.3 A cognitive model of praxis 49

5.4 Neuroanatomical mapping 51

5.5 The simulated tasks and the generative model 52

5.7 Disconnections and Bayes optimal pathology 57

5.8 In silico brain damages of the praxis model 58

5.9 Discussion 66

Chapter 6: 3rd Architecture: Cognitive control in active inference 67

6.1 Introduction 68

6.2 Computational model of cognitive control during driving 70

6.3 Neuronal dynamics of cognitive control 79

6.4 Simulations of cognitive control in the driving task 80

6.5 General discussion 90

Chapter 7: Discussion 94

Acknowledgements 96

References 97

Introduction

In the second half of the previous century, the leading approach in theoretical and experimental psychology shifted from behaviorism to information processing (Barsalou, 2014). This shift was driven by developments in other disciplines, such as Cybernetics (Ashby, 1957), the science where information processing and adaptive control are interpreted as adjustments to perturbations, information theory (Shannon, 1948) which provided a mathematical formalism to quantify information, and the science of computation, since von Neumann's first design of intelligent machines (Von Neumann & Burks, 1966). The idea that humans can also be described as information-processing machines emerged from all these new insights. This made the applications of machine learning and cognitive science close to the other (Ackley et al., 1985; Newell & Simon, 1972). Furthermore, in the last decades, embodied cognition in neuroscience, enactivism in philosophy, and applications in robotics pushed for a more grounded approach to cognition, by showing its roots in motor behavior (Barsalou, 2008; Lungarella et al., 2003; O'Regan, 2011). These approaches emphasized the concept that cognition typically involves acting with a physical body immersed in an environment, leading toward a pragmatic turn in cognitive science (Engel et al., 2013, 2016).

This thesis aims to offer a unified perspective under a unique framework, *the free energy principle* and its main corollary, *active inference*. The free energy principle starts from dynamical systems theory and Bayesian statistics and applies them to the study of self-organization (Haken, 2004; Kauffman & Kauffman, 1993). In this way, many aspects of the brain, such as how we perceive and behave, become a necessary, self-evident and emergent feature of the physics of life (Friston, 2019; Kirchhoff et al., 2018). Active inference, on the other hand, represents the corollary of how the free energy principle is realized. Under this setting, biological agents are cast as probabilistic inference machines that embody a generative model that continuously generates predictions to maintain their existence by following a sole imperative: the minimization of a tractable proxy of self-information, called variational free energy, via engaging in perception and action (Da Costa et al., 2020; Friston et al., 2017; Parr et al., 2022).

This unique mathematical solution accounts for different psychological functions, such as perception, learning, and decision-making, by treating them as interdependent inference processes. Under active inference, cognition and behaviour are formalized as belief-based schemes with several implications on physiology, neurobiology and neural correlates (Parr et al., 2022). The assumption is that specific cognitive processes and the associated neurobiological substrata can be derived from the principle of minimizing variational free energy under a generative model. Thus, in the studies presented, this idea will be used to show how organisms solve their adaptive problems, via agent-based simulations. More practically, these simulations rely on the design of a specific generative model in discrete state spaces that recapitulates aspects of cognition and brain dynamics and how they are translated into behavior, realized by a probabilistic specification of the trajectories of desired future observations. The usefulness of this stands in the fact that active inference furnishes the rules and the constraints required to simulate, predict and understand behaviour: as mentioned by Richard Feynman

“What I cannot create, I do not understand”. Indeed, creating simulations represents one of the novel directions taken by cognitive science to depart from more narrative formulations typical of the previous century (Parr et al., 2022). The falsifiability and the subsequent experimental investigation are then associated with testing a specific generative model and the associated process theories following the so-called computational phenotyping trend (Schwartenbeck & Friston, 2016; Smith et al., 2021).

The free energy principle stems from straightforward results from statistical mechanics. In a way, it provides a Bayesian interpretation of the principle of least action (Friston, 2010, 2019; Ramstead et al., 2022). In the next three chapters I provide a theoretical background by sketching up the fundamental assumptions of the principle: I start from the definition of any ‘thing’ provided by the Markov blanket and how this connects to a physical, dynamical system interpretation of a system that organizes itself. Then I show how a new kind of physics can provide a probabilistic description of the system, taking into account the concept of surprise and model evidence. This is based on a key formalism: generative models. These formal tools ‘look as if’ they are embodied by every self-organizing agent: the notion of action and perception emerge automatically when the creature minimizes free energy (Friston, 2019). Crucially, when we speak about cognitive architectures and more intelligent designs, the process becomes prospective as the behaviour of an allostatic agent is directed towards the future (Corcoran et al., 2020). I show how this is put into practice by providing simulations in discrete state spaces under partially observable Markov decision processes. The subsequent three chapters implement these formalizations in three simulations studies, presented in article format, that aim to provide a biologically plausible account of relatively complex behaviors and cognitive processing: they include visual sampling, hierarchical encoding of action representations, motor execution and mechanisms of the reciprocal influence between enacted and observed movements, the underlying neuropsychological processes and the functional impairments in case of brain lesions, habitization, the engaging of mental effort and attentional control, and more. In all these characterizations the aim is to provide links between the computational formalism, the cognitive representation, and the underlying neurobiology. The aim is to show that cognitive processing can be regarded as an emergent property of the process of minimizing variational free energy at different temporal scales. At the same time, I highlight how disfunctions emerge from optimal inferences but with suboptimal parameters. I also hint to the idea of how neural dynamics and brain organization must reflect the cognitive processing and the shapes of the generative models. The idea is to test if and how active inference can offer a way of understanding cognition and behavior starting from the dynamics of self-organization and Bayesian inference.

Chapter 1

The architecture of life: self-organization under the free energy principle

1.1 The Markov blanket

Consider the following problem. How does a biological system resist the natural tendency to disorder and dissipation in an unpredictable and constantly changing world? The fluctuation theorem generalizes the second law of thermodynamics by demanding the long-term increase of a system's entropy. In other terms, it mandates thermodynamic equilibrium or, in a word, death (Evans & Searles, 2002). Yet, living organisms seem to violate this law as they are capable of existing for long periods of time, maintaining a precise set of physiological states (Friston & Ao, 2012). The free energy principle (FEP) aims to account for this observation by answering these questions: if something does exist, what kind of processes, dynamics or features must it possess? What must a living being do to exist?

The FEP starts by providing a particular statistical definition of living things and then displays the behavior and the dynamics that satisfy that definition (Friston, 2019). The first step is to separate things that exist from those that do not. To do this is necessary to rely on the notion of a spatial boundary (Pearl, 1988). If cast in a statistical term, the probabilistic boundary is called *Markov blanket*. Consider a state space with all the states of the universe; here, we take some of those states and call them *internal states*. Their *blanket states* would be the ones that separate the internal states from the influence of *external states*. Therefore, the external states are conditionally independent from the internal states. More specifically, the blanket states are constituted by the parents of an internal state, its children, and the parents of the children. See figure 1.1 for an exemplification. The fundamental feature of the Markov blanket is the Markov property: considering the causal relationship of these three sets of states, in order to have a statistical interpretation of the dynamics of the internal states given the external states, the only thing required to know are the blanket states. In a real-world setting, the blanket states can be seen as the statistical partitions that separate the living being from the outside world, the external states. For example, to understand what is going on in the environment (the external states), the brain (the internal state) builds a statistical interpretation based on the blanket states, meaning the stimulations of my sensory epithelia, such as the physiological state of my photoreceptors when I am observing the world (Friston, 2019).

Furthermore, as shown in figure 1.1., the blanket states can also be divided into: 1) active states that can change external states and are influenced only by internal states; 2) sensory states which influence the internal states and are influenced by external states. From this characterization, it emerges how, in a dynamical setting, the organism and the environment are coupled: active states, constituted by effectors, actuators and reflexes, act upon the world by influencing the external states which, in turn, they generate an impression on the sensory states, such as sensory organs and receptors, that updates internal states which determine further active states in a perpetual action-perception cycle (Friston, 2019; Parr et al., 2022). In real-world settings, the Markov blanket characterization usually entails a multiscale perspective, which can be viewed as a system of *nested Markov blankets*. For example, the membrane of a cell is insulated but put in interaction with a network of other cells which constitutes a creature that, in interactions with other creatures, becomes a social network or a society (Hesp et al., 2019; Palacios et al., 2020).

Importantly, the statistical partition provided by the Markov blanket suggests a symmetry in this interaction, as the influence between external and internal states because one becomes informative of the other through the blanket. This leads to a form of synchronization or *dynamical coupling*, an emergent property of two coupled systems predicting each other (Friston & Frith, 2015). This phenomenon is easily exemplified in the 1600s experiment from the Dutch scientist Christiaan Huygens where two oscillators, such as two pendula or metronomes, converge to a synchronous ticking when they share a common mechanical support. Through this support, indeed, the metronomes can communicate by 'acting' through mechanical forces. This synchronization is one of the most pervasive emergent properties in nature, and complex systems in particular (Friston et al., 2021; Strogatz, 2012; Strogatz & Stewart, 1993), as it explains the dynamics of brain and heart cells, swarms, fireflies, circadian rhythms, and several other systems (Matheny et al., 2019; Varela et al., 2001). See figure 1.1 for an exemplification.

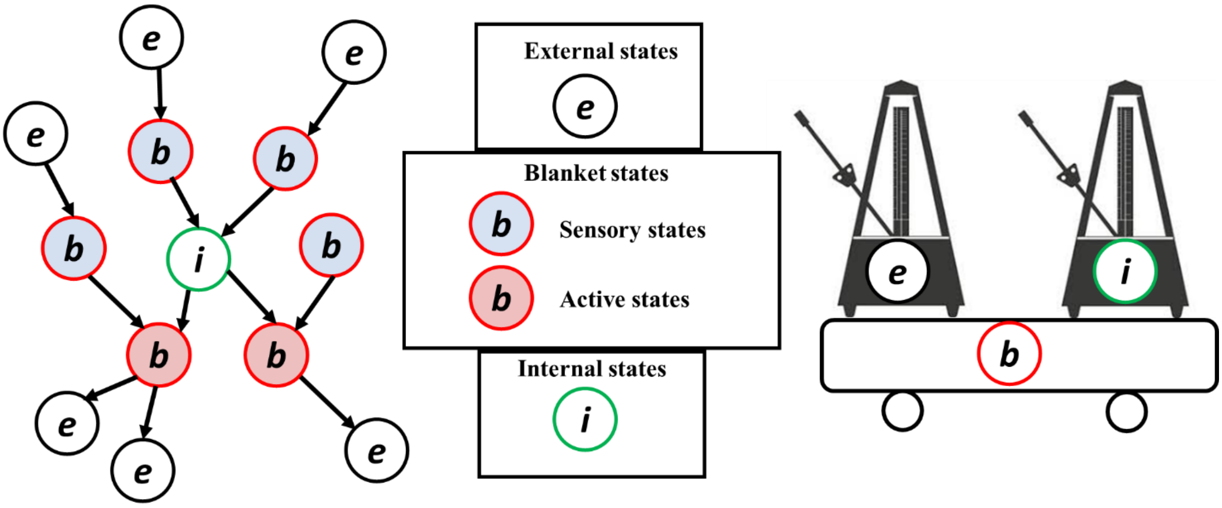

Figure 1.1. On the left, a schematic in the form of a Bayesian network depicts the states of the universe and, given an internal state, we can individuate its Markov blanket by taking the parents of the internal state, its children, and the parents of the children. Each arrow indicates conditional dependence; thus, the internal state is insulated from the influence of the external states. These external states influence the sensory states of the internal state, which can influence the blanket's active states. Only through action via active states can the system influence the external states. On the right panel, the dynamical coupling between internal and external states through the blanket. The reciprocal influence of these two leads the two oscillators to synchronize.

As I will show, by looking at this synchrony between internal and external states, it looks as if they are representing and acting upon each other. The internal states are modelling and influencing the environment but at the same time, the environment is modelling and influencing the internal states. I will show how these dynamics can be cast in terms of Bayesian mechanics.

1.2 From non-equilibrium steady state dynamics to Bayesian mechanics

I have shown how the FEP defines something that exists as statistically separated from the environment. How, however, is the Markov blanket kept intact? Further characterization involves describing a dynamical system that entails a specific set of states (Friston, 2019). These kinds of systems can be described by stochastic differential equations such as the Langevin equation:

$$\dot{\mathrm{x}} = f(x) + \omega$$

This equation describes the rate of change (dot notation indicates the derivative with respect to time) of some states x, in terms of their flow $f(x)$ and random fluctuations ω. The next step is to describe the shape of the probability density of the states *p*(x) which can be interpreted as a likelihood of finding the system (e.g., an organism) in question in a specific (e.g., physiological) state at any point in time with a state sampled at random. Under this view, systems that gather themselves together, such as living organisms, can be cast as a random dynamical system that possesses a Markov blanket under the definition of the probability density p(x) which takes the name of *non-equilibrium steady state density* (NESS density). In other words, an open system at non-equilibrium steady state entails a pullback attractor, such as an attracting set of physiological states, and the trajectories of the system, over time, will revisit that attracting set (Arnold, 1995; Crauel & Flandoli, 1994; Friston & Ao, 2012). To describe the evolution in time of this steady-state density, an interpretation is provided in terms of *density dynamics* by using the *Fokker-Planck equation*:

$$\dot{\mathrm{p}}(x \mid m) = \nabla * (\Gamma\nabla - f)p$$

This equation describes the evolution in time of the probability density function $\dot{\mathrm{p}}(x \mid m)$ considering a state x and a Markov blanket m. The change of this density is described by random fluctuations and the flow of the states through the state-space where ∇ is the gradient, Γ is the variance of the fluctuations and f is the flow. Crucially, we want to consider systems that are defined by and thus restrict themselves to a particular number of states (Friston, 2019). For example, an organism that seeks to maintain a certain body temperature or to restrict the levels of blood glucose within a certain range. Therefore, we can set this probability density as not changing, meaning with a rate of change ẋ equal to 0. In this case, the Fokker-Planck equation entails a steady-state solution. This allows us to rearrange the equation and to intuitively observe that, when the density is not changing, the flow of the states must be countering the random fluctuations $-f(x) = \omega$.

$$\dot{\mathrm{p}}(x \mid m) = 0 \Rightarrow f(x) = (\Gamma - Q)\nabla \ln p(x \mid m)$$

A further step forward allows us to rewrite the solution to this equation according to Helmholtz's decomposition: we can express the flow of the states as a solenoidal flow Q (a conservative, rotational, divergence-free component) and a gradient flow Γ (a dissipative, irrotational, curl-free component), with respect to the steady-state density (Friston, 2019; Friston & Ao, 2012). These flows can be easily exemplified by picturing a stream of water flowing down a sinkhole where the basin's altitude is given by the log of the probability of occupying any state (see figure 1.2).

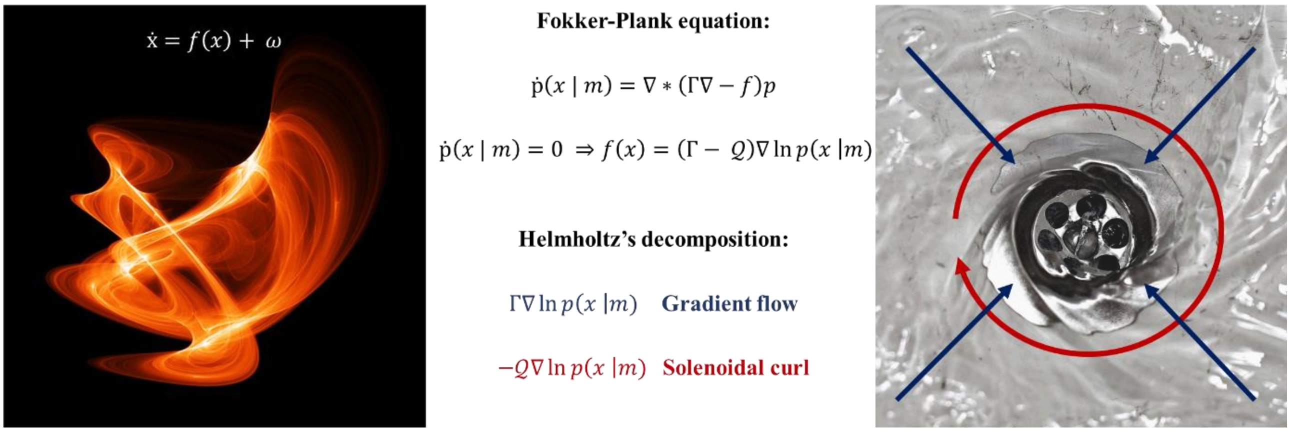

Figure 1.2. The left panel represents the trajectory of the random dynamical system with a NESS through the state space. Here we associate a probability density p(x), with that trajectory. The right panel represents the shape of the dynamics on p(x) in terms of the Helmholtz decomposition as two flows: one sinusoidal flow Q, in red, that goes around the contours of the log probability of the states and one curl-free gradient flow Γ, that counteract the random fluctuations by performing an ascent on the log probability. The two representations are connected by the Fokker-Plank equation.

In simpler terms, the states that an organism keeps revisiting are enclosed in a very limited set. The system is flowing towards and oscillating around those attracting states and this revisiting defines the steady state density and the kind of organism itself. Crucially, the attracting set corresponds to the conditions that allow for the organism's homeostatic integrity or, in other words, the *agent's phenotype* as the description of the characteristics compatible with its long-term existence (Corcoran et al., 2020; Friston, 2010; Parr et al., 2022). Therefore, living beings self-organize towards a specific set of physiological states as they constantly come back to that same neighborhood of states. This, however, doesn't mean that they aspire to be at thermodynamic equilibrium: there is complexity in this itinerancy as the dynamics of the internal states that generate behavior are not random (Friston, 2019). The solenoidal curl represents this constant oscillation, and it is manifest at every level of self-organization: from slow to fast electrophysiological oscillations in the brain, to circadian cycles, to hunger and satiation states and biological rhythms. On the other hand, the gradient flow is necessary to restrict the number of states as it counteracts the random fluctuations. These fluctuations would result in the dissolution of the blanket, and the system itself would cease to exist. Crucially, the system is always trying to flow in a way that 'looks as if' it is minimizing the negative log probability of the blanket states x (Friston, 2019). This quantity has a key meaning in information theory, known as surprisal or self-information (Shannon, 1948).

$$-\ln p(x)$$

This means that everything that exists, by possessing the attracting set of states, must look as if it is trying to minimize surprisal by gradient descent. From the perspective of the living being, this surprisal scores how much a sensory state differs from the attracting states that define where the agent's physiology can endure. For example, a surprising state for me would be my body temperature reaching 38° C as my attracting state for my phenotype is set at 37°.

From this observation, a 'new kind' of physics can be derived: Bayesian mechanics (Friston, 2019; Ramstead et al., 2022). According to it, the internal states of the Markov blanket can be seen as parametrizing a belief distribution over the external states. In other words, it 'looks as if' the system is performing a form of inference. In this is implicit a fundamental point: the system and the NESS density can be interpreted as a generative model since the highest probability region of the state space of the system will be a joint distribution over all the variables of the system. Under the Markov blanket, the generative model is the joint probability distribution over external and blanket states. This is supported by the statistical generalization of the *complete class theorem*, which states that, under certain assumptions, any process that follows a cost function in a finite sample space is Bayes optimal with respect to some prior (Andrews, 2021; Friston, 2019; Friston et al., 2022). Therefore, any dynamical system that minimizes some cost function (surprisal) according to some decision procedure is considered Bayes optimal under a generative model. This allows casting self-organization in terms of Bayesian inference (Friston, 2019).

Chapter 2

Self-evidencing architectures: generative models for perception and action

2.1 Generative models

In the most basic formalization, generative models are probabilistic models that can describe how (e.g., sensory) observations are generated or caused by events and objects in the environment that can't directly be accessed, thus often called *hidden states*. The simplest form of a generative model is the one typically described by Bayes' theorem where s represents the possible hidden states and o the possible observations. Thus, $p(s)$ is the prior belief over those states and the likelihood $p(o|s)$ specifies how the states generate the outcome.

$$p(s, o) = p(s) * p(o|s)$$

To understand how causes lead to consequences, the model has to be inverted: namely, the Bayes rule is applied to update prior beliefs into posterior beliefs, by inverting the likelihood mapping, starting from the probability of observations given the causes to infer the probability of causes given the observations.

$$p(s|o) = \frac{p(s) * p(o|s)}{p(o)}$$

On the denominator, the marginal likelihood or Bayesian model evidence $p(o)$ is necessary to obtain a distribution of just one of the two variables from the joint distribution. It is a measure of the model's goodness of fit as it represents the probability of generating the data from the prior. Crucially, Bayesian model evidence is linked to surprisal: they are inverse functions, thus, maximizing one corresponds to minimizing the other.

Minimizing average surprisal (entropy), correspond to maximizing the evidence for the agent model of the sensory exchanges with the world. This process is often called self-evidencing (Hohwy, 2016) and recapitulates one previously mentioned feature that stems from the FEP: every biological organism is a model of the environment in which they are immersed. For example, the phenotype of an agent represents the model of the environment, optimized by evolution, while the brain represents the optimized model of the econiche of one specific agent optimized through neurodevelopment and learning (Parr et al., 2022). Organisms are indeed inference machines that optimize the predictions of their models given the sensory data from the world.

The idea of the brain as a statistical machine that makes inferences about its sensations is ancient and was first articulated by the physicist Helmholtz. The problem of perception has been explicitly associated with a form of hypothesis testing (Dayan et al., 1995; Gregory et al., 1980) which in recent interpretations are formalized under the Bayesian brain hypothesis and predictive coding (Clark, 2013b; Knill & Pouget, 2004; Rao & Ballard, 1999). Under these views, perception corresponds to the application of the Bayes rule to obtain the posterior probability distribution of the hidden states of the world. A generative model like the one described allows formalizing perception as a probabilistic characterization of the sensory consequences of hidden states. However, I overlooked a crucial step: in most cases, evaluating surprisal or model evidence is an intractable problem. Indeed, exact Bayesian inference is usually mathematically unfeasible as it requires an impossible marginalization or integration problem over every possible state of the world, which is physically unrealizable (Beal, 2003; Friston et al., 2006; Parr et al., 2022). This is where the concept of variational free energy comes into play: *variational inference* or *variational Bayes* is a form of approximate Bayesian inference in which it is assumed a functional form for the posterior beliefs. This implies that it will depart from the form of the true posterior distribution and this departure renders this technique an approximation. However, thanks to this, tractability is earned.

2.2 Perception as variational inference

Under the FEP, perception is the process that minimizes an approximation of surprisal called variational free energy by updating beliefs about the hidden states of the world given sensory observations. This is obtained by performing gradient descent on variational free energy under a generative model, a process also known as variational inference (Bastos et al., 2012; Friston, 2005, 2010). Introduced by Richard Feynman (Feynman, 1972) in quantum physics, the concept of free energy is used to solve complex integration problems when computing model evidence and it has been exploited in statistics and machine learning (e.g., Neal & Hinton, 1998). The key feature is that, instead of trying to minimize something that cannot be measured, a measurable and always positive bound is created. Then, by minimizing the bound, the unknown quantity can be minimized too. The free-energy bound is created by use of a nonnegative quantity called Kullback–Leibler divergence.

This KL divergence measures the difference between two probability distributions: the first is an arbitrary distribution called *variational density* or *recognition density q(s)* and the second is the posterior density, which is the probability of causes after having observed their consequences. When Free energy is minimized, the variational density is rendered maximally similar to the true posterior probability *p(s|o)*. This can be done by optimization processes such as the iterative updating of *q* sufficient statistics through a gradient descent algorithm (e.g., variational message passing, Parr et al., 2022).

$$\underbrace{\mathrm{F}(s)}_{Variational\,FE} = \underbrace{\mathrm{D}_{KL}[q(s)||p(s|o)]}_{Divergence} - \underbrace{\ln p\,(o)}_{Surprisal}$$

This accounts for perception or state estimation as a form of belief updating, as the minimization of the discrepancy between the model and the world. This can be operationalized in terms of prediction errors where the optimal explanation of sensory data is the hypothesis that minimizes the prediction errors the most while diverging as less as possible from prior beliefs. By rearranging the terms, free energy can be written as:

$$\underbrace{\mathrm{F}(s)}_{Variational\,FE} = \underbrace{\mathrm{D}_{KL}\,[q(s)\,||\,p(s)]}_{Complexity} - \underbrace{\mathrm{E}_{q(s)}[\ln p(o|s)]}_{Accuracy}$$

The generative model must generate, via perception, the most possible accurate predictions (thus with minimal prediction errors) with the least possible complexity (as an Occam's razor that entails efficiency and prevents overfitting, Parr et al., 2022). However, perceptual inference by itself is unable to reduce surprisal as it only reduces the difference between free energy and surprisal by minimizing the bound. To counteract random fluctuations, self-organize e maintain homeostatic integrity, it is necessary to influence the external states of the world through action and movement.

2.3 Active inference

Under the FEP, the fundamental step further is to consider generative models that include action and the active sampling of sensory data: a biological agent can minimize surprisal by changing the things that are predicted, the data that are sampled. This is active inference (Parr et al, 2022). In this setting, our priors furnish the data that we expect to experience and we selectively sample the observations that realize that prior. In other words, an agent will act upon the world to make its predictions come true (Friston et al., 2013; Friston & Ao, 2012; Parr et al., 2022). Indeed, the generative model of an organism does not simply represent the dynamics of the environment but it also specifies the preferred conditions that support the agent's existence. These conditions correspond to the attracting set of states and are encoded in the prior beliefs of the generative model. Implicitly, these priors constitute preferred observations that are considered more likely to occur and entail less surprisal (Parr et al., 2022). Practically, the generative model of an active inference agent entails a prior over the trajectories of active states. These trajectories are known as policies π and consist of sequences of future observations. The generative model that also includes policies π is described by the following equation:

$$p(s, o, \pi) = p(o|s, \pi) * p(s|\pi) * p(\pi)$$

Crucially, this leads to a prospective form of inference about the trajectories of future observations. Thus, the generative model must consider the free energy of the future known such as the *expected free energy*. The model entails the prior beliefs about the sensations that will be encountered in the future, and the system will choose the most probable path to realize them (Parr et al., 2022). This, fundamentally, corresponds to a form of planning as the selection of action sequences known as *planning as inference* (Botvinick & Toussaint, 2012). The agent predicts future observations by performing 'what-if' simulations by generating possible futures (one for each policy π) by scoring each policy in terms of expected free energy and selecting the one that is expected to minimize free energy in the future. The expected free energy associated with each policy π considers the prior preferences of the agent (extrinsic or pragmatic value) and the expected information gain about states of the world (intrinsic or epistemic value). These terms can be expressed in terms of risk and ambiguity: risk is the (KL) divergence between anticipated outcomes given a policy $q(o|\pi)$ and preferred outcomes *(*$p(o)$. Ambiguity is the expected uncertainty (i.e., conditional entropy H) about outcomes, given the model's likelihood $p(o|s)$:

$$\underbrace{\mathrm{G}(\pi)}_{Expected\ FE} = \underbrace{\mathrm{D}_{KL}\left[q(o|\pi)\,||\,p(o)\right]}_{Risk} + \underbrace{\mathrm{E}_{q(S|\pi)}[\mathrm{H}[p(o|s)]]}_{Ambiguity}$$

These two components of expected free energy ensure that the agent plans and adaptively balances exploitation (i.e., preference-seeking) and exploration (i.e., information-seeking) as the agent is driven to both maximize utilitarian aspects of behaviour and resolve uncertainty by pursuing epistemic affordances (Parr et al., 2022). This furnishes self-evidencing with an enactive aspect and a minimal form of cognition (Corcoran et al., 2020). With active inference, then, self-organization starts to look like autopoiesis (Maturana, &. and Varela, 1980); Thompson & Varela, 2001).

The theoretical strength of active inference is that both perception and action serve the same objective function: they both minimize free energy by updating my model in the present and by changing the world to make it more similar to my prior beliefs in the future. This also addresses how perception itself is not a passive process but a constructive (generative) and active process where a perceptual hypothesis about the causes of the sensation is confirmed or disconfirmed by the sensory observation (Blake & Yuille, 1992; Friston, Adams, et al., 2012; Yuille & Kersten, 2006). Another strength of active inference resides in its biological plausibility. For example, in contrast to reinforcement learning, where the definition of an arbitrary reward function is required, in active inference, value is determined by first principles as it is an emergent propriety of self-organization (Parr et al., 2022). Furthermore, it incorporates epistemic drives, artificial curiosity and information seeking under a unique solution, without the need for supplementary information gain bonuses. At last, as I will show, more complex aspects of cognition can emerge when the generative model becomes hierarchical by generating predictions about components of the generative model itself.

Chapter 3

Hierarchical architectures: active inference in deep temporal models

3.1 Active inference in discrete state spaces: partially observable Markov decision processes

In active inference, a task is represented by simulating a cognitive agent that makes observations in an environment, takes decisions and acts upon it. This environment is represented as a discrete state space where sets of hidden states *s*, called *states factors*, are defined as beliefs or discrete probability distributions over unobserved task variables, such as the agent's position in the environment, and set of outcomes *o*, called *outcome modalities*, are the observable stimuli, from which the agent can infer hidden states. The temporal evolution of the states is discretized in a number of time points. This kind of modelling takes the name of *partially observable Markov decision process* (POMDP) and entails two general features. The first is 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, FitzGerald, et al., 2017b). POMDP can be represented by using graphical models: figure 3.1 depicts a *Bayesian network* where the circles, called nodes, correspond to variables (e.g., observations, hidden states, and policies), while the arrows connecting the nodes entail edges that explain the kind of dependencies, such as the probabilistic relations, among the variables. For example, the arrow that goes from the node state at timepoint 1 to the node outcome at timepoint 1 means that the value of the latter depends on the value of the former according to the probabilistic mapping described by the edge **A**, which encodes a likelihood matrix.

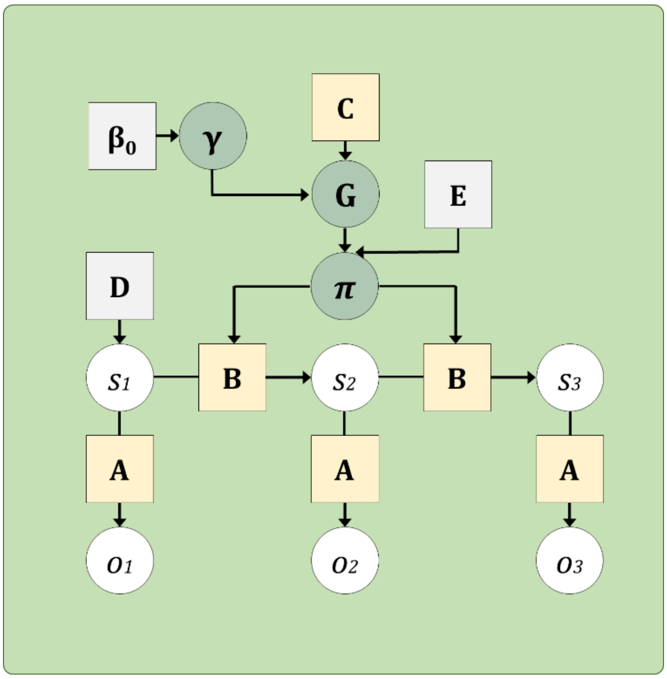

Figure 3.1 Bayesian network of an active inference partially observable Markov decision process. The circles correspond to variables: observations, hidden states, and policies; the arrows connecting the nodes are equipped with edges that entail the probabilistic relations among the variables. In this example, the model entails 3 timesteps, thus, starting from an initial state s_1 the policy π sets the trajectory for two transitions among states.

More specifically, the **A** matrix represents the (likelihood) mapping between states and outcomes. The **B** matrix (transition function) encodes the probability of moving from one state to another. The **C** matrix encodes prior beliefs about observations, which in active inference reflect prior preferences. The **D** vector encodes the prior about the initial hidden state. The **E** vector encodes priors about policies. The **G** denotes expected free energy. The γ is a precision parameter associated with expected free energy which encodes confidence in the current **G** estimates; Finally, $\boldsymbol{\beta_0}$ is a prior over the precision. See also Table 1 for further details.

Symbol	Mathematical formulation
π: beliefs about policies	$\pi = \sigma(\ln E - \gamma G)$
G: Expected free energy (used to score policies during planning), in terms of risk and ambiguity in matrix notation.	$\mathbf{G}_\pi = \underbrace{\mathbf{A}s_{\pi,\tau} * (\ln \mathbf{A}s_{\pi,\tau} - \ln \mathbf{C}_\tau)}_{Risk} \underbrace{-diag\,(\mathbf{A}^{\mathrm{T}} \ln \mathbf{A}) * \mathbf{A}s_{\pi,\tau}}_{Ambiguity}$
A: likelihood matrix that encodes the statistical relations between hidden states	$p(o \mid s)$

and outcomes. It maps how states (causes) generate outcomes (consequences)	
B: transition matrix that encodes the statistical relations between subsequent states (e.g., beliefs about how states change over time).	$p(s_{t+1} \mid s_t)$
C: matrix encoding prior preferences over observations. Observations with higher probabilities are desired, e.g., the attracting set	$p(o\|\mathbf{C})$
D: vector encoding the prior belief about hidden states at the first time point	$p(s_1\|\mathrm{D})$
E: vector encoding the prior belief about policies. Beside planning, some policy can have higher probability	$p(\pi)$
γ: precision parameter about expected free energy **G** which encodes confidence about the current **G** estimates.	γ
β_0: hyperparameter that determines the initial value of γ	$\beta = 1 / \gamma$

At every timepoint a policy is selected according to the following equation, which represents a (prior) distribution over the policies:

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

This distribution comprises two components. The former is the prior over the policies (an **E** vector) and the latter is the expected free energy **G**. The balance between these two terms is determined by a precision term (γ), which encodes the confidence of beliefs about **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. Subsequently, policy-dependent transition probabilities encoded in the **B** matrix, or the **D** vector if it is the first time point, generate hidden states which then generate observations at each time point according to the likelihood **A** matrix. After receiving an observation, the equation also incorporates the variational free energy **F**, in the posterior over policies:

$$\pi = \sigma(\ln \mathrm{E} - \mathrm{F} - \gamma \mathrm{G})$$

At last, an action u and the associated transition between states is selected according to the following equation, where $\boldsymbol{\alpha}$ represents a precision parameter that controls the randomness of the action selection mechanism:

$$p\ (u\ |\alpha)\ =\ \sigma(\alpha * \ln p(u|\ \pi))$$

As I will show in the studies, the structure of these models can be extended to include a hierarchical level that makes inferences about lower-level state transitions. These are called deep temporal models.

3.2 Hierarchical active inference: deep temporal models

Every biological agent embodies a specific generative model that accounts for the her adaptive needs and cognitive capabilities at various levels of complexity (Corcoran et al. 2020). Very simple generative models account mostly for reflexive and reactive behavior by being equipped with a target setpoint (such as an expected physiological state) that reacts to a prediction error (Adams et al., 2013). Generative models that entail a minimal form of cognition, planning, or agency, as the one described in the previous section, are equipped with the prior belief that they will minimize free energy in the future and they entertain the possibility of alternative (or counterfactual) trajectories of behavior. This, as I said, is at the basis of allostatic control, autopoiesis or *anticipatory homeostatic control* (Corcoran et al., 2020; Parr et al., 2022; Pezzulo et al., 2015). However, these models afford inference at only one timescale and are too shallow to account for complex cognitive functions in a more biologically plausible way. On the other hand, hierarchical generative models allow for inference at multiple timescales. These kinds of models, also known as deep temporal models, entail a higher hierarchical level that represents the dynamics of the lower level (Friston, 2008; Friston et al., 2018; Kiebel et al., 2008; Pezzulo et al., 2018). The higher level informs and updates the lower level while the lower level observes the upcoming state transitions. Consider the example of reading: the encoding of letters occurs at a faster rate and subordinates the encoding of the word, which in turn is faster and subordinates the encoding of the sentence. This kind of processing highlights the necessity of a temporal hierarchy which allows to go beyond the immediate demands of homeostasis regulation and further extend the agent's autonomous horizon (Friston et al., 2018; Hesp et al., 2021; Pezzulo et al., 2018).

In practice, a hierarchical active inference model consists of two models, one placed above the other and then linked in a bidirectional manner. Specifically, at a higher level, each time-point furnishes the initial states for the first time-point of the level below. On the other hand, the posterior belief at the last timepoint of the level below provides an observation for the associated timepoint at the higher level. The lower level operates at a slower timescale than the higher-level model because each observation at a higher level must ‘wait’ for the whole sequence of time points at the lower level. See figure 3.2 for a hierarchical Bayesian network.

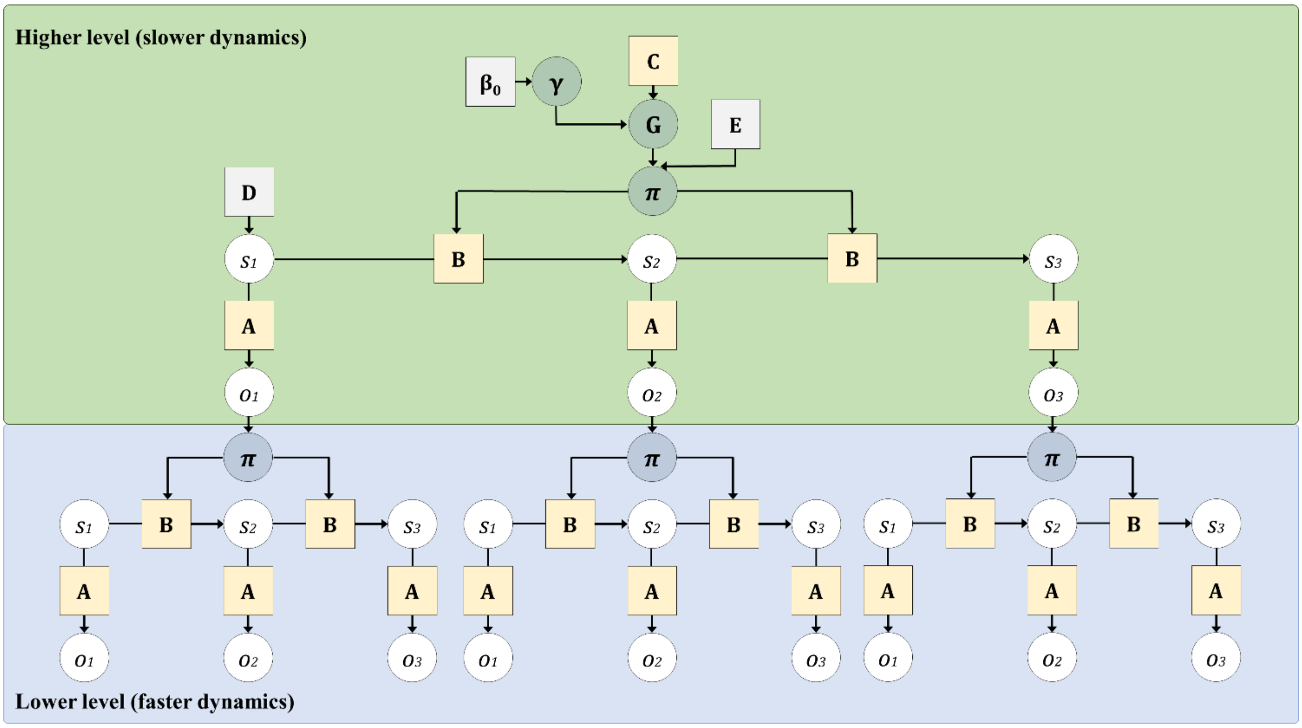

Figure 3.2. In this hierarchical Bayesian network, every timepoint at the higher level generates prior that influences the dynamics at the lower level. On the other hand, the posterior beliefs at the lower level are propagated bottom-up and become observations at the higher level (Friston et al., 2018). This provides bidirectional interactions between variables at different levels of the hierarchy because the inference of each variable is simultaneously influenced by messages passed from different directions.

This kind of deep inference allows us to infer events that change at different timescales, from slower dynamics at higher levels to faster dynamics at lower hierarchical levels and to simulate agents that entail more sophisticated strategies for planning and hence to represent more complex cognitive phenomena (Parr et al., 2022). These kinds of architecture are necessary to encode perceptual events with nested dynamics where a set of percepts must be recognized before the regularities among those percepts can be detected (Smith, Friston, et al., 2022). Furthermore, through these architectures, narratives and hierarchical representations such as those typical of language can be encoded as the higher level of the model generate predictions about the deeply structured sequence of state transitions at the lower level. In the next chapter, I will introduce the first application of this idea, in the context of action representations.

Chapter 4

1^{st} Architecture: an active inference model of hierarchical action understanding, learning and imitation[1]

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 explain the neurobiological underpinnings of visuomotor cognition, including the multiple routes for action understanding in the dorsal and ventral streams and mirror mechanisms.

4.1 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 & 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 & Gergely, 2007; Jacob & 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:

$$p(goal \mid movement) = p(goal) * p(movement \mid goal) / p(movement)$$

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

---

[1] Paper submitted to *Physics of Life Reviews*

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., 2006, 2009; 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 & Russon, 1998; Grafton & 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 4.1. 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

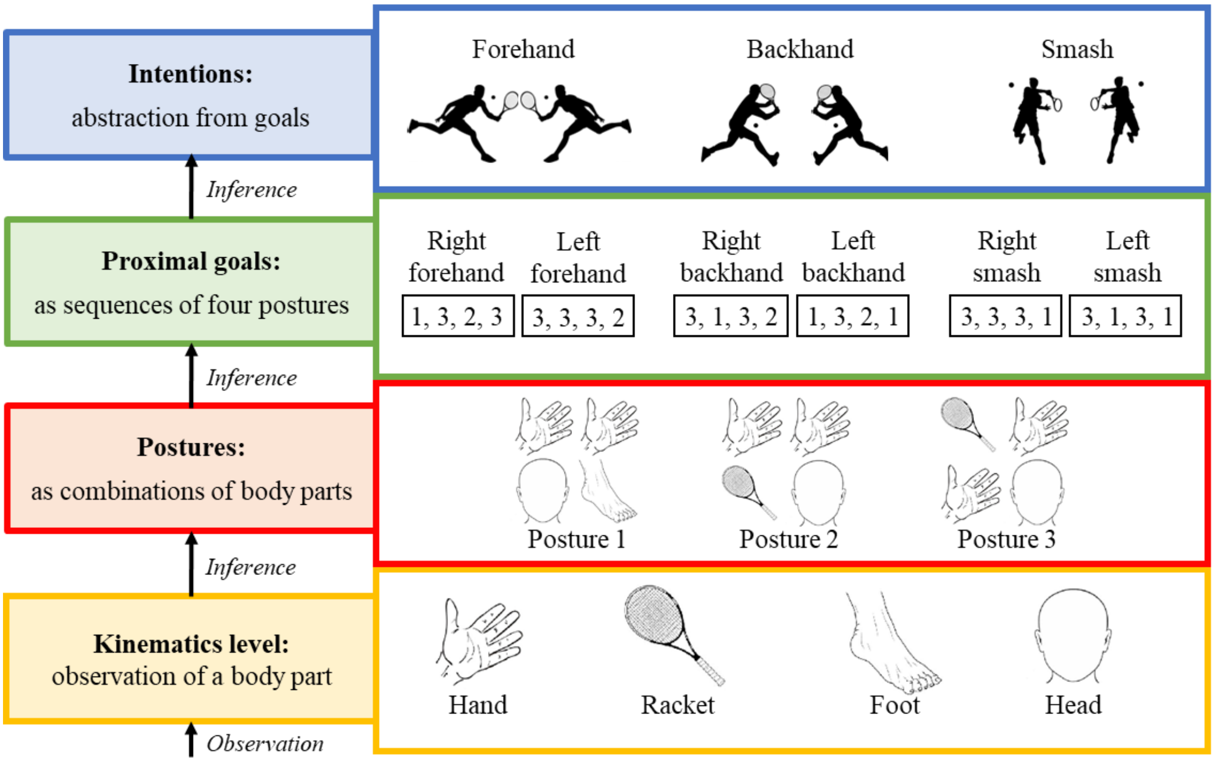

Figure 4.1. This schematic illustrates the hierarchical representation of an action repertoire in the example scenario that we adopt in this paper: playing tennis. The hierarchy comprises three levels of hidden variables (postures, proximal goals, and intentions) and an observation layer (kinematics). 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 4.1 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 & Chater, 2005).

This study aims 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 how 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.

4.2 Active inference simulations during action understanding and imitation

4.2.1 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 formally describes how the learner performs action understanding and imitation, by observing how the teacher's body movements unfold over time.

As illustrated in Figure 4.2D, 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.

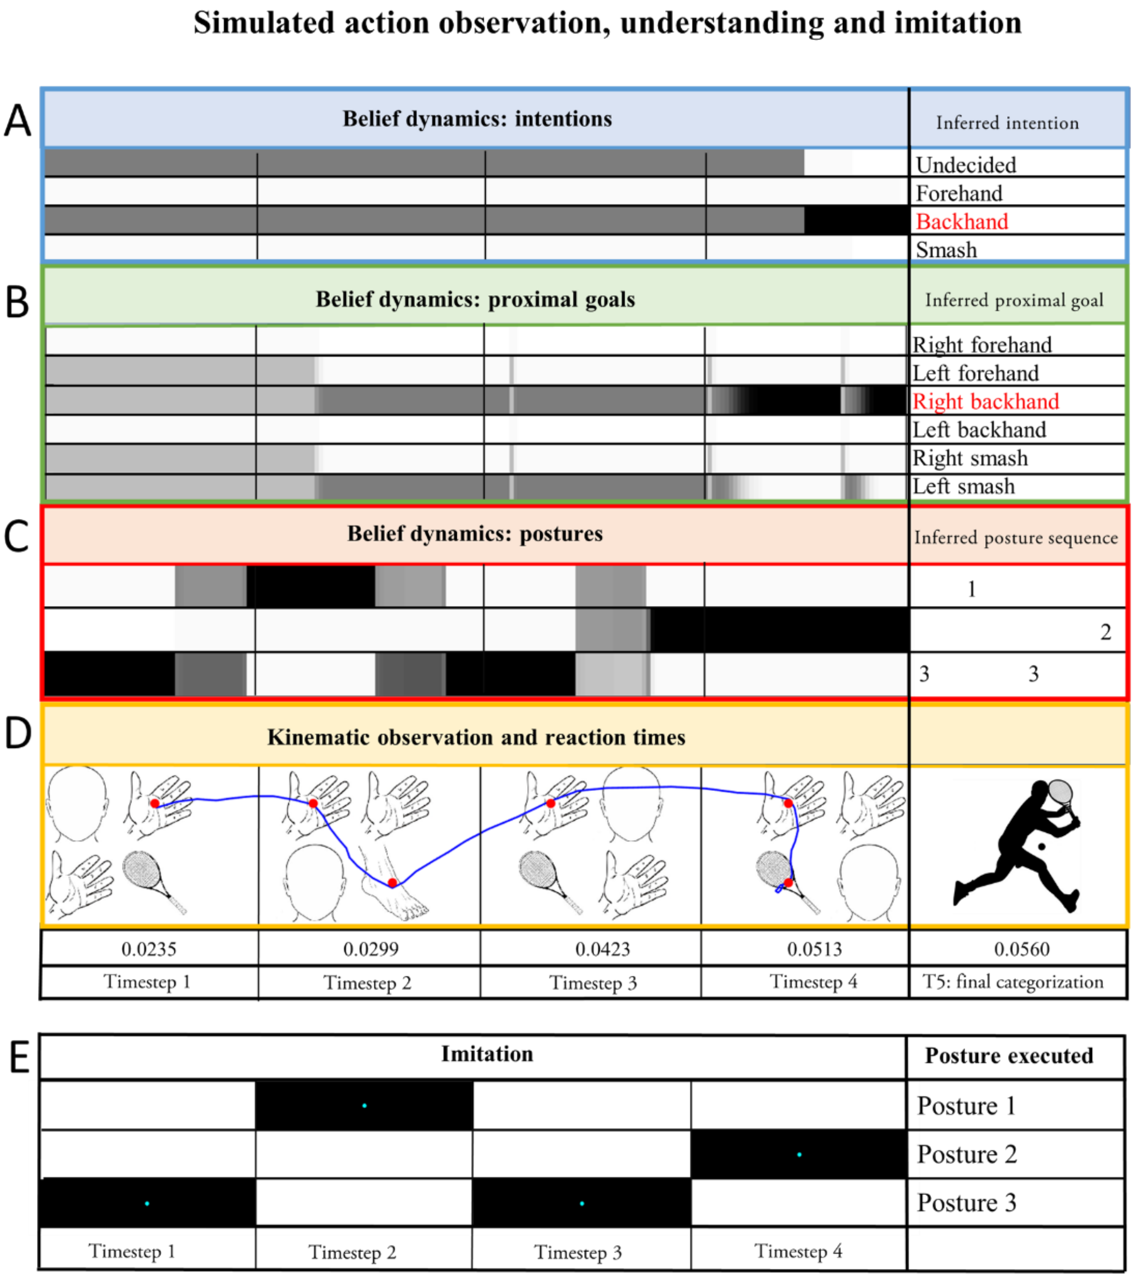

Figure 4.2. Results of Simulation 1: action observation, belief updates during action understanding and imitative motor response. The Panels ***A-D*** *correspond to the four hierarchical levels of the model shown in Figure* 4.1. *Here, we describe them in the reverse order, starting from the bottom (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 inference of the three hidden variables (intentions, proximal goals, and postures) based on observation of movement (body part kinematics). In turn, this hierarchical inference can be conceptually divided into three kinds of inferences, see Figure 4.2. 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 4.2D 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 4.2C 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 4.2B shows the belief dynamics associated with the second kind of inference, which regards a deeper aspect of the observed movements: the one that concerns *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 4.2B 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 4.2A shows the third kind of inference, which regards the most general aspect of the observed movement: the intention. The specific example illustrated in Figure 4.2A regards the inference of the *backhand* intention. Finally, Figure 4.2E 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.

4.2.2 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, Adams, et al., 2012; Friston et al., 2015; Parr & Friston, 2017). Salience is determined by predicting where to perform the next fixation to help reduce 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 gathering information to reduce uncertainty (Blake & 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 4.3). In one condition, illustrated in the left panels of Figure 4.3, 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 4.3, 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 the most salient location to gaze depends on the learner's hypotheses or prior beliefs (which are updated after gathering novel observations). As an effect, the patterns of eye movements and belief updating of the left panels of Figure 4.3 and the right panels of Figure 4.3 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.

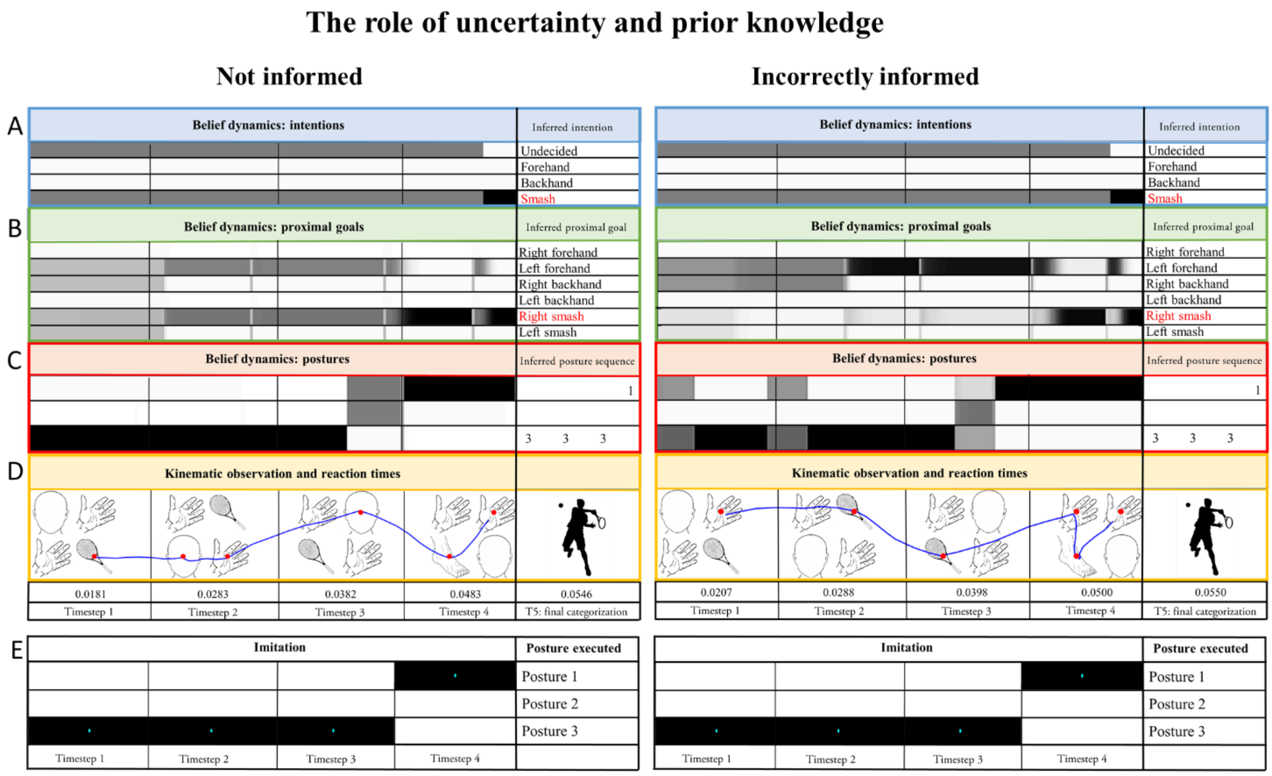

*Figure 4.3. Results of Simulation 2: the dynamics of action understanding of two learners, with no prior knowledge (left learner) or incorrect prior knowledge (right learner). The **D** panels show the different oculomotor dynamics for the two cases, respectively. The **A**, **B** and **C** panels show the different belief dynamics for the two cases, with the same notation as in Figure 4.2. 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, Adams, et al., 2012; Friston et al., 2015). 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 & Baldi, 2009; Itti & Koch, 2000). Using this method, a salient location corresponds to the one expected to reduce uncertainty, if a saccade were performed towards it (Parr & Friston, 2017). Empirical research shows that this notion of Bayesian surprise characterizes well which aspects of a visual scene capture human attention (Itti & Baldi, 2009; Mirza et al., 2016; Posner & 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, based on the learner's 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 motor knowledge (Flanagan & 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., 2005, 2006; Cross et al., 2006; Maranesi et al., 2014, 2019; Pedullà et al., 2020; Pezzulo, 2013; Rizzolatti & Craighero, 2004). Later, we will elaborate further on this topic.

4.2.3 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 first 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 4.2) and novelty (Figure 4.4) 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 does not know 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 4.4, 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 & Oudeyer, 2018; Schwartenbeck et al., 2019).

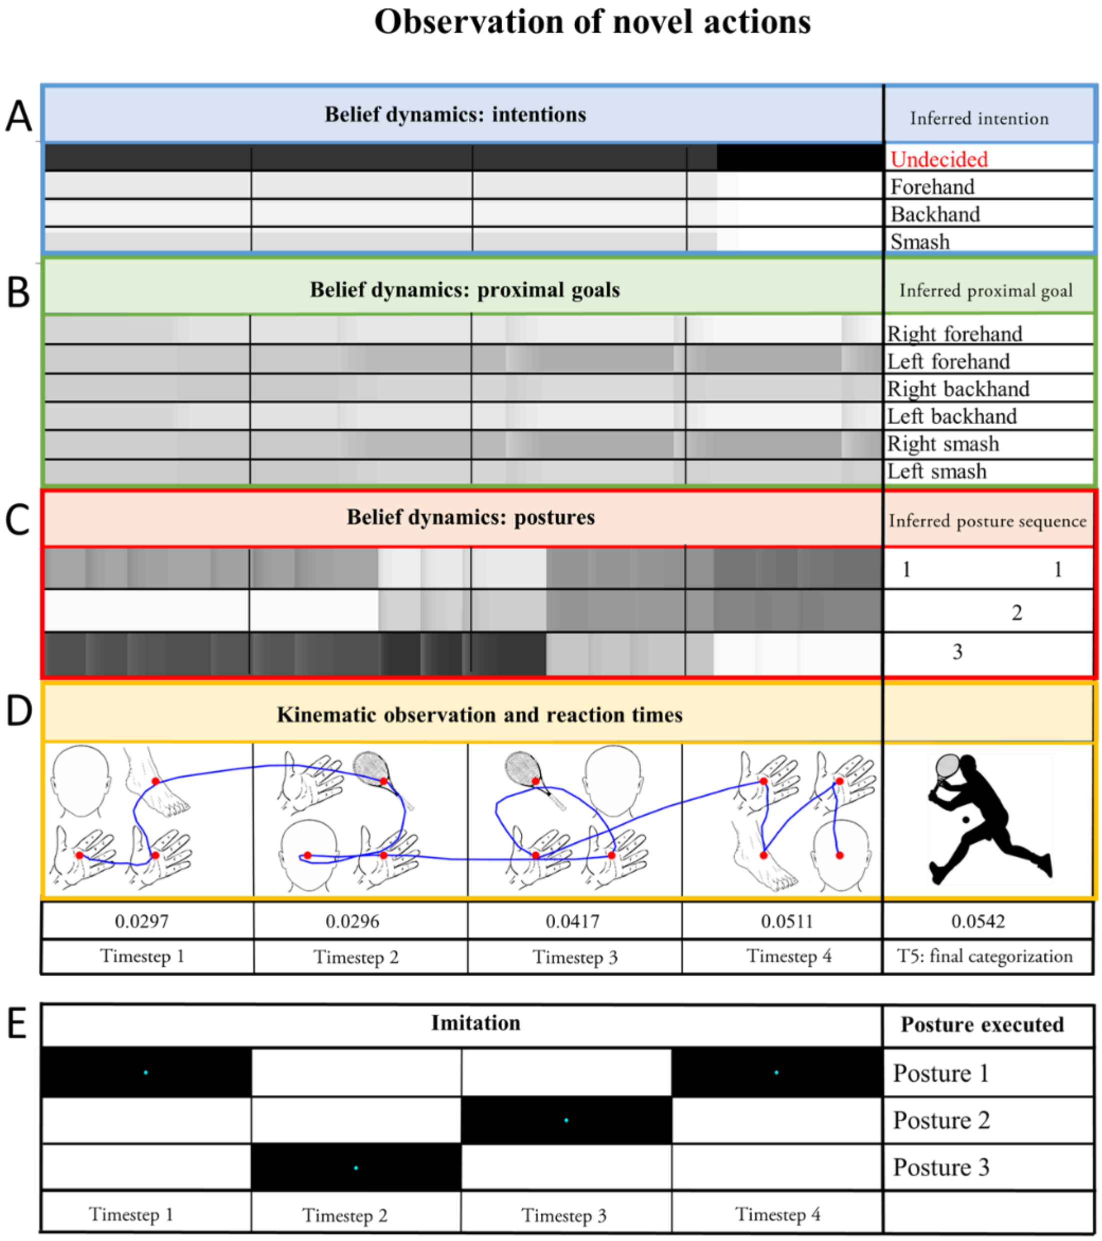

Figure 4.4. Results of Simulation 3: observation and imitation of a novel action. In panel ***D*** *are shown the simulated oculomotor dynamics when the learner observes an unknown action, here corresponding to the sequence of postures 1-3-2-1. In this case, visual attention is driven by both surprise 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 panels* ***C*** *and* ***B****, it can be noted how the probabilistic*

beliefs are more spread across the hidden states resulting in a much larger amount of uncertainty. This leads the student to an "undecided" intention categorization in panel ***A****.*

Figure 4.5 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 4.5 encodes the probabilistic relations between the levels of the postures and the proximal goals. Figure 4.5A 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 quadrant location 1; the second three columns represent the postures at 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 corresponds to the fact that they have a high probability). Figure 4.5B shows the learner's *generative model* before learning (to put it simply, her knowledge about the environment's statistics). This generative model is analogous to the generative process, but the row corresponding to the *left backhand* is grey, which indicates that the learner does not know the sequence of postures 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 4.5C 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.

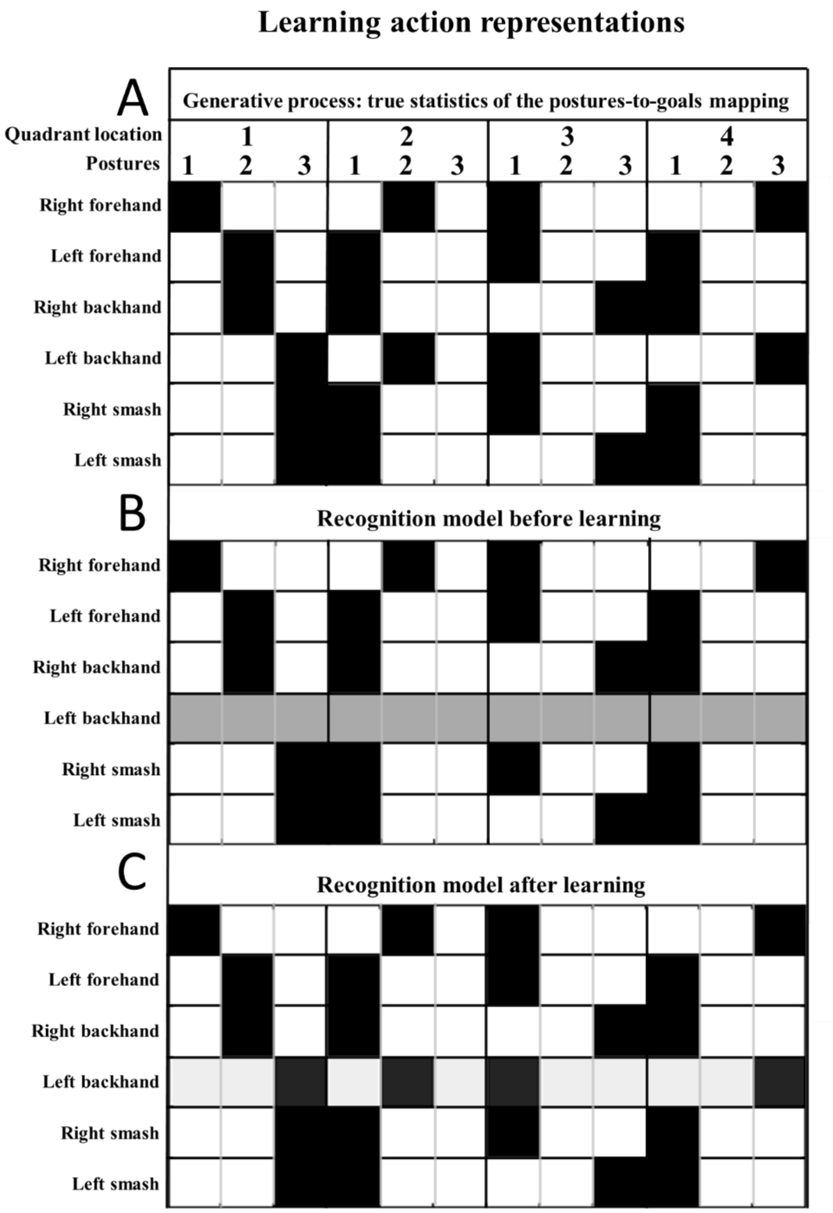

*Figure 4.5. Results of Simulation 3: how the generative model is updated when a novel action is learned. In this example, the part of the generative model that is updated is the (likelihood) mapping between postures and goals. For each panel (**A**, **B**, **C**), the rows represent the goals and the 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). 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. Panel B represents 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, since probabilities are spread). 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).

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 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 4.1. 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 & Damasio, 2008; Galetzka, 2017; Meteyard et al., 2012), as shown in the context of language (Borghi et al., 2019; Borghi & Binkofski, 2014; Borghi & Cimatti, 2010; Davis & 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.

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.

4.2.4 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 & 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., 2005, 2006; 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 & Sinigaglia, 2010); on a "motor simulation" of the observed action in one's nervous system (Aglioti et al., 2008; Calvo-Merino et al., 2005, 2006; Cross et al., 2006; Gergely & 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 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 4.6), 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 4.6), the learner is instructed to observe the left smash action but to produce an incongruent action (a left forehand).

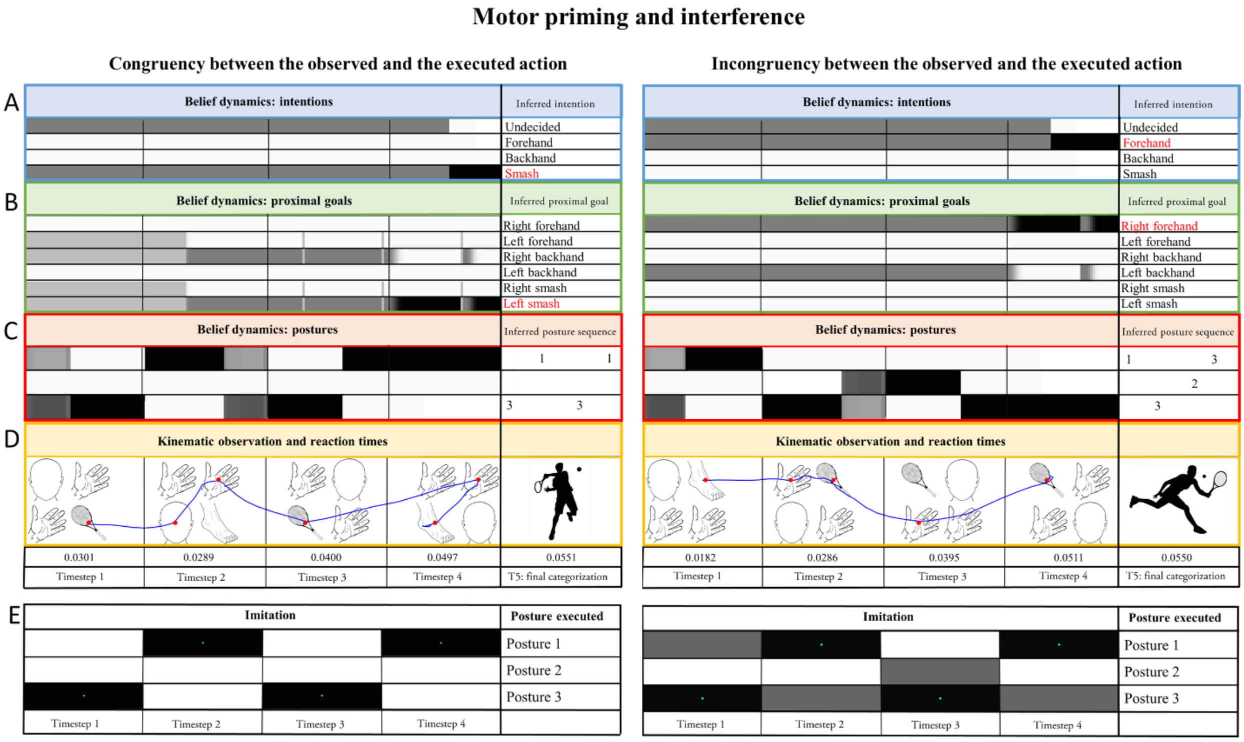

Figure 4.6. Results of Simulation 4: the dynamics of motor priming (left panels) and interference (right panel) for two learners. The left learner observes a left smash and imitates the same action (motor priming). The right learner observes a right forehand but executes a left smash (motor interference). 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 correspond to the observed sequence 1-3-2-3.*

Figure 4.6 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 the 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.

4.2.5 Summary and discussion of the simulations

The simulations described in this section help illustrate the key mechanisms of active inference that are relevant to modelling 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 & de C. Hamilton, 2007; Jacob & 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, 2005, 2008; Friston, FitzGerald, et al., 2017b; Parr et al., 2022; Pezzulo et al., 2018a, 2021).

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 (particularly from the expected free energy equation, see the Methods). In both cases, the observation policy maximizes information gain; however, the type of uncertainty differs when familiar and novel actions are observed. 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). In contrast, 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 inferring the hidden states that generate the outcomes. In contrast, parameter exploration predicts that the learner actively seeks *novel* combinations of hidden states and outcomes, because this determines the learning of how outcomes are generated. Only via this novelty-seeking behaviour can the learner 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 as 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 without an integrated action representation, but this would not correspond to a deep semantic representation in our account.

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 as 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.

4.3 Neurobiological underpinnings of the proposed model

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

4.3.1 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, which is particularly 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 acquiring representations (as we showed), surprise relates to improving 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 to 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. Studies support our prediction 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., 2006, 2007, 2021, 2022; Tessari & 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., 2016, 2019; Bartolo et al., 2001; Cubelli et al., 2000; Goldenberg & 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, Nitschke, et al., 2016a; Martin, Dressing, et al., 2016; Tessari et al., 2007, 2021). 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 & Buxbaum, 2013; Hoeren et al., 2014; Martin, Dressing, et al., 2016; Mengotti et al., 2015; Tessari et al., 2007, 2021). 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, Nitschke, et al., 2016a; Martin, Dressing, et al., 2016; Rijntjes et al., 2012; Rumiati et al., 2005; Tessari et al., 2021; Weiller et al., 2009, 2011b), 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 us to infer the possible functions from structure (Binkofski & Buxbaum, 2013; Goldenberg & Hagmann, 1998; Heilman et al., 1997) and can discover the alternative functions of familiar tools (Heilman et al., 1997; Roy & 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 4.7 for a graphical illustration of this proposed taxonomy.

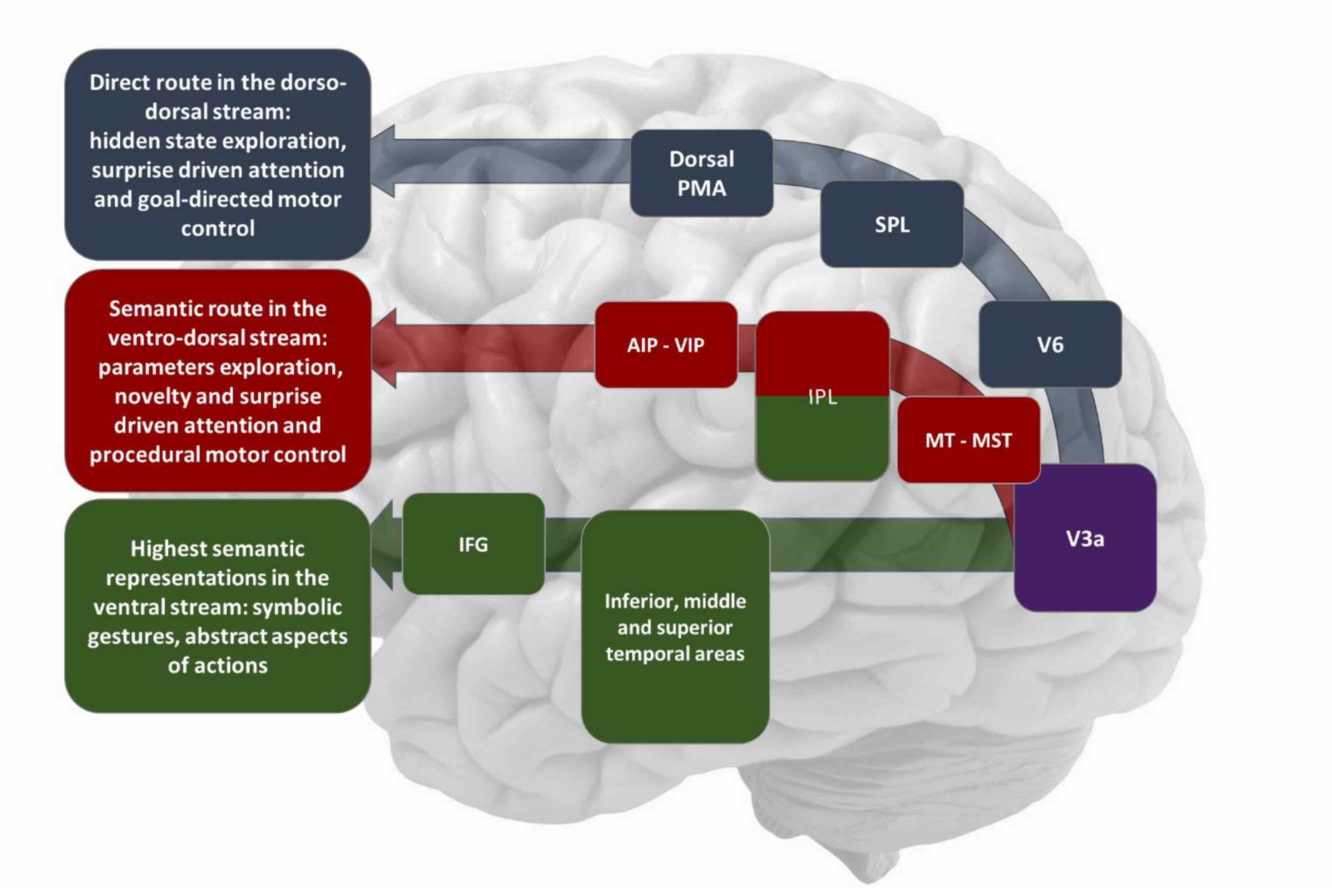

Figure 4.7. A plausible neural implementation of the proposed hierarchical model (this is a modified version of the neuro-cognitive model proposed by Tessari, Proietti & Rumiati, 2022, on Cognitive Neuropsychology). Visuomotor cognition is distributed onto multiple routes in the dorsal and ventral pathways. 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 an explanation.

4.3.2 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 & 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 & 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(goal | 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 predicts the sensory input related to that action (Friston et al., 2011; Kilner et al., 2007). Hence, predictive action monitoring is the functional role associated with the mirror neuron system. Motor resonance is not achieved by a direct matching mechanism but by emulative action inference or motor, embodied, simulation (Csibra & Gergely, 2007; Gallese & Goldman, 1998; Jeannerod, 2001, 2006; Pezzulo, Barsalou, et al., 2013). 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.

4.4 Methods

4.4.1 The hierarchical 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 4.8. The hierarchical arrangement of the generative model shown in Figure 4.8 reflects – and formalizes – the hierarchical organization of action representations shown in Figure 4.1; see Friston et al., 2018; Mirza et al., 2016 for related models in active inference and Botvinick, 2008; Friston, 2008; Kiebel et al., 2008, 2009; Lee & Mumford, 2003 for a broader view of hierarchical models in the brain.

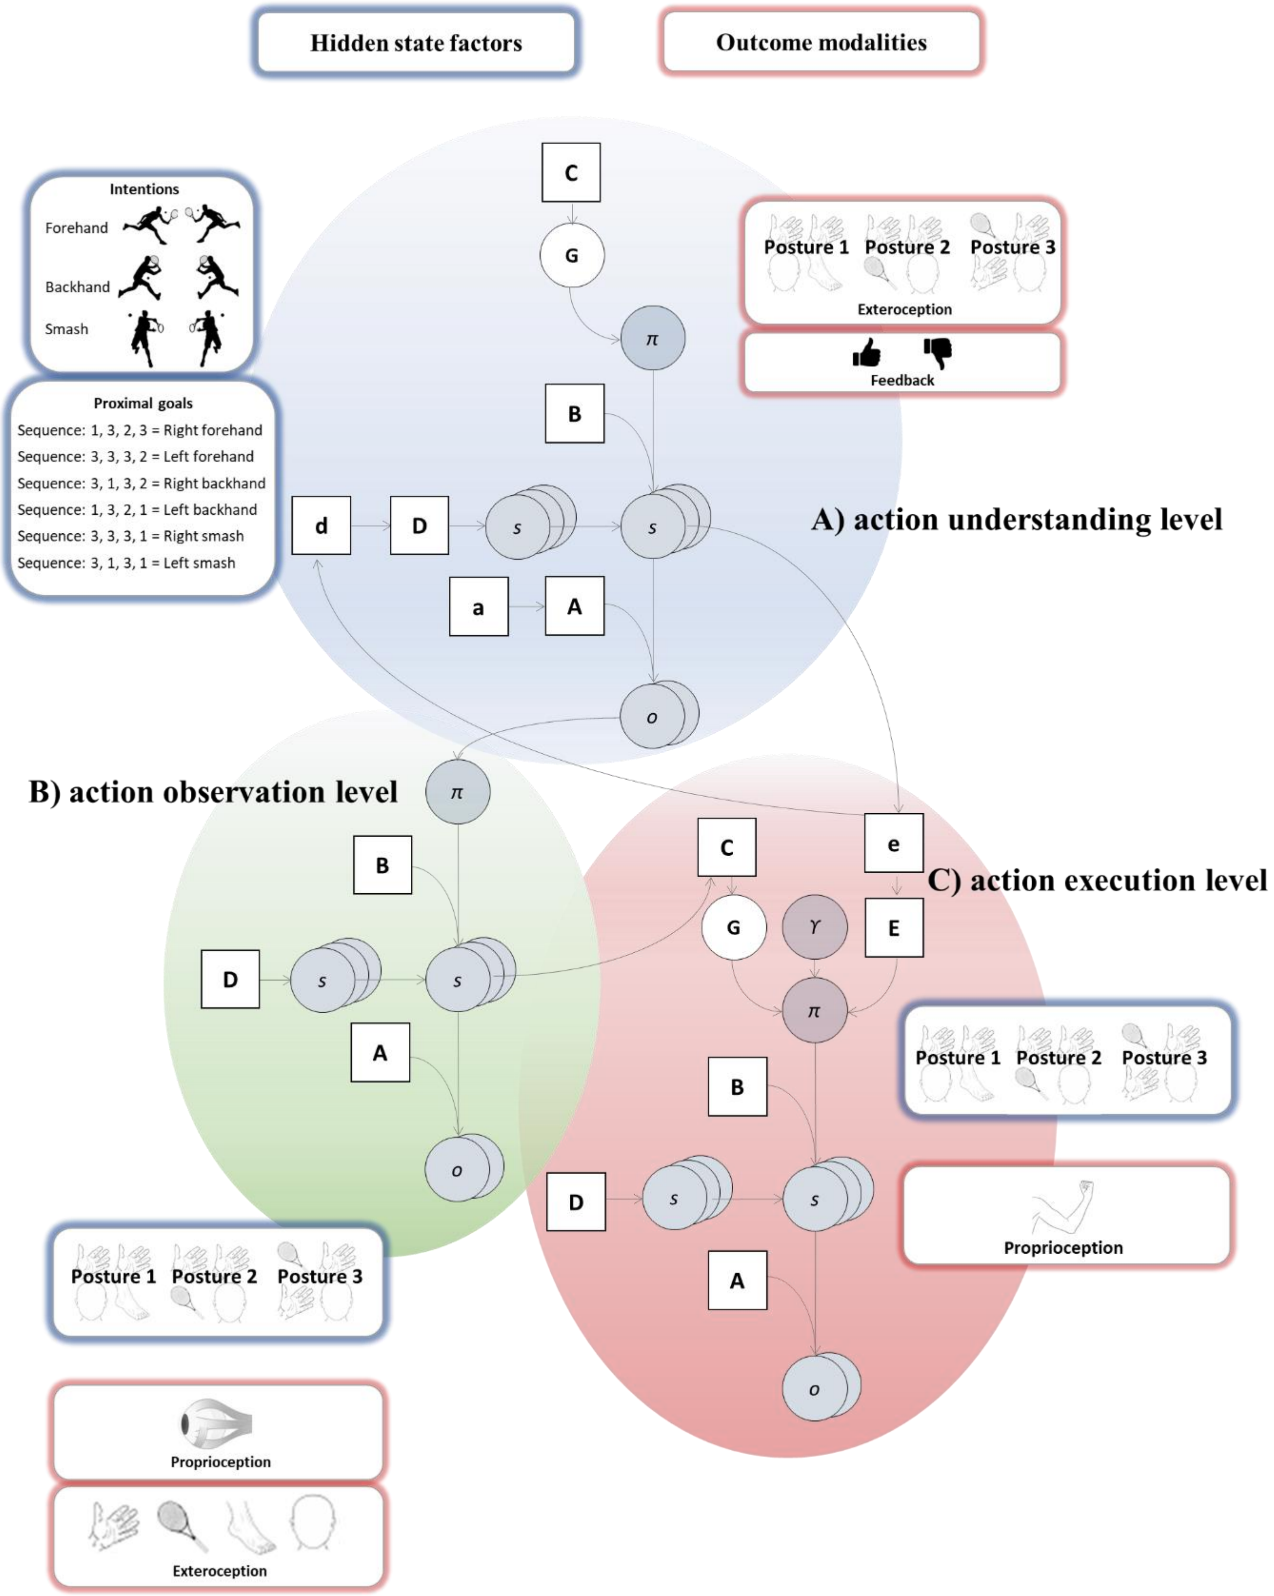

Figure 4.8: Schematic of the hierarchical generative model used in this article. (A) The Action Understanding level. At this level, there are two hidden state factors and two outcome modalities. The first factor is the proximal goal performed by the teacher (the sequence of postures) while the second factor is the categorization of the intention. The two outcome modalities are exteroception and feedback. (B) The Action Observation level. At this level, two hidden state factors and two outcome modalities exist. 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, FitzGerald, et al., 2017b). 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), and* ***E*** *(prior about policies) matrices. The* ***G*** *denotes expected free energy. The lowercase letters* ***a****,* ***d*** *and* ***e*** *denote the 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 4.8, *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's 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.

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 4.1). The Action Observation level models the recognition of movement kinematics and postures (hence formalizing the functioning of levels 1 and 2 in Figure 4.1) 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 state 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, each of which 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 among all 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 is not temporally linked to the other levels.

The generative model shown in Figure 4.8 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** matrices), 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). As described in the introduction, 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 $p(s|o)$ optimizing it by the iterative updating of its sufficient statistics through, for example, a gradient descent algorithm. Furthermore, 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, FitzGerald, et al., 2017b; Friston & 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 4.2-4.6 can be equated to the dynamics of neural activity of these groups of neurons, which evolve in a way that minimizes free energy.

4.4.2 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 sequences of postures that it has seen most often in the past (see below) and that are encoded into prior beliefs over the initial states **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, Adams, et al., 2012; Parr & 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 (**a** parameters of the) **A** matrix that links hidden states (proximal goals) and outcomes (postures) at the Action Understanding level. Crucially, when the model is allowed to learn (**a**) parameters, the expected free energy equation has to be expanded to include a new (novelty) term:

$$\underbrace{\mathrm{G}(\pi)}_{Expected\ FE} = \underbrace{\mathrm{D}_{KL}\left[q(o|\pi)\ ||\ p(o)\right]}_{Risk} + \underbrace{\mathrm{E}_{q(s|\pi)}[\mathrm{H}[p(o|s)]]}_{Ambiguity} - \underbrace{\mathrm{E}_{p(o|s)q(s|\pi)}[\mathrm{D}_{KL}[q(\mathbf{A}|o,s)\ ||\ q(\mathbf{A})]]}_{Novelty}$$

The novelty term in this expanded equation scores how much the beliefs in the **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 form of active learning, during which the active inference agent will preferentially look at (novel) locations that it expects to change the **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 (ambiguity), pursuing novelty is useful to reduce uncertainty about model parameters (sometimes called ignorance).

4.4.3 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 **A** matrix. This has been done by setting a column that encodes one of the proximal goals (*left backhand*) in the **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 of 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 the sequence of postures, by accumulating Dirichlet distributions in the **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.

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

The model determines the selection and control of imitative responses by two components. The first component is the posterior belief about the proximal goals, as encoded by the prior beliefs over the policies (**E**) and the associated **e** parameter at the Action Execution level. The **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 (**C**) and hence, in turn, the expected free energy (**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 & Zimmer, 1984, 1985, 1989) as it uses cached (**E**) values that are learned over time and the latter a *goal-directed* form of control as it uses expected free energy (**G**) computations that take the agent's prior preferences (**C**) into account.

The specific action executed depends on the weight given to the procedural (**E**) and goal-directed (**G**) components. The relative influence of the two components is determined by the parameter γ (gamma) which corresponds to the precision of beliefs about the expected free energy. This process of precision weighting corresponds to attentional modulation at the cognitive level and synaptic gain at the neurophysiological level (Clark, 2013a; Feldman & Friston, 2010; Parr & 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 (**E**) of the action sequence, while goal-directed control generated by **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, Niv, et al., 2009; Dezfouli et al., 2014; Dezfouli & Balleine, 2013; Miller et al., 1960). In active inference, this sequential knowledge is encoded in a (**V**) matrix that relies on 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 imitative responses are more driven by procedural knowledge in the former case, 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 often been related to a difference between model-based and model-free methods of reinforcement learning (Daw et al., 2005; S. W. Lee et al., 2014; Maisto et al., 2019; Pezzulo, Rigoli, et al., 2013). More of this is shown in chapter 6.

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 to execute imitative actions. Hence, it would trigger an imitative response that interferes with the correct action; see the right panel of Figure 4.6.

4.4.5 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 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 & Heyes, 2005).

4.5 General discussion

We have shown 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, pursuing 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 acquiring 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. All these proprieties emerge naturally from the process of active inference as vastly described in the literature. What is unique in our proposal is the generative model's particular (deep) structure and how messages are propagated from one block to another. On this account, an element of novelty is how the posterior beliefs from both levels of understanding contribute to action selection at the Action Execution level and how the posteriors about the executed actions are propagated back to the understanding levels (see the method for a detailed description of the model). This organization of the model recapitulates and supports several experimental findings while offering a formal perspective under a unique solution. Therefore, the model permits the formulation of some empirical predictions that can 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 contributions to the attentional capture of novelty and surprise (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 clue to further investigations into the phenomenon of mirroring or motor resonance as the effect seems to be stronger for lower levels of representations (e.g., object-oriented actions) and disappears as the action becomes more symbolic and abstracted (Csibra & Gergely, 2007). Another set of model predictions regards the neuronal dynamics that we would expect to observe during action observation and imitation, if they correspond to inferential processes, as our active inference model assumes. The belief dynamics plotted in Figures 4.2-4.6 can be easily mapped into (simulated) neuronal population dynamics, using the methods illustrated in Friston et al., 2017. The systematic testing of the model predictions in experimental conditions is an objective that we intend to pursue in the future. Furthermore, the present model can be involved in neuropsychological studies on patients following the trend of computational phenotyping (Krakauer & Shadmehr, 2007; Schwartenbeck & Friston, 2016), where a “digital version” of the patient is simulated by fitting the patient’s action into the model and recover the parameters of the model that generated them. Specifically, a prior or a likelihood matrix can be damaged to resemble neuropsychological disorders. See the next chapter for developments 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 states spaces as shown in Parr et al., 2021 and Parr & Friston, 2018.

Chapter 5

2nd Architecture: limb apraxia and active inference in the visuomotor pathways

This study provides a computational account of apraxia and the visuomotor system under the active inference framework. In this context, the computational aspects of the brain are explained as a process of Bayesian inference to minimize an approximation of the surprisal associated with sensory observations, called variational free energy. The aim is to show how the neurofunctional organization of the visuomotor system becomes an emergent property of free energy minimization, under the appropriate generative model. The model proposed suits the current neuropsychological literature as it recapitulates cognitive models based on the two visuomotor pathways. By applying virtual lesions to the model, we provide the profiling of five apraxia syndromes that affect different cognitive aspects of motor behaviour. These new hypotheses about the neurocomputational basis of the pathology can offer a more formal quantitative approach in future clinical research.

5.1 Introduction

In this study, we aim to provide a unified perspective on the visuomotor system and its deficits by jointly describing them in neuropsychological, computational and neuroanatomical terms under the framework of active inference, where brain structures are considered to entail a probabilistic generative model that explains the causes that generate sensory observations, and brain functions are cast as a process of active Bayesian inference to minimize an approximation of the improbability of those sensations. By adopting this formalization, we provide a computational account of apraxia with a generative model that suits the neurofunctional evidence of previous cognitive and neuroanatomical models. We start with an overview of limb apraxia. Then we propose a descriptive cognitive model of the visuomotor system that, starting from previous historical proposals, provides an updated perspective that can be coherently mapped with the current neuropsychological and neuroanatomical literature. We then propose a generative model able to account for the tasks executed during clinical assessments coherently. We refer to some key theoretical constructs that are useful to motivate the structure of the generative model and how it recapitulates the previously shown cognitive model. Next, we explain how the artificial lesions are applied to the model and how they are justified when pathology is cast as optimal inference with suboptimal priors. Lastly, we proceed with profiling five syndromes affecting the visuomotor system due to brain damage by providing a simulation of behaviour and a formal illustration of how specific pathological priors may characterize the deficits. We aim to account more formally for the predictions generated by the cognitive model by addressing how behavioural impairment may be produced at the neurocomputational level. In each simulation, the impairment can be observed both at the behavioural level, observing how the simulated patient is able or unable to perform the task correctly, and at the level of probabilistic beliefs. This approach is inspired by the one of computational phenotyping, which originated in psychiatric research (Schwartenbeck & Friston, 2016), where the computations performed by the brain are associated with measurable behaviours (Krakauer & Shadmehr, 2007; Mirza et al., 2016; Testolin & Zorzi, 2016).

5.2 Limb apraxia

Limb apraxia is a deficit of the intentional execution of actions, firstly described by Steinthal (1881) as an impairment in planning, executing actions and moving a body part because of a brain lesion. Crucially, the deficits are caused by neither a lack of comprehension of the task instructions nor a deficit of the sensorimotor system, object recognition problems or frontal inertia (De Renzi et al., 1986; Heilman & Rothi, 1993). Furthermore, the deficit affects bilateral actions as the impairments manifest when either the affected or the unaffected hand is used (Heilman & Rothi, 2003). The assessment of apraxia usually involves imitation, instead of the verbal probing of actions, since the cognitive-motor deficit often coexists with language comprehension deficits (Goldenberg, 2014; Renzi, 1985). In particular, imitation of both meaningful and meaningless actions is required as the latter are often impaired and are not represented in long-term memory (Goldenberg et al., 1996; Mengotti et al., 2013; Rumiati & Tessari, 2002a; Tessari et al., 2007). Meaningless actions carry no meaning as they do not produce an identifiable effect on the environment, either material or communicative, and their interpretation is related to recognizing only kinematic aspects (Goldenberg, 2014; Proietti et al., 2021). On the other hand, in meaningful actions, their meaning can be recognized and reproduced from the stored representation in long-term memory. Moreover, they can be distinguished into transitive actions, involving actual manipulations (such as actual tool use or the performance of complex sequences of actions, or pantomime of tool use) and intransitive actions, which involve symbolic gestures (such as the military salute or waving goodbye).

5.3 A cognitive model of praxis

The first cognitive model of praxis was proposed by Rothi and colleagues (Rothi et al., 1991) based on the models of language production (Patterson & Shewell, 1987) and has been later implemented by many authors (Buxbaum & Randerath, 2018; Cubelli et al., 2000; Rumiati & Tessari, 2002a). The model describes the cognitive elaboration involved in different motor tasks, such as recognizing and producing actions by imitating, pantomiming on verbal command or at object view, and the possible processing failures that can occur at different steps of the elaboration based on input and output processes, type of stimulus and modalities requirements (Rothi et al., 1991). In this study, we propose a further elaborated version of the cognitive model and a mapping with neuroanatomical substrates and computational mechanisms.

Our model entails three routes that map visual inputs into motor outputs. A first direct, sub-lexical, route based on a visuo-motor conversion mechanism parses the observed action in smaller and simpler motor components without involving stored action representations. Such a route is the only one involved in the reproduction of meaningless actions (nevertheless it can also be involved in the reproduction of meaningful actions). A lexical route is specific for meaningful actions that already belong to a person's repertoire and can be further divided into semantic and non-semantic subroutes. The non-semantic route involves an action and tools-related recognition system which drives a motor output (action output system) by reproducing only over-learned, meaningful actions, encoded as action engrams (Rumiati & Tessari, 2002a; Tessari et al., 2021). This route mostly relies on structure-related components of actions and tools (Cubelli et al., 2000; Rothi et al., 1991; Tessari & Rumiati, 2004). On the other hand, the semantic route involves more abstract components that include a vocabulary of actions (action semantics) and higher conceptual representations. It provides a mechanism for motor programs storage that is driven by an explicit gesture identification process and by semantic representations of tools, such as their function (Bartolo et al., 2001; Buxbaum & Randerath, 2018; Cubelli et al., 2000; Mengotti et al., 2013; Rumiati & Tessari, 2002a). All routes converge into a motor working memory buffer that holds temporarily both meaningful or meaningless motor plans for being performed (Buxbaum & Randerath, 2018; Cubelli et al., 2000; Rumiati & Tessari, 2002a) and allows learning new actions (Ottoboni et al., 2021; Tessari et al., 2006). See Figure 5.1.

The routes are richly interactive, and in most instances, they all contribute to visuomotor processing, although their contribution depends on the nature of the task and the environment (Buxbaum & Kalénine, 2010; Tessari & Rumiati, 2004). However, under some circumstances, such as brain damage, the different contributions of the routes to action reproduction can emerge. For example, double dissociations in patients' performance occur as selective deficits in imitating either meaningful or meaningless actions, suggesting relative functional independence between the direct and the lexical routes (Bartolo et al., 2001; Cubelli et al., 2006; Mengotti et al., 2013; Tessari et al., 2007). As we will show, the specific cognitive and behavioural deficit depends on the locus of the impaired processing stage (Bartolo et al., 2001; Heilman et al., 1982; Sirigu et al., 1995).

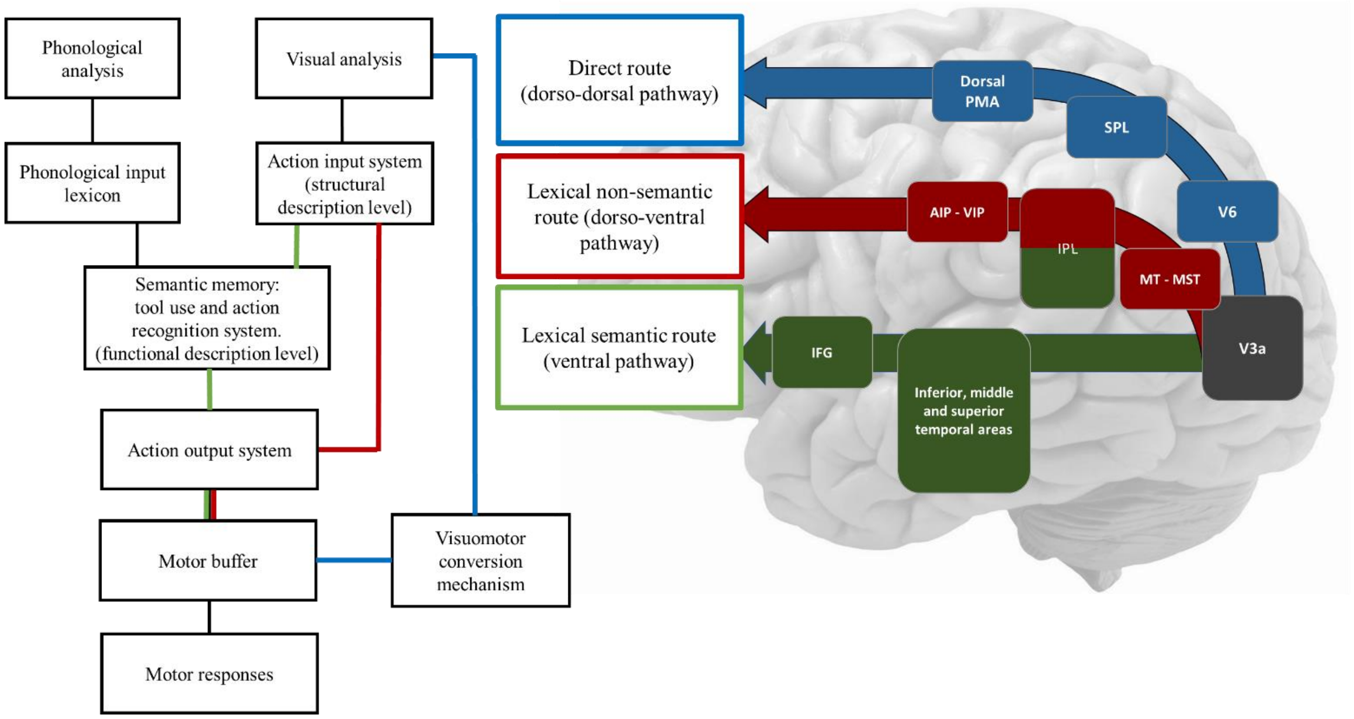

A functional model of the visuomotor system

Figure 5.1: A functional model of visuomotor cognition and its neuroanatomical mapping: from the top of the model on the left, visual (but also phonological) information is taken as input. This information is then processed in parallel by three routes that entail a different level of hierarchical elaboration, from a more sensorimotor representation to a more conceptual one. The different processing stages are linked in the form of three routes: a direct route (in blue), a lexical non-semantic route (in red) and a lexical semantic route (in green). Damages can occur at different processing stages along the routes, producing differentiated behavioural deficits. On the right, these three routes are mapped on the cortical neuroanatomy of the dorso-dorsal pathway (in blue for the direct route), of the dorso-ventral pathway (in red for the lexical non-semantic route) and of the ventral pathway (for the lexical semantic route). This figure is an implemented version from Tessari, Proietti, & Rumiati, (2021). Cognitive Neuropsychology, 38(7-8), 515-530.

5.4 Neuroanatomical mapping

At a neuroanatomical level, the cognitive model can be mapped onto distinct subdivisions of the visuomotor pathways in the brain (Binkofski & Buxbaum, 2013; Rizzolatti & Matelli, 2003) starting from the historical contribution of (Goodale & Milner, 1992; Jeannerod et al., 1995; Mishkin et al., 1983).

A dorso-dorsal stream is specialized for online sensorimotor representations of the postural alignment of different body parts and converts physical object properties, such as location or size, into appropriate motor commands for the control of reaching and grasping. On the other hand, a ventro-dorsal stream contains long-term structural representations about skilled actions, known motor programs, the manipulation of known tools and the spatiotemporal aspects of actions (Binkofski & Buxbaum, 2013; Dressing et al., 2018; Hoeren et al., 2014; Niessen et al., 2014). At last, the ventral stream processes the categorical relationships between perceptual and semantic elements (Dressing et al., 2018; Klein et al., 2013; Lambon Ralph, 2014; Musso et al., 2015; Patterson et al., 2007; Rijntjes et al., 2012; Weiller et al., 2011), including conceptual, functional and semantic aspects of tool-use and gestures and the more symbolic ones (Dressing et al., 2018; Heilman et al., 1982; Niessen et al., 2014; Ottoboni et al., 2021; Rijntjes et al., 1999; Tessari et al., 2007; Vry et al., 2015). This stream has some overlapping with areas involved in communicative tasks (Finkel et al., 2018).

Anatomically, the dorso-dorsal pathway runs from V3a to V6 to V6a and the medial intraparietal area (MIP) in the superior parietal lobule (SPL), and from there, via the superior longitudinal fascicle II, to premotor areas (F2vr and F7-non-SEF1) (Dressing et al., 2018; Hoeren et al., 2014; Martin, Dressing, et al., 2016; Martin et al., 2016; Martin, Nitschke, et al., 2016a; Vry et al., 2015). The ventro-dorsal pathway runs, via arcuate and superior longitudinal fascicle III, from the medial superior temporal area (MT/MST) to the inferior parietal lobule (IPL), and then to the ventral premotor cortex (AIP – F5 and VIP – F4) (Binkofski & Buxbaum, 2013; Kalénine et al., 2010; Kreher et al., 2008; Vingerhoets, 2014; Vry et al., 2012). The ventral pathway runs through the extreme capsule, the inferior fronto-occipital fascicle and the uncinate fasciculus, connecting the inferior, middle and superior temporal area and the IPL to anterior IFG (Catani et al., 2002, 2005; Hamzei et al., 2016; Makris & Pandya, 2009; Rijntjes et al., 2012; Saur et al., 2008; Weiller et al., 2011, 2021). Some authors suggest that also the lateral occipitotemporal cortex (LOTC) should be included in the ventral stream (Lingnau & Downing, 2015).

The direct route and the processing of new movements for meaningless gesture imitation have been associated with the dorso-dorsal stream (Binkofski & Buxbaum, 2013; Hoeren et al., 2014; Martin, Dressing, et al., 2016; Rumiati et al., 2005; Tessari et al., 2007) while the processing of known gestures and the lexical routes have been related with regions belonging to both the ventral and the ventro-dorsal streams (Dressing et al., 2018; Kleineberg et al., 2018; Rijntjes et al., 2012; Weiller et al., 2009, 2011b; Wurm & Caramazza, 2022). Specifically, the ventral stream might decode the meaning of a movement at the conceptual level, with a focus on social and intransitive gestures and non-motor aspects and task-irrelevant objects properties (Bracci et al., 2017; Cubelli et al., 2000; Petreska et al., 2007; Wurm et al., 2017; Wurm & Caramazza, 2022) and the ventro-dorsal stream the structural descriptions of tools and recognition of gestures as meaningful (see also, Dressing, et al., 2018; Tessari et al., 2021).

5.5 The simulated tasks and the generative model

Once again, to characterize the praxis disorders our formalization is based on an active inference model. From a neuropsychological perspective, a functional model of cognitive processing must be extracted from how a given task is executed. Thus, to be able to infer how cognitive processing may be carried out, what must be modelled is the task and not the function. This is crucial because dissociations are found in the aspects of the task and not in cognitive functions. We then assume that the task recapitulates the function. Given this, our profiling of simulated patients is based on how they execute different visuomotor tasks and commands.

The setting for these simulated tasks involves a "clinician" and a "patient". The clinician can perform a gesture that must be observed, verbally recognized, or imitated by the patient. Alternatively, the clinician can verbally command the patient to execute a gesture. More specifically, based on the clinical assessments used by Cubelli and colleagues (Cubelli et al., 2000), the "simulated patient" can be required to execute four kinds of tasks:

- Imitation of familiar actions: the clinician performs a gesture without objects that has to be imitated by the patient. The gesture consists of a tool-use pantomime such as the pantomime of hammering.
- Imitation of novel meaningless actions: the clinician performs a meaningless gesture that has to be imitated by the patient. This task allows for assessing the functioning of the direct pathway.
- Execution of meaningful actions from verbal commands: the clinician verbally pronounces a meaningful gesture that the patient has to reproduce (e.g., "wave goodbye").
- Recognition of meaningful actions: the clinician performs a meaningful gesture, and the patient has to label it verbally.

In practice, each of these tasks is simulated via a generative model in a partially observable Markov decision process. The generative model we propose is an extended and modified version of the one illustrated in Chapter 4. See also Figure 5.2 for a detailed account.

In the tasks, the clinician performs or commands an action that is cognitively processed by the patient at four levels of hierarchical representations as probabilistic beliefs (Grafton & de C. Hamilton, 2007; Wurm & Lingnau, 2015):

1) a kinematics analysis level, which entails the encoding of kinematic features of movements (in our case, represented as elementary motor engrams, meaning body parts and their relations);

2) a body schema representation level, where the relations between body parts are described as a posture;

3) a structural description level, where a form of recognition occurs and the representations are organized in more complex engrams, stored in memory (in the model as sequences of postures);

4) a functional-conceptual description level, where a form of identification occurs as the representations are generalized at a conceptual level and become less constrained by motor aspects.

A more detailed description of the functioning of the model is provided in chapter 4. For this study's purposes, it is important to remind that the model structure comprises three sections: a higher hierarchical inference level (*action understanding level*) that accounts for a semantic interpretation of the observed action, encoding the structural description level and the functional-conceptual description level; a lower level of inference (*action observation level*) for the observation and the recognition of kinematics aspects encoding the kinematics analysis level and the body schema representation level; a motor response structure (*action execution level*) linked to both levels, that reproduce the observed or commanded action. The translation from the kinematic patterns to lexical and semantic interpretations of the clinician's actions occurs between the action observation level and the action understanding level. The translation into motor responses is prompted by both levels. See Figure 5.2.

Although the task is not representative of a practical neuropsychological assessment, the structure of our model is sufficiently complex to highlight the functional rules and constraints provided by active inference and message passing. Therefore, we can provide an operative example of how a disconnection in one or more nodes affects the overall functioning of the visuomotor system, associating the breaking of specific links with different profiles of limb apraxia.

How are the tasks prompted in the generative model? Computationally speaking, in a *meaningless gesture* the mapping between state and outcomes (likelihood function, A matrix) at the structural description and functional-conceptual description levels is highly entropic. This means that the simulated participant only has access to a kinematics analysis and a body schema representation to translate the observed action into movement (Buxbaum et al., 2000). On the other hand, when this mapping is precise, the patient knows the meaning of the gesture (i.e., she is in front of a *meaningful gesture*). In short, the knowledge of the meaning of a gesture is associated with having precise likelihood mapping. Verbal commands are prompted by setting a higher prior belief on the representation of a specific action. Verbal recognition is assumed when the action is identified at the functional-conceptual description level. The imitative responses are generated by taking the posterior belief about the states of the *action understanding level* and the *action observation level* and are used to inform prior preferences and the prior policies at *the action execution level.*

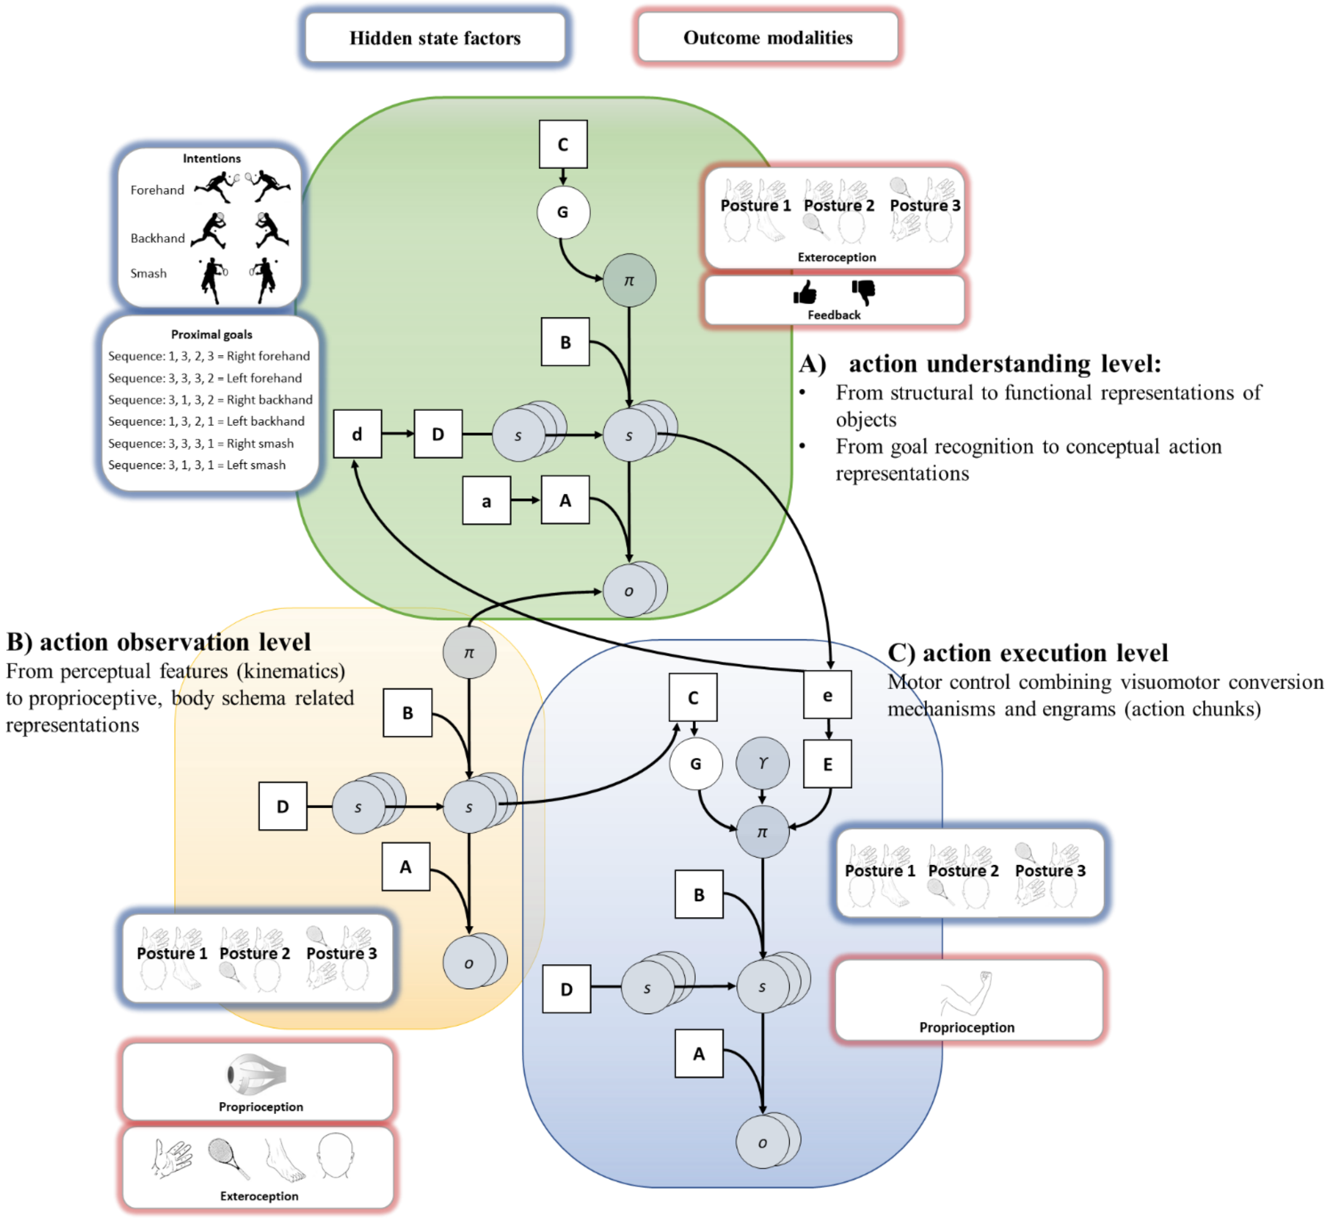

A) action understanding level:

- From structural to functional representations of objects
- From goal recognition to conceptual action representations

B) action observation level

From perceptual features (kinematics) to proprioceptive, body schema related representations

C) action execution level

Motor control combining visuomotor conversion mechanisms and engrams (action chunks)

Figure 5.2: The generative model has been adapted from chapter 4. Here, hidden state factors correspond to different states of the world, such as the location (where) and category (what) of an element. Outcome modalities correspond to the possible observations. Multiple outcome modalities account for parallel sources of sensory input, such as visual inputs and proprioceptive sensations. The generative model produces actions both during observation to gather the most informative piece of data and during the execution of the motor response. The most probable policy is selected using a softmax function of policies expected free energy to do this. The ***B*** *matrixes determine a single piece of actions (as state transitions), while* ***A*** *matrixes constitute recognition as perceptual inference (observation of the outcomes). Initial prior beliefs are represented by* ***D*** *vectors (probability of initial states), while* ***C*** *matrices are a prior that biases the amount of expected free energy* ***G*** *to determine a priori which outcomes are preferred.* ***E*** *constitutes a prior over the policies defining a priori which policies are more probable to be selected.*

5.6 Modules, hierarchies and degeneracy in the generative model

In this section, we highlight three features pertaining to the structure of the generative model that are particularly relevant for the anatomical organization of the brain and that become even more relevant when considering its lesions.

When variational free energy is considered among competing models for explaining certain data, its minimization is achieved by the process of Bayesian model selection, where the model with the most evidence is the one that provides the most accurate explanation of the data and, at the same time, is the less complex (e.g., its structure entails the least number of parameters or degrees of freedom). Therefore, optimizing the structure of a generative model to make it less complex is another way of minimizing free energy and this is associated with the neuroanatomical organization (e.g., structure learning, Friston et al., 2017; Hobson & Friston, 2012; Tononi & Cirelli, 2006). Crucially, this complexity can be reduced when the conditional dependencies are known so that only the marginal distributions need to be encoded (e.g., the internal states of a Markov Blanket only require knowing the blanket states (Friston, 2021; Parr et al., 2020).

The hierarchical organization of the brain represents one of the structural proprieties reducing model complexity (Hasson et al., 2015; Kiebel et al., 2008; T. S. Lee & Mumford, 2003; Murray et al., 2014; Vidaurre et al., 2017). To support this, our generative model is equipped with a deep temporal feature: the *Action understanding level* generates the hypotheses that are tested by engaging saccades at the *Action observation level*. At the same time, the posterior beliefs about postures at the Action observation level are passed as ascending messages, becoming observable outcomes at the Action understanding level (Proietti et al. 2021). This also highlights how the higher the hierarchical level the more actions are represented in an abstract way. Moreover, the temporal dynamics in the different levels runs at different scales (Friston et al., 2017; Kiebel et al., 2008, 2009). Intuitively and in similitude with the language domain, the evolution of the states related to kinematics aspects of an action (as the meaning of a word) entail a faster time scale than the one associated with the intention (the meaning of the whole sentence). However, the different levels of the hierarchy interact bidirectionally as the inference of each latent variable is simultaneously influenced by messages passed from different directions.

Furthermore, to account for complexity in the model structure, the brain also entails functional segregation or neuronal populations, also known as modules (O’Keefe & Recce, 1993; Zeki & Shipp, 1988). According to Parr and colleagues (Parr et al., 2020), this modularity can be mathematically expressed in the form of factorization of a probabilistic system. This can be achieved following the mean-field theory (Fasoli, 2013), where posterior beliefs can be approximated as the product of statistically independent factors (Beal, 2003). These factors are not completely independent but send messages to one another via their mean-fields (their average values). These mean-field dynamics can be interpreted as a description of the message passed among the neural modules (Parr et al., 2019; van de Laar & de Vries, 2019). This help to formalize their functional specialization, such as in the classical account of how the brain can encode ‘what’ and ‘where’ information independently (Friston & Buzsáki, 2016; Parr et al., 2020). This also comes useful because this factorization relates to the notion of transfer learning or context invariance (Bengio, 2011; Parr et al., 2020; Rojas-Carulla et al., 2018). In our context, this relates to the fact that the meaning of a gesture can be inferred independently from the context, e.g. when the gesture is observed from different visual perspectives (in the schemes, when the disposition of body parts is flipped). Although we make no claims about which specific approximation techniques may be implemented in the brain, the notion of beliefs factorization supports the interpretation of the brain as a system that entails sparse dependency structures (Parr et al., 2020) and therefore, multiple visuomotor neural streams.

The third feature in our model relevant to structure complexity is *degeneracy*: this concept in biological systems refers to the ability to use alternative structures for the same task (Edelman & Gally, 2001; Tononi et al., 1994, 1999). This subsumes a many-to-one mapping where different physical structures can yield the same functional outcomes or behaviour. In the brain, this means that a cognitive function can be supported by more than one neural system (Friston & Price, 2003; Noppeney et al., 2004; Price & Friston, 2002; Sajid et al., 2020). This is a crucial topic in neuropsychology where the same behavioural output (e.g., spelling words) can be achieved by two different strategies (phonological vs lexical/semantic). Indeed, it provides a perspective of how different pathways may support functional recovery after brain damage. Acknowledging that neuronal structures entail various degrees of degeneracy and pluripotentiality (Noppeney et al., 2004) is crucial for designing more accurate lesion-deficit models (Sajid et al., 2020). In our model, a degenerate function emerges from the fact that the agent motor control (*Action execution level*) is driven by the posterior belief of 3 factors, the body schema representation, the structural description level and the functional-conceptual description, which are then combined and weighted as a form of hierarchical motor control (Pezzulo et al., 2018b). Our agent's ability to execute actions is then spared even if one of these three motor control inputs is disconnected. Our simulated profiling will provide a more detailed account of degenerate function in visuomotor control.

In summary, through our model, we show how neuropsychological data support the functional relevance of a modular and hierarchical organization showing how patients manifest selective deficits in cognitive functions supported by different neural streams and at different levels of abstraction, from motor to more semantic. We also show how a degenerate function can be maintained (e.g., imitation of a gesture) even if hierarchical structures that support it at different levels of categorization are lesioned (Humphreys and Forde 2001).

5.7 Disconnections and Bayes optimal pathology

The structure of the generative model that the brain implements to make inferences is the key factor that determines anatomical connectivity (Parr & Friston, 2018a). We will show that our scheme appears to be consistent as, when disrupted, it gives rise to the same deficit observed in experimental and clinical domains. But how is this disruption motivated and performed? Our neuropsychological characterization is motivated by the notion of Bayes optimal pathology (Parr et al., 2018), which argues that any neuropsychological disorder entails a set of prior beliefs that would make a patient's behaviour appear Bayes optimal. In other words, a patient's non-adaptive behaviour is not explained by a broken inference process, but it is due to damages over the biological substrata that encode priors and parameters. In this context, one can speak of optimal Bayesian inference with suboptimal priors (Parr et al., 2018). The question we pose, then, is what kind of prior beliefs are set to generate that optimal behaviour and which neural structure may encode those priors.

Lesions in our model are performed in a way that resonates with the notion of disconnection syndrome (Catani & Ffytche, 2005). A disconnection syndrome may arise from the disruption of the likelihood mapping between two neural populations (e.g., the mean fields or factors). This likelihood is encoded in the Dirichlet parameters of A matrices. These parameters encode synaptic connection strengths providing an account of Hebbian neural plasticity. Furthermore, a disconnection may arise when the propagation of posteriors from one factor or hierarchy to another is interrupted, resembling the disruption of axonal connections. In our simulations, we will formally illustrate the pathological priors for each of the different profiles that characterize the neuropsychological syndromes.

This illustration is provided by associating our generative model with the descriptive cognitive model shown in Figure 5.1. In this way, the cognitive model becomes operationalized in that the links and the nodes in the cognitive model are reflected in the message passing among the links and the nodes in the generative model. Therefore, we can map the computation performed in the generative model with the cognitive function described by the cognitive model. In the following sections, we provide the profiling of five clinical syndromes affecting the visuomotor system, originally proposed by Cubelli and colleagues (2000), by simulating brain damages in different nodes (priors and parameters) of the generative-cognitive model. We will show how each profile and lesion corresponds to a specific impairment in cognitive and behavioural functions.

5.8 In silico brain damages of the praxis model

5.8.1 Simulation 1: profiling pantomime agnosia

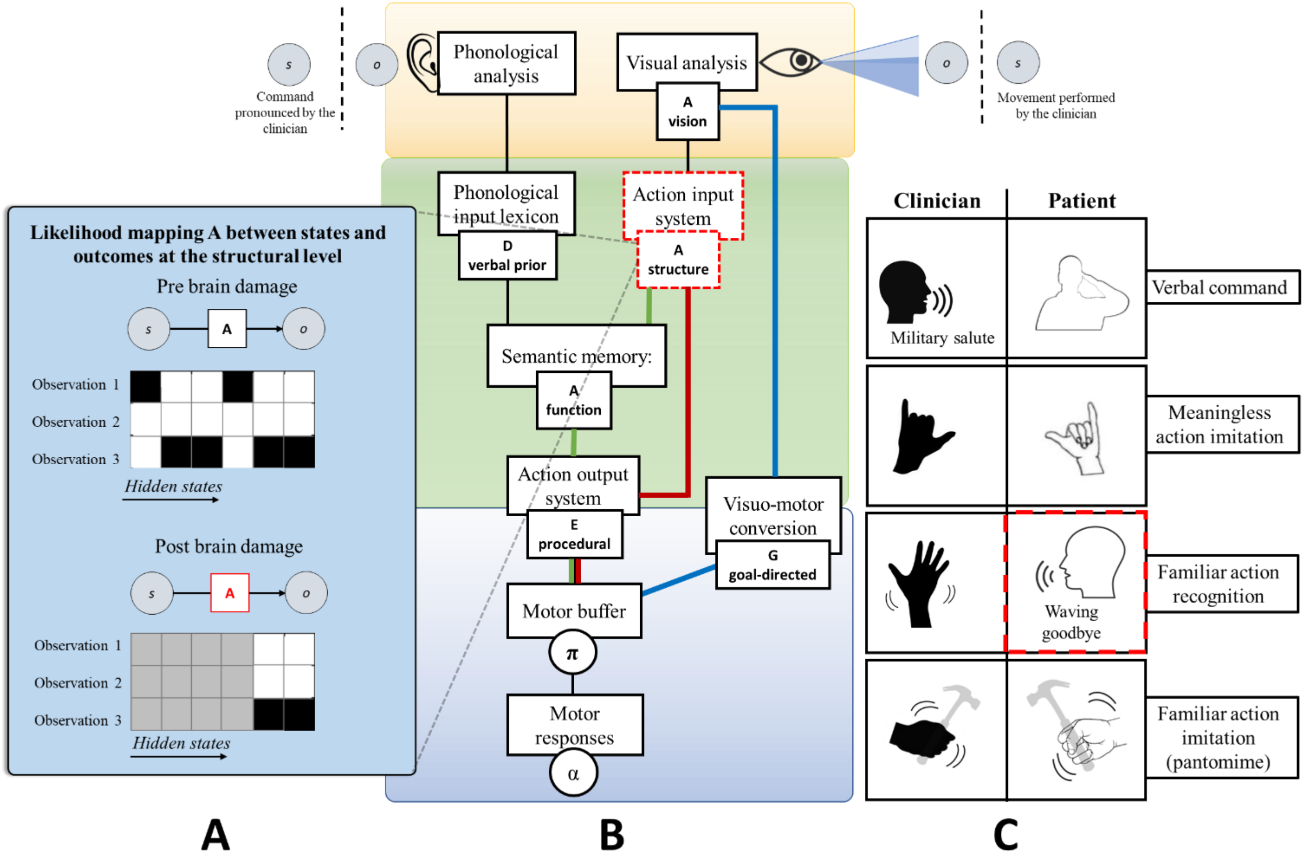

Figure 5.3: to simulate damage over the semantic memory, we lesion the likelihood matrix A that encodes the mapping between hidden states and outcome at the structural level of representation as represented by the red dotted line in the cognitive model in panel B. In the generative model, this corresponds to an entropic distribution in some of the relations between states and observations as shown in panel A where the darker the tint, the stronger the probabilistic relation. At the behavioural level, this translates into impairments in the recognition of familiar actions as indicated by the red dotted line in panel C.

The damage is performed on the action input system by setting flat beliefs (highly entropic belief, e.g., a flat distribution) about the probabilities distribution that encodes the likelihood function (**a** parameter) for the structural description level. This damage can be then seen as a loss of lexical (structural) knowledge (Buxbaum 2001). Thus, the simulated patient manifests the clinical profile of gesture and pantomime agnosia where there is impaired comprehension and discrimination of the observed gestures with a spared ability to imitate or execute them on verbal command (Rothi et al., 1986). The first clinical report of this disorder was provided by Rothi and colleagues (1986). Here, patients' visual agnosia, defined as impaired visual recognition, is not determined by a deficit in visual acuity, attention, or general cognitive ability. The reported patients could imitate pantomimes they could not recognise (Rothi et al., 1986). A pantomime agnosic patient entails an impairment due to a severe loss of visual knowledge about tools structure. The observations on this deficit support the idea that the stored knowledge about tools' structure and function, and gesture identities, such as recognition and understanding, is not necessary for processes subserving object use and imitation

(Buxbaum et al., 1997; Lauro-Grotto et al., 1997; Negri et al., 2007). Identifying an object and recognizing the object-associated pantomimes is dissociated from the ability to use tools and imitate their pantomime. In short, motor production processes related to object use are dissociated from the action or object recognition (Negri et al., 2007; Rumiati et al., 2009; Tessari et al., 2007).

5.8.2 Simulation 2: profiling conceptual apraxia

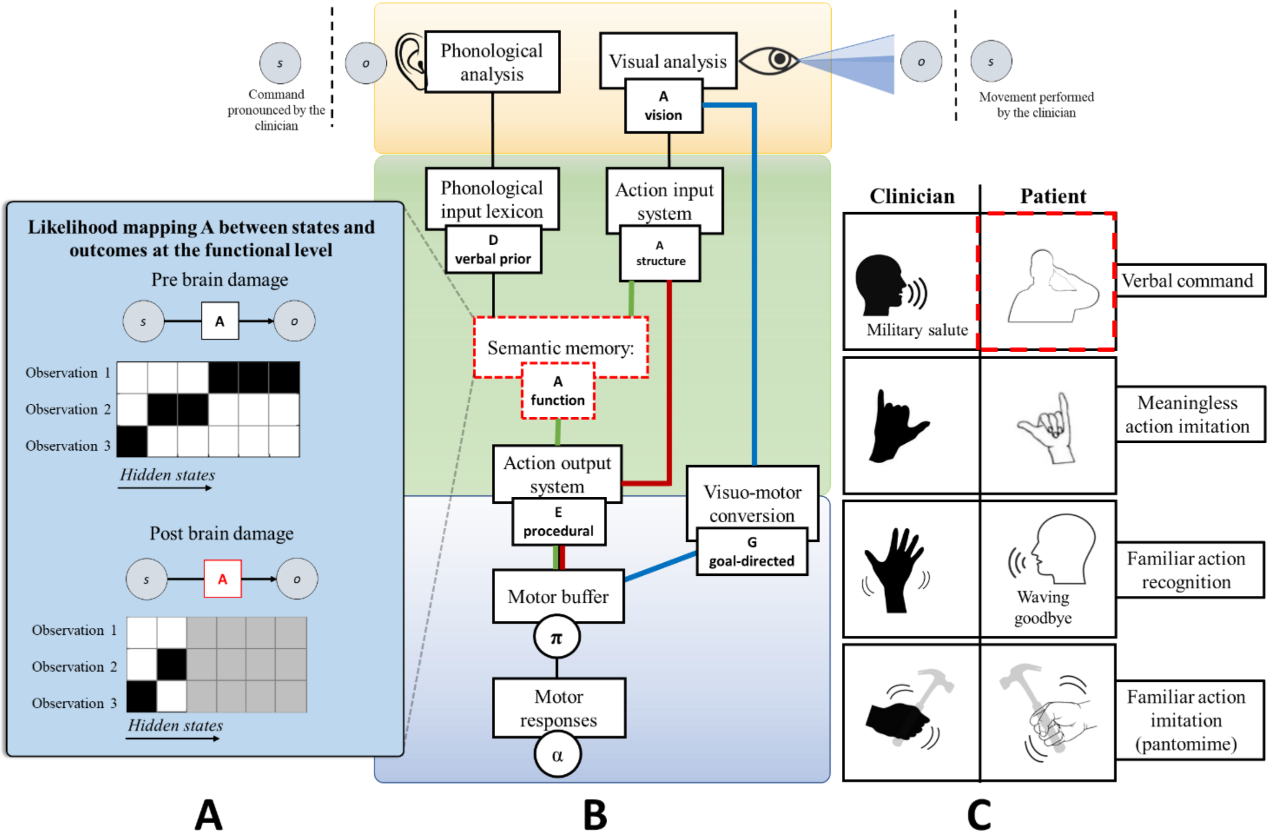

Figure 5.4: to simulate damage over the semantic memory, we lesion the likelihood matrix A that encodes the mapping between hidden states and outcome at the highest level of representation, the conceptual-functional level. This is represented by the red dotted line in the cognitive model in panel B. In the generative model, this corresponds to an entropic distribution in some of the relations between states and observations as shown in panel A where the darker the tint, the stronger the probabilistic relation. At the behavioural level, this translates into impairments in executing actions on verbal commands as indicated by the red dotted line in panel C.

The generative model is damaged in the locus of the highest level of action and tools representation, where they are represented at a semantic and conceptual level. This is done by setting flat beliefs in the **a** parameter for the functional description level. Consequently, it also impairs the system (phonological input lexicon) that biases action engrams via external verbal cues (d parameter on prior beliefs) to be translated into motor programs. In the simulated patient, imitation is spared but there are impairments in the execution of action on verbal command. There are also problems in attributing meaning to gestures via explicit verbal identification; however, visual recognition is intact (Ochipa et al., 1992; Heilman et al., 1995). This occurs because the system is still able to provide a structural description of the input supported by the non-conceptual ventro-dorsal pathway but the multimodal and nonmotor aspects of action knowledge are impaired, which are supported in the posterior IPL and the ventral occipital temporal cortex (VOTC) (Bi et al., 2016; Leshinskaya & Caramazza, 2014)). This recapitulates the ability to disentangle familiar from unfamiliar gestures and to tell apart well-executed from clumsy gestures but not to engage in higher-level identifications. This corresponds to the clinical profile of conceptual apraxia also defined as a semantic type of ideational apraxia (Cubelli 2000) where familiar gestures are recognized as such (structural description is spared), but identification is impaired. These deficits occur in the absence of ideomotor apraxia and semantic language impairments (Bayles et al., 1987): while ideomotor apraxic patients typically make production errors, such as temporal or spatial errors (Heilman and Rothi, 1985), semantic ideational apraxic patients typically make content errors (Ochipa et al., 1989). The loss of knowledge concerns the appropriate actions associated with the tool: the patient shows impairment in selecting the appropriate tool for a particular task and understanding and solving mechanical problems based on intuitive physics. (Ochipa, et al., 1992). This supports the idea that the impairment locus is at a conceptual, functional description level and the neural separability of perceptual and conceptual representations (Leshinskaya & Caramazza, 2016). These patients are able to imitate the known actions, that they do not recognize, properly via the visuo-motor conversion system, i.e. the direct route.

5.8.3 Simulation 3: profiling procedural apraxia

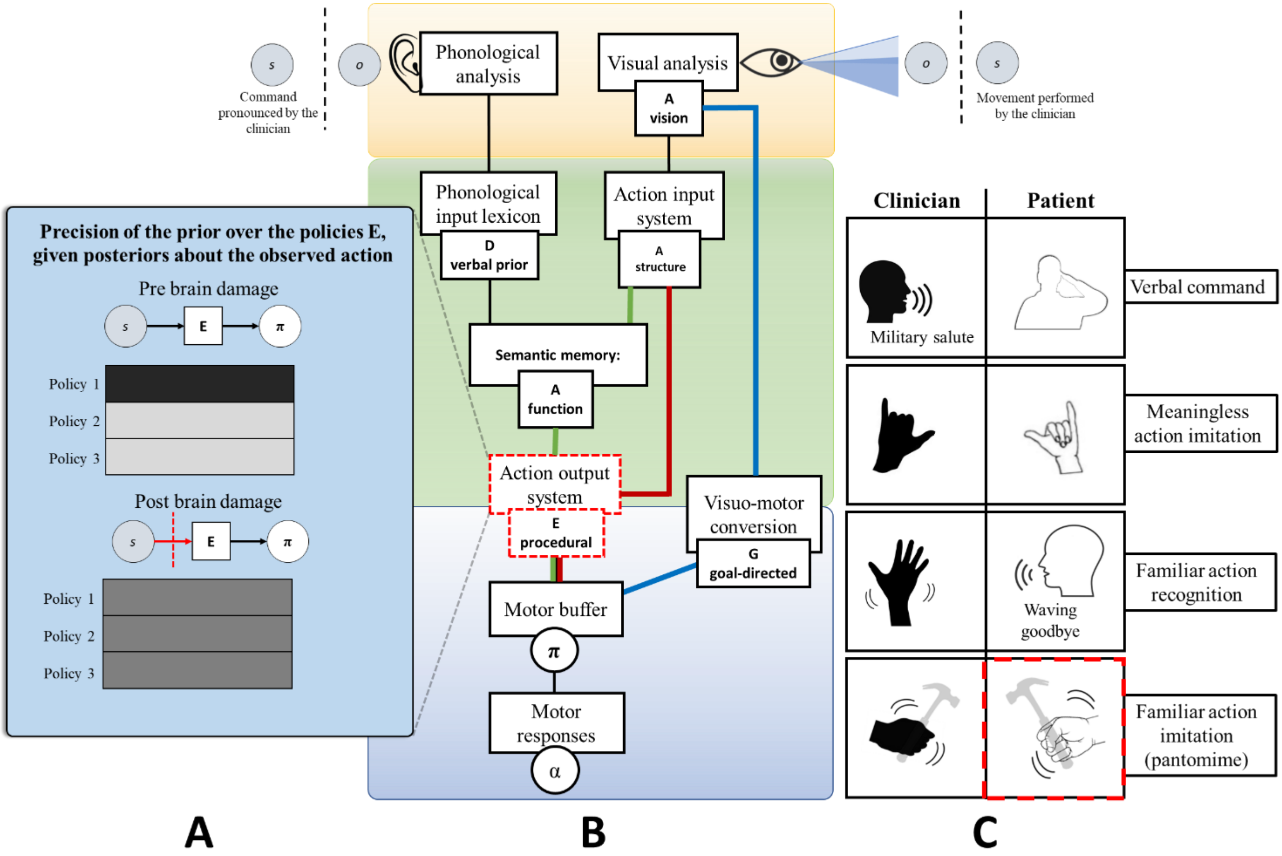

Figure 5.5: to simulate a damage in the action output system, we lesion the passage of posterior belief about the observed action to prior policies E. This is represented in the cognitive model in panel B. In the generative model, this corresponds to an entropic distribution of the prior over the policies E as shown in panel A in the scheme showing the probability of selecting a policy a priori (the darker the tint the stronger the prior). At the behavioural level, this translates into an impairment in imitating familiar actions that require the engaging of a previously learned procedural plan, as indicated by the red dotted line in panel C.

The generative model is damaged in the link that propagates beliefs from lexical (both semantic and non-semantic) representations to the skilled procedural motor controller **E**, representing the action output system that drives motor responses. In other words, the motor program stored in memory disconnects from procedural-skilled motor control that would drive the execution of the action as a whole action program. Therefore, the simulated patient manifests the clinical profile of a procedural type of apraxia where familiar gestures are recognized and comprehended, but their execution is impaired; on the other hand, meaningless gestures are correctly imitated through the direct route (Tessari et al., 2007). The clinical profile is supported by the observation that deficits in movement programming, which include wrong action sequences, inappropriate use of objects and matching tasks, are dissociable from more action semantics errors (Hermsdörfer et al., 2012). The anterior IPL, the posterior middle temporal gyrus, the insula and the extreme capsule are associated with the encoding motor knowledge as action engrams (Buxbaum et al., 2005; Heilman & Rothi, 2003; Hermsdörfer et al., 2013; Kalénine et al., 2010) to detect invariant spatiotemporal features of skilled movements shaped during motor learning (Hoeren et al., 2014; Vry et al., 2012; Weisberg et al., 2007).

5.8.4 Simulation 4: profiling conduction apraxia

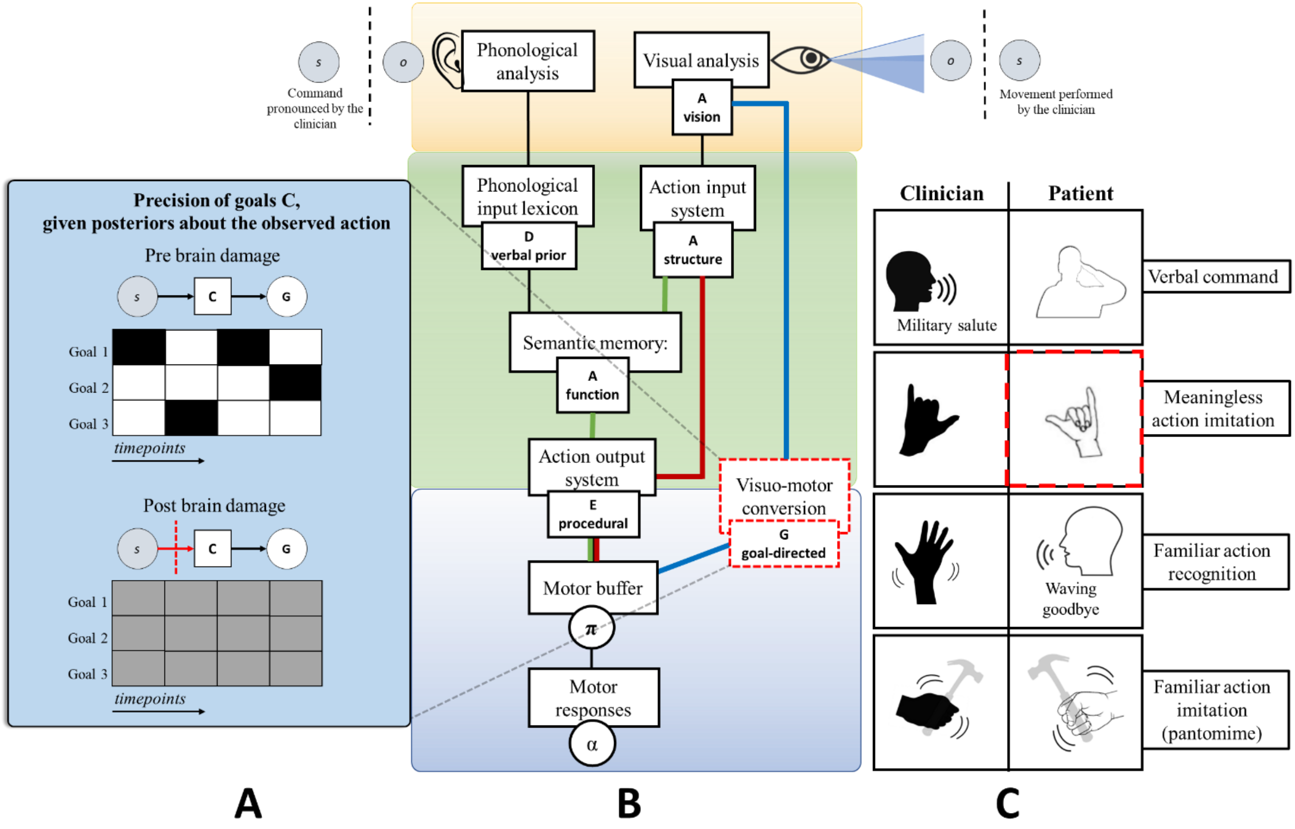

Figure 5.6: to simulate damage in the visuo-motor conversion mechanism, we lesion the passage of posterior belief about the observed action to prior preferences about action execution. This is represented in the cognitive model in panel B. In the generative model, this corresponds to an entropic distribution of preferred observations as shown in panel A in the matrix that encodes the precision of goals (the darker the tint the more desired the observation). At the behavioural level, this translates into impairments in imitating meaningless actions that require online sensorimotor encoding, as indicated by the red dotted line in panel C.

The generative model is damaged in the system that translates low-level, sensorimotor representations about the observed action into the motor plan. This system represents a visuomotor conversion mechanism that sets motor goals for the action execution level, given the posterior belief of the action observation level (Pezzulo et al., 2018, Proietti et al., 2021). Our cognitive model entails a disconnection in the direct route, mapped anatomically into the dorso-dorsal stream. The simulated patient manifests the clinical profile of conduction apraxia where the impairment is specific for meaningless actions while familiar gesture recognition, comprehension and execution are spared (Ochipa et al., 1994). The performance of movement is often clumsy and rigid (Heugten, 1998; Shelton & Knopman, 1991). More specifically, an impaired imitation of meaningless gestures is associated with a dysfunction of the body schema (Buxbaum et al., 2000) which encoding is based on the integration of visual and proprioceptive input about the position of body parts that supports spatial limb representations encoded in the SPL (Hagura et al., 2007; Lacquaniti et al., 1995; Seelke et al., 2012) and the visuospatial transformations involved in mapping observed postures into 'somatesthetic spatial code' performed by the IPS (Buxbaum, 2001; Creem-Regehr et al., 2007; Watson et al., 1986). This is recapitulated in our generative model as the disruption is localized in the *kinematics analysis* and *body schema representation* levels. Numerous clinical cases of patients able to correctly execute tools pantomime and symbolic gestures but unable to imitate meaningless actions have been reported and some of them also reported double dissociations between imitation of meaningful and meaningless actions (Bartolo et al., 2001; Goldenberg & Hagmann, 1997; Ochipa et al., 1994). These studies support that the imitation of meaningless gestures relies on a low-level sensorimotor pathway that rapidly maps perception to action execution, independently from semantic memory. Most impairments have been observed after left brain damage and specifically in the inferior parietal cortex, which is more commonly associated with apraxia (Bartolo & Ham, 2016; De Renzi et al., 1980; Goldenberg, 1995; Goldenberg & Hagmann, 1997; Heilman & Rothi, 1993). Indeed, meaningless hand gesture processing involves predominantly brain areas in the left hemisphere required in action planning. The network includes the temporo-occipital junction (BA 19/37), the inferior and superior parietal cortex (BA 40and BA 5/7), with a critical role of the inferior parietal lobe for mapping perception to action (Dressing et al., 2018; Martin, Dressing, et al., 2016; Martin et al., 2016; Martin, Nitschke, et al., 2016a).

5.8.5 Simulation 5: action working memory apraxia

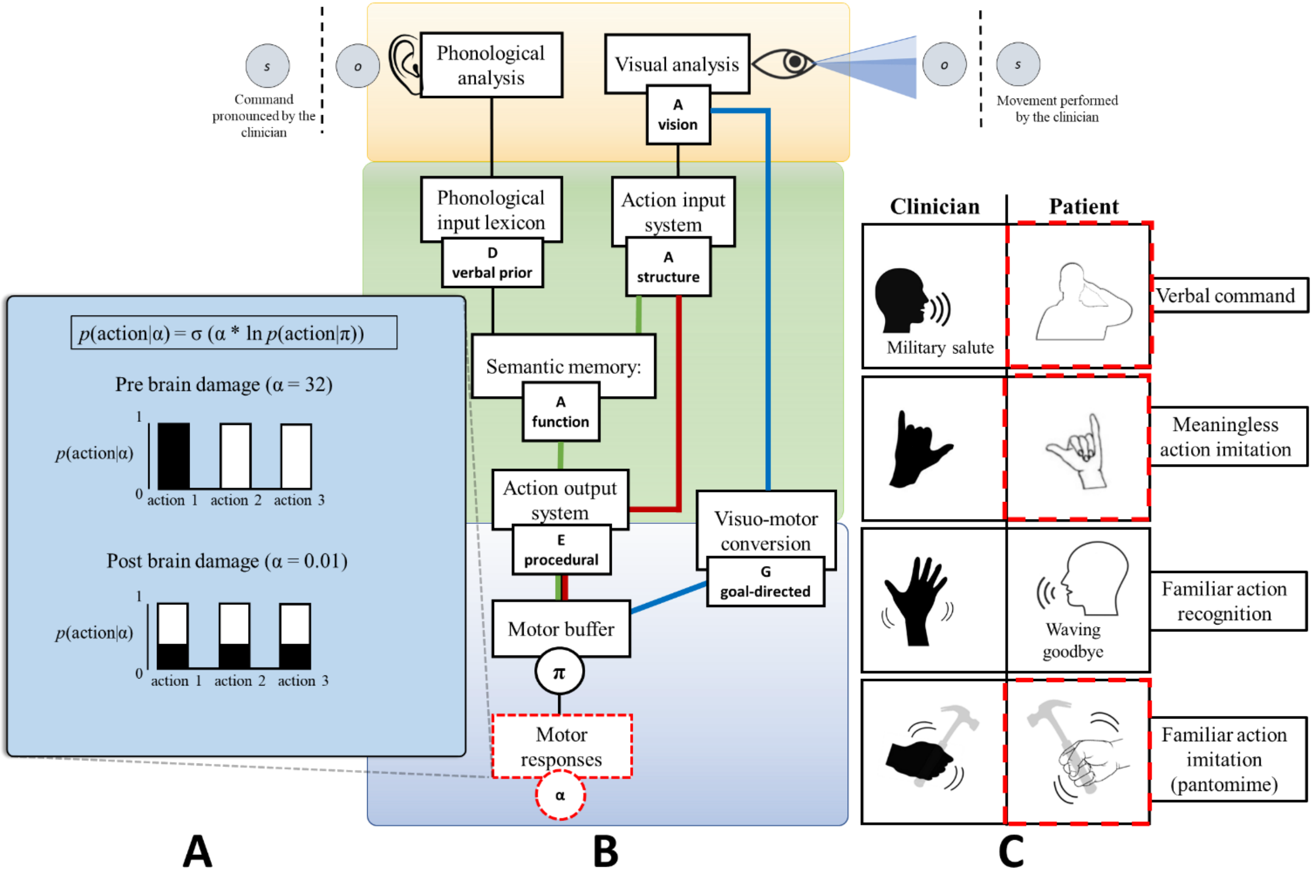

Figure 5.7: to simulate a damage to the motor working memory, we lesion the buffer for motor responses in the cognitive model. In the generative model, this corresponds to the alpha precision parameter for action selection. The red dotted line shows the locus of the lesion (panel B). At the computational level, this means that introducing

a low value of the alpha precision parameter in the softmax function used to calculate the control state given a policy, will render the distribution for the selection of the possible controls highly entropic in comparison to high values that renders it precise (panel A). At the behavioural level, this translates into impairments in every task that requires action execution, as indicated by the red dotted line in panel C.

The generative model is damaged in the precision parameter (the alpha parameter) for the system that translates the policy in a sequence of control states, resulting in a high amount of randomness in selecting these controls. This parameter for action selection can indeed be interpreted as a memory buffer for actions that holds the motor plan ready to be implemented. The higher its value, the more precisely the planned actions are implemented. On the other hand, very low values can be seen as an inability to hold the current plan and translate it into actions. Thus, when this mechanism is damaged, all motor responses are impaired. Indeed, the simulated patient manifests impairments in executing gestures either on verbal command or in imitation tasks (both meaningful and meaningless actions). The working memory subsystem is dedicated to processing motor information not associated with verbal or visuospatial working memory and is specific for encoding, retaining, and recalling goal-directed actions and body configurations (Galvez-Pol et al., 2020). The existence of a working memory component for actions has been supported in several behavioural studies (Rumiati & Tessari, 2002; Smyth et al., 1988; Smyth & Pendleton, 1990; Wood, 2007) where a decreased memory span has been observed on a primary motor task that is being executed concomitantly to a secondary sensorimotor task in comparison to a verbal or spatial secondary task (where this decrease has not been observed). A recent study has reported neuropsychological evidence for a distinct motor working memory subsystem: left hemisphere stroke patients with apraxia showed a more pronounced deficit in motor working memory tasks (involving pictures of object-oriented, meaningless and meaningful actions) in comparison to left hemisphere patients without apraxia and controls (Bardakan et al., 2022). At the neuroanatomical level, the maintenance of this motor information appears to be associated with the frontoparietal praxis network. Specifically, increased functional connectivity and activation in the frontoparietal network (Cai et al., 2018) has been found during a greater memory load of motor information. Furthermore, the left middle and right inferior frontal gyri and the superior and inferior parietal lobule of the left hemisphere are involved in the maintenance of biological motion (Lu et al., 2016). In the context of tool use, activations in the left inferior frontal gyrus and the left ventral premotor cortex have been observed during the maintenance of manipulable tools compared to non-manipulable ones (Mecklinger et al., 2002). In our model, the lower the precision parameter (the alpha parameter), the less precisely the planned actions are implemented, increasing the apraxic deficit as in the neuropsychological evidence provided by Bardakan and colleagues (2022), who showed that the overall performance in the motor WM tasks predicted the severity of apraxic deficits in clinical tests of gesture imitation.

5.9 Discussion

Literature on apraxia may result complex and confusing. As suggested by Petreska and colleagues (Petreska et al., 2007) this might be due to a lack of systematicity in the apraxia assessments that leads to contradictory and inconsistent results, ambiguous use of terminologies and high dimensionality in models' parameters that blurs statistical approaches. With this paper, we aim to provide a first principled approach to unify perspectives under the framework of active inference. Indeed, when working under a normative principle, the key to understanding brain function is specifying the appropriate generative model (Parr et al., 2020; Parr & Friston, 2018a). We have shown as the cognitive model is recapitulated by the structure of the generative model and show and lesion applied in one node of the former corresponds to computational mechanisms of the latter, such as a sub-optimal prior or a change in a parameter value. In our simulations we have provided five apraxia profiles: pantomime agnosia entailing a loss of structural knowledge of actions and tools, conceptual apraxia where the loss of knowledge pertains to the highest functional representation, procedural apraxia as the inability to translate a perceived movement into a pre-existent motor program, conduction apraxia where the failed translation into movements pertains lower level sensorimotor capabilities and, at last, action working memory apraxia where the deficit is due a decrease memory span in motor tasks. These profiles emerge automatically by lesioning the model at different processing stages and as manifested by the consequent behavioural deficits. Therefore, the scheme we propose appears to be consistent as, when disrupted, it gives rise to the same deficit observed in experimental and clinical domains. Furthermore, compared to previous proposals, our model argues for the existence of a third route inspired by evidence from the context of language (e.g., Coltheart et al., 1983; Coslett, 1991; Law et al., 2005; Wu et al., 2002) but also proposed by neuropsychological literature (Dressing et al., 2018; Martin et al., 2016; Martin, Nitschke, et al., 2016a). This may help explain contradictions in the apraxia literature since lesions are often extensive, and the semantic and non-semantic subdivisions of the lexical route may be ambiguously affected.

Although our profiling is purely descriptive (as it does not involve direct measurements), we aim to set a direction for future research. Setting generative models that recapitulate actual clinical assessments designed to fit better model parameters may lead the way for a formal and systematic approach. This would allow for quantitative computational phenotyping (Schwartenbeck & Friston, 2016) and the assessment of statistical proprieties such as degeneracy as shown in (Sajid et al., 2020).

Chapter 6

3^{rd} Architecture: Cognitive control in active inference

We advance a novel formulation of cognitive control and the deployment of optimal attention for action selection, within the active inference framework. The theory proposes that cognitive control amounts to optimising a (precision) control signal, which balances deliberative and habitual components of action selection. This underwrites goal-seeking behaviour while minimizing cognitive effort. To illustrate the theory, we simulate a scenario in which a driver aims to reach her destination safely and with optimal cognitive control in different conditions. When the environment is predictable, she can get to her destination using habitual strategies that engage minimal attentional resources. However, when the driver faces unexpected and potentially dangerous circumstances, she deploys cognitive control to engage in deliberative policies anticipating bad outcomes. Our simulations highlight the information processing dynamics of an active inference agent that deploys cognitive control in conditions of high or low environmental predictability. These numerical experiments show that selecting deliberative policies engages several “cognitive control” processes to increase monitoring resources and learning rates, devaluate habits and resolve cognitive conflicts. However, despite their apparent heterogeneity, all these processes stem coherently from free energy minimization, which is the core claim of active inference. Finally, we discuss the putative neurobiology of cognitive control. We argue that the neural activity in the dACC reflects the specification of an optimal (precision) control signal that balances deliberate and habitual components of action. In contrast, neural activity in the dopaminergic system reflects the further optimization of the precision control signal in relation to the current observations.

6.1 Introduction

It is common wisdom that "practice makes perfect". From a cognitive perspective, practice also entails a gradual passage from an effortful or controlled mode – during the performance of novel tasks – to more automatic information processing during the performance of familiar tasks (Anderson, 1982; Shiffrin & Schneider, 1977). For example, while a novice driver must devote significant cognitive resources to each aspect of the driving task, an experienced driver can often drive effortlessly, paying little “attention to action”.

A common explanation of this phenomenon is that the control of skilled and novel actions is associated with two fundamentally different types of brain processes or controllers, which are called automatic versus controlled (or intentional), habitual versus deliberate, or procedural versus goal-directed processes, or system 1 versus system 2 (Balleine & Dickinson, 1998; Daw et al., 2005; Kahneman, 2011; Norman & Shallice, 1986; Stanovich & West, 2000). The same dichotomy between separate action selection mechanisms, or controllers, recurs across lower-level motor control studies and higher-level decision-making studies (Milli et al., 2021). Hence, to generalize, below, we use the word *actions* to refer to the products of these two controllers, irrespective of whether they are movements, decisions or mental actions.

In habitual or procedural control, the task is initiated and executed without deliberate attention and is performed automatically, without tapping into limited processing resources and without awareness (Kahneman et al., 1983; Posner, 1978; Shiffrin & Schneider, 1977). This type of control is engaged for relatively simple or well-learned cognitive and motor tasks. Two hallmarks of habitual-procedural control are the fact that environmental stimuli can directly trigger action initiation and that action execution can directly engage a preconfigured behavioural plan, such as a sequence of motor acts (Anderson, 1982; Taatgen & Lee, 2003) or action chunks (Dezfouli & Balleine, 2012; Rumiati & Tessari, 2002b; Tessari et al., 2006, 2021). The combination of these two factors ensures that tasks are generally executed faster (because actions are recalled automatically) and with little expenditure of cognitive resources (because the agent needs to monitor only the final outcome of the preconfigured plan, not each constituent motor act). This corresponds to the usual definition of a habit as a skilled action, which can be engaged with minimal processing resources (K. J. Miller et al., 2018). Of course, these advantages come at the expense of flexibility: this type of control is only appropriate when the situation is predictable but can fail in novel or unforeseen circumstances (Moors & De Houwer, 2006; Schneider & Chein, 2003).

On the other hand, deliberate or goal-directed control involves the formation of novel action plans, the careful online monitoring of their outcomes and the capability to counteract maladaptive habitual responses and temptations. In turn, these are all cognitively demanding tasks, which is why deliberative control is associated with mental effort and the engagement of attentional resources (and perhaps conscious processing). At the same time, deliberative control is flexible and permits dealing more effectively with complex and unforeseen circumstances, novel contingencies and volatilities (Balleine & Dickinson, 1998).

Since habitual-procedural and controlled-deliberate processes have complementary strengths and weaknesses, the crucial question is how they are selected and/or combined during action selection. Some proposals assume a “competition” or “arbitration” between two separate behavioural controllers based on their relative uncertainties (Daw et al., 2005). However, more recent proposals assume that the different controllers can cooperate and be combined, therefore providing a continuum of solutions (S. W. Lee et al., 2014; Pezzulo, Rigoli, et al., 2013) and perhaps might be arranged hierarchically (Dezfouli et al., 2014; Pezzulo et al., 2015). Furthermore, there is growing appreciation of the fact that the arbitration or combination of habitual-procedural and controlled-deliberate processes is a cost-benefit computation that should consider not just the pragmatic benefits but also the cognitive effort required to engage those processes (Daw et al., 2011; Dolan & Dayan, 2013; Maisto et al., 2019; Pezzulo, Rigoli, et al., 2013). This leads us to formal theories of cognitive control, which formalize the decision of how much cognitive effort to deploy to engage in more costly (controlled-deliberate) or less costly (habitual-procedural) processes.

A model of attention that provided an early attempt to characterize this cost-benefit computation is called *attention to action* by Norman and Shallice (1986). Accordingly, two processes operate complementarily to select and control actions (although their joint operation could be disrupted under certain maladaptive conditions). First, a *contention scheduling* selects amongst the possible action schemas – which in modern terms could be called a policy selection mechanism – that selects among policies or action sequences (Friston, et al., 2017; McClelland & Rumelhart, 1981; Parr et al., 2022; Rumelhart & Norman, 1982; Sutton & Barto, 1998). Second, a s*upervisory attentional system* controls the selection of action schemas by exerting some additional activation (or inhibition) to an action schema to bias its selection in the contention scheduling mechanisms. This second mechanism provides some cognitive control by deploying attention, which requires cognitive effort, especially if the schema to be selected is unfamiliar (Cooper & Shallice, 2000; Shallice & Burgess, 1993).

A more recent, neurobiologically grounded theory of cognitive control and the deployment of cognitive effort is the *expected value of control* (EVC) by Shenhav and colleagues (2013). This theory proposes that the allocation of control is based on a cost-benefit evaluation of the payoff one obtains by engaging a controlled process and the cost (cognitive effort) required to engage the sufficient amount of control to achieve the payoff. The theory identifies three key processes of cognitive control. First, the *regulation process* describes the capacity of a control mechanism to influence lower-level information processing mechanisms. Regulation is achieved by a control signal that changes the parameters and functioning of lower-level mechanisms, and which has two fundamental features: identity and intensity. Identity specifies which lower-level parameters are targeted or which behaviour is up-weighted and which one is inhibited. Conversely, intensity represents the strength of the signal, such as the degree to which the lower-level parameters are displaced from their default value. Second, the *specification process* is responsible for the decision of whether or not to pursue a controlled process and (in case) the actual selection of the most appropriate control signal, which specifies which of the possible action plans should be engaged and how intensely (e.g., accurately) they should be pursued. Third, the *monitoring process* ensures that the cognitive system has essential information for signal specification, which includes information about current circumstances and whether the current behaviour is affording progress towards goals. An ample literature suggests that monitoring processes could consider various sources of information, such as response conflict, response delays, errors, and negative feedbacks, which might indicate a need for cognitive control (Botvinick et al., 2001; Laming, 1968; Rabbitt, 1966; Shenhav et al., 2013). Finally, at the neural level, EVC theory ascribes to the dorsal anterior cingulate cortex (dACC) roles in monitoring and specification, and the lateral prefrontal cortex (lPFC) a role in regulation.

Here, we advance a formal account of cognitive control and its associated cost-benefit computations, and cognitive and neuronal processes within the active inference framework. We formalize cognitive control as the optimization of a single control signal: a (*precision*) parameter that balances habitual and deliberate components of action selection in (hierarchical) active inference (Pezzulo et al., 2015, 2018). One can also describe the optimization of the precision parameter as a form of *meta-control* because the precision parameter controls (or weights) the contributions of habitual and deliberate controllers (Botvinick et al., 2019; Doya, 2002; Silvetti et al., 2018). We will show that optimizing the precision parameter provides an optimal solution to the key problem of cognitive control: namely, ensuring accurate action selection at the lowest computational cost.

Our account conceptually relates to previous proposals that cast attention to action (Norman & Shallice, 1986) and cognitive control (Shenhav et al., 2013) as optimization problems. However, unlike previous proposals, we derive control equations from the normative scheme of active inference. Besides its normative appeal, our proposal reconciles two separate streams of research that focus on aspects of cognitive control that depend on reward-related factors (Shenhav et al., 2013) and epistemic factors, such as environmental uncertainty and ambiguity (Behrens et al., 2007). Active inference uses a free energy minimization scheme that considers both pragmatic (goal or reward achievement) and epistemic (uncertainty minimization) imperatives during action selection, hence explaining the two facets of cognitive control. Furthermore, casting free energy minimization in terms of gradient descent allows one to simulate neuronal dynamics in the dACC, the locus coeruleus and the dopaminergic system.

In the following paragraphs, we detail our formal implementation of cognitive control within active inference and the putative neuronal correlates of control computations. We subsequently describe three simulations that exemplify active inference in driving tasks with distinct (lower to higher) control demands. Finally, we summarize and further discuss how our formulation explains the rich phenomenology and neurobiology of cognitive control.

6.2 Computational model of cognitive control during driving

6.2.1 Simulation scenario: the driving task

In this numerical study, we simulate a driver who aims to travel safely from home to her office and allocates an optimal level of cognitive control to complete the task successfully without wasting cognitive resources. The driving scenario comprises 32 trials or “road-steps”; see Figure 6.1. At each road-step, the agent receives a sensory cue from the environment and selects one of two possible policies: she can either *drive forward* along the road or *dodge to the left*. During our simulations, most of the time, there are no danger cues; hence, the driver can safely follow the default policy to drive forward, which becomes increasingly habitual. However, sometimes a danger (i.e., a deer on the road) appears and —to avoid a collision — the driver needs to select the alternative, dodging policy. In turn, this requires engaging cognitive control to inhibit the habitual policy (i.e., *drive forward*) and select the alternative policy (i.e., *dodge to the left*) [2].

6.2.2 Generative model for the driving task

The agent solves this task using the generative model shown in Figure 6.1. As in the previous chapters, the generative model is illustrated with the formalism of Partially Observable Markov Decision Processes (POMDP), with nodes s1, s2 and s3 denoting hidden states (i.e., beliefs or discrete probability distributions over unobserved task variables, such as the agent's position and the presence of dangers); nodes o1, o2 and o3 denote observations (observable stimuli, from which the agent can infer hidden states); the node π denotes beliefs about the policies (or action sequences); and edges represent probabilistic relations among state variables (the letters **A**, **B**, **C**, **D**, **E** within the squares that mark the edges specify the probabilistic mappings among variables).

As explained already, when simulating active inference, variational free energy (F) is computed for each allowable policy or action sequence (π) and comprises two terms; see Equation 1. The first is a complexity term that scores a (Kullback Leibler or KL) divergence between the posterior beliefs about states of an auxiliary distribution (called a *variational density* q(s|π) in variational inference), and the posterior beliefs about states of the generative model (called *posterior density* p(s|π)). The second term is accuracy, which scores the expected (logarithm of) the probability of observations given beliefs about unobservable states (ln p(o|s)). These two terms jointly ensure that the agent engages in a continuous perception-action cycle, by updating its (posterior) beliefs about the states to better fit its observations and by selecting courses of action that actively sample the observations predicted by the model. This means that perception and action fulfil the same (free energy minimization) imperative.

$$\underbrace{\mathrm{F}(\pi)}_{Variational\ FE} = \underbrace{\mathrm{D}_{KL}\,[q(s|\pi)\ ||\ p(s|\pi)]}_{Complexity} - \underbrace{\mathrm{E}_{q(s|\pi)}[\ln p(o|s)]}_{Accuracy} \qquad (1)$$

---

[2] It could be argued that dodging to the left is an automatic response for obstacle avoidance, not a deliberate plan. However, in this simple setting, we omit automatic responses from the active inference model and treat the dodging behaviour as a deliberate policy.

Again, as already explained, the *expected free energy* associated with each policy π considers the prior preferences of the agent (i.e., extrinsic or pragmatic value), the expected information gain about states of the world (intrinsic or epistemic value) and the expected information gain associated with the parameters of the model (novelty). The first two terms can be rearranged into risk and ambiguity. Risk is the (KL) divergence between anticipated outcomes given a policy (q(o|π)) and preferred outcomes (p(o)). Ambiguity is the expected uncertainty (i.e., conditional entropy H) about outcomes, given the model's likelihood p(o|s). However, since in this model we implement learning, a third component is added, the *novelty* term that scores how much the beliefs in likelihood parameters (a mapping between states and outcomes) are expected to change after a new observation:

$$\underbrace{\mathrm{G}(\pi)}_{Expected\ FE} = \underbrace{\mathrm{D}_{KL}\ [q(o|\pi)\ ||\ p(o)]}_{Risk} + \underbrace{\mathrm{E}_{q(s|\pi)}[\mathrm{H}[p(o|s)]]}_{Ambiguity} - \underbrace{\mathrm{E}_{p(o|s)q(s|\pi)}[\mathrm{D}_{KL}[q(\mathbf{A}|o,s)\ ||\ q(\mathbf{A})]]}_{Novelty} \qquad (2)$$

For simplicity, we assume that the agent's generative model (shown in Figure 6.1) represents faithfully the task variables and the statistical relations between them, although some aspects of the world need to be learned (specifically, the **A** matrix). The **A** matrix represents the (likelihood) mapping between states and outcomes. The **B** matrix (transition function) encodes the probability of moving from one state to another. The **C** matrix encodes prior beliefs about observations, which in active inference reflect prior preferences. The **D** vector encodes the prior about the initial hidden state. The **E** vector encodes priors about policies. The **G** denotes expected free energy. The lowercase letters **a** and **e** denote the (Dirichlet) concentration parameters of the **A** and the **E** matrices, respectively, and which are used to model learning by accumulating probabilities. The $\boldsymbol{\gamma}$ is a precision parameter associated with expected free energy; it plays a particularly important role in this setting, as it represents the control signal optimized through **G(γ)**. Finally, $\boldsymbol{\beta_0}$ is a prior over the precision, used to simulate dopaminergic responses.

As shown in Figure 6.1, the generative model comprises four sets of hidden states. These represent a warning signal (warning signal present or absent), the driver's current behaviour (drive forward or dodge to the left), a danger (deer on the road present or absent), and the safety of the situation (safe travel if there is no car crash or injured deer if the car has crashed into the animal). For each set of hidden states, there is a corresponding set of observations.

Finally, in this task, we consider the choice between two very simple policies $\boldsymbol{\pi}$ having length 1 (i.e., that comprise only one action): namely, *drive forward* and *dodge to the left*). This reflects the fact that for simplicity, the active inference agent only makes one decision per trial or "road-step" (i.e., they do not plan across multiple timesteps). For each trial, the agent starts from a (dummy) initial state, receives a sensory observation, selects a policy and makes a transition to another state. Then, a new trial begins.

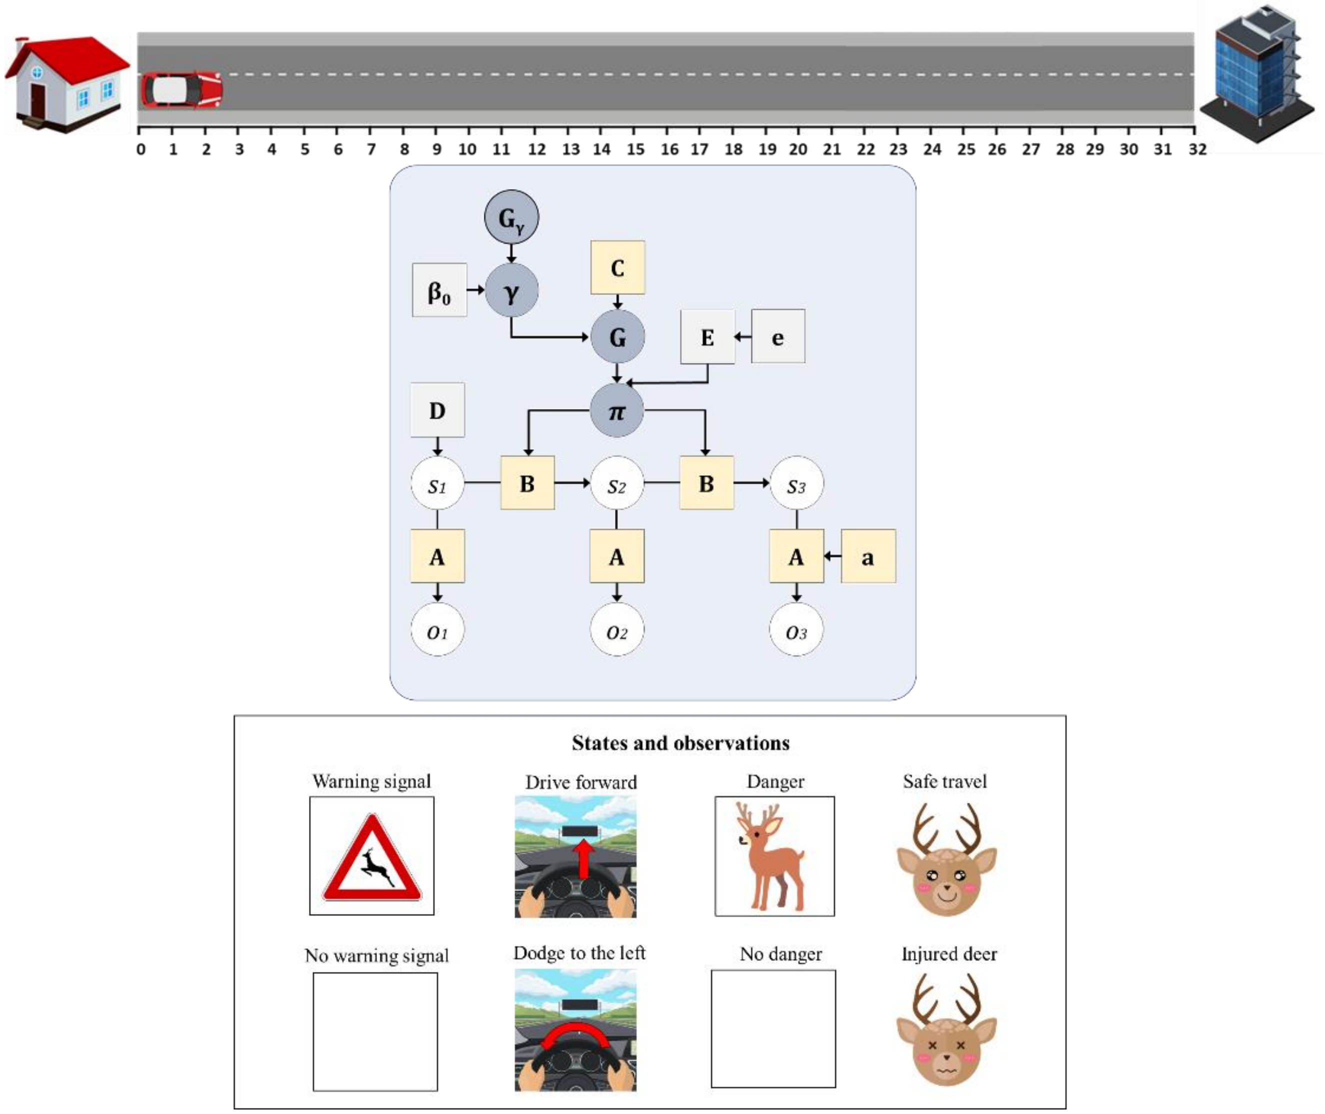

Figure 6.1: The driving scenario. The top panel represents the scenario, with 32 trials or "road steps". The centre panel shows the generative model of the active inference agent, expressed using the formalism of Partially Observable Markov Decision Processes (POMDP). The nodes represent the agent's beliefs about states and task variables, which are encoded as discrete probability distributions. The edges represent the statistical relations between these variables, which are specified in the **A**, **B**, **C**, **D** *and* **E** *matrices. The lowercase letters* **a** *and* **e** *denote the (Dirichlet) concentration parameters of the* **A** *and the* **E** *matrices, respectively, and which are used to model learning by accumulating [co]occurrences. Crucially, the modulation of cognitive control corresponds to the specification of a precision parameter* γ*, which is continuously updated at each road-step of the task and balances habitual and deliberate components of action selection. The bottom panel illustrates the four sets of hidden states of the model and the corresponding observations. See the main text for an explanation.*

6.2.3 Implementing cognitive control in active inference

Having described the generative model of the active inference agent, we focus on the part of the model dedicated to cognitive control. We illustrate how active inference characterizes the cognitive control processes (namely, *regulation*, *specification* and *monitoring*) identified in previous studies (Shenhav et al., 2013) and how it explains their associated neural computations in the dACC and dopaminergic systems.

Regulation of control. The key circuit for the regulation of control is shown in Figure 6.2 (note that this figure simply "zooms in" on the parts of Figure 6.1 that are relevant for cognitive control). Regulation here refers to the capacity of a control mechanism to influence lower-level information processing, i.e., a meta-control. Figure 6.2 shows that the driver's behaviour is determined by balancing deliberate (**G**) and habitual (**E**) components of action. In the figure, the deliberate component of action (**G**) corresponds to the online scoring of the quality of the different policies, i.e., their expected free energy. In contrast, the habitual component of action (**E**) corresponds to a prior over the policies. This prior is learned over time by accumulating the statistics of policy occurrences (i.e., the more a policy is selected, the greater its prior becomes). Note that this accumulation is a straightforward consequence of minimising variational free energy with respect to the underlying Dirichlet parameters **e**.

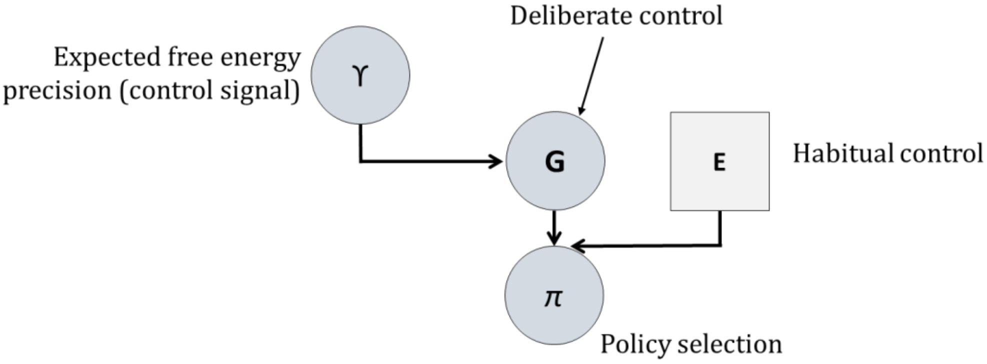

Figure 6.2: Regulation of cognitive control in active inference. This figure zooms in on the components of the architecture in Figure 6.*1 that regulate cognitive control. See the main text for an explanation.*

Crucially, the weighting of deliberate and habitual components of action depends on the **γ** parameter, which constitutes a *precision* estimate of beliefs about expected free energy (i.e., the precision of **G** or confidence placed in deliberative planning). This precision parameter plays the same role as *attentional resources* (Cooper & Shallice, 2000; Shallice & Burgess, 1993) and the *control signal* (Shenhav et al., 2013) in cognitive theories of cognitive control: it determines the weighting of habitual (**E**) and deliberate (**G**) components of action selection. In general, the higher the precision γ of **G**, the greater the weight of deliberate control over habitual control. However, the stronger the habitual component (**E**), the higher the γ parameter needs to be to induce some deliberative control over action selection.

In our driving example, the driver's habitual behaviour (**E**) corresponds to the *driving forward* policy. This habit will become increasingly precise with time if no perils are encountered. Selecting the alternative (*dodge to the left*) policy — in the presence of perils — requires cognitive control by increasing the precision γ of the deliberate component (**G**). Notably, the final policy selection results from a weighted combination of habitual and deliberate components. The relative contributions of these two components are transformed into a (prior) probability distribution over policies through a normalized exponential (softmax) function, as illustrated in Equation 3:

$$\pi = \sigma(\ln \mathrm{E} - \gamma \mathrm{G}) \tag{3}$$

Specification of a control signal. As discussed above, the precision parameter γ plays the role of a control signal that balances habitual and deliberate components of action. Determining the optimal value of γ requires balancing the benefits of deploying cognitive control (e.g., the possibility of obtaining preferred outcomes) versus the required cognitive cost, which we assume to be proportional to γ.

To specialise the generative model for cognitive control – and to render it able to specify a control signal – we introduce a prior over the precision of beliefs about policies. This prior G(γ) can be regarded as an expected free energy, not of policies, but of the precision of beliefs about policies. We will refer to this quantity as a *control cost*; namely, the surprisal or negative log likelihood of the precision taking a particular value, see Equation 4:

$$\underbrace{\mathrm{G}(\gamma)}_{Control\ cost} = -\ln p(\gamma) = \underbrace{\mathrm{D}_{KL}[p_G(\pi|\gamma) \,||\, p_E(\pi)]}_{Complexity} - \underbrace{\mathrm{E}_{q(o|\gamma)}[\ln p(o)] + \mathrm{E}_{q(o|\gamma)}\big[\mathrm{H}[p(o|s)]\big] - \mathrm{E}_{p(o|s)q(o|\gamma)}[\mathrm{D}_{KL}[q(\mathrm{A}|o,s) \,||\, q(\mathrm{A})]]}_{Expected\ cost}$$

(4)

$$p_G(\pi|\gamma) = \sigma(\gamma\mathbf{G}), \qquad p_E(\pi) = \sigma(\mathbf{E}),$$

$$q(o|\gamma) = \sum_{s,\pi} q(o,s,\pi|\gamma), \quad q(o,s,\pi|\gamma) = q(o|s)q(s|\pi)q(\pi|\gamma)$$

$$q(\pi|\gamma) = \sigma(\ln\mathbf{E} - \gamma\mathbf{G})$$

$$q(o|s) = \mathbf{A}$$

$$q(s|\pi) = \mathbf{B}(\pi)s$$

$$\ln p(o) = \mathbf{C}$$

This control cost has two terms. The first term (*complexity*) ensures that the most likely precision minimises the divergence between beliefs about deliberative and habitual policies. One can regard this divergence as *mental effort*; namely, the Kullback–Leibler (KL) divergence between beliefs about the policy to pursue under deliberate control, given the control signal γ, and the policy to pursue under habitual control. Intuitively, the habit is assumed to be a "default policy" or an initial bias about how to act that can be engaged at no cost, reflecting the assumption that decision-makers are intrinsically biased towards low-effort options (Botvinick et al., 2009; Jimura et al., 2010; Kool et al., 2010; Kool & Botvinick, 2014). Conversely, in conditions of *cognitive conflict* between a habitual and a deliberate policy, the agent might engage mental effort to pursue a deliberate policy that differs from the habit (see Figure 6.3 for a clarification of the difference between *cognitive conflict* and *mental effort*). Therefore, the complexity term reflects how much an agent diverges from the habit when engaging in deliberate behaviour (Rubin et al., 2012; Todorov, 2009; Zénon et al., 2019). Since the relative weight of deliberate behaviour depends on the strength of the control signal (i.e., the value of γ), selecting a stronger control signal to ensure that the action is deliberative induces higher mental effort. Rather, selecting a weaker control signal causes behaviour to be dominated by habitual priors and is associated with lower mental effort.

The second term (*expected cost*) represents the utility of selecting a specific value of the γ parameter and ensures that the policies that result from the most likely γ are valuable. Note that – as usual in active inference (Friston, Samothrakis, et al., 2012; Friston et al., 2014) – the notion of a “valuable” policy encompasses three terms. The first term is an extrinsic value that depends on prior preferences. The second and third terms score the resolution of uncertainty associated with the hidden states (i.e., ambiguity-avoidance) and resolution of uncertainty associated with the model's parameters (i.e., novelty-seeking), respectively. The fact that active inference considers both the preferred outcomes and the resolution of uncertainty will become important later when we explain cognitive control dynamics (and associated brain responses in the dACC) when pursuing both rewarding (Shenhav et al., 2013) and informative (Behrens et al., 2007) outcomes.

It is important to remark that the expected free energy G(γ) in Equation 4 (that optimizes the control signal γ) is different from the expected free energy G(π) of Equation 2 (that scores policies), in at least two ways. First, G(γ) includes a complexity term that scores mental effort. Second, it includes a notion of expected cost that also depends upon the precision – because the predicted outcomes depend upon subsequent states, that depend upon the policies that depend upon precision.

In summary, the specification of a control signal corresponds to the minimization of expected free energy G(γ) – or the optimization of the precision parameter γ that balances habitual and deliberate components of action selection. In this setting, G(γ) (i.e., control cost) can be seen as a hierarchical inference level that plays the role of a prior over precision that ensures the achievement of the most valuable and informative outcomes — for example, that aversive outcomes are avoided — subject to the constraint that cognitive effort is minimised; in the sense that beliefs about deliberative policies do not diverge markedly from beliefs about habitual policies.

Therefore, the precision parameter γ plays the role of a control signal and its optimization permits defining its "identity" and the "intensity" (Shenhav et al., 2013). Signal identity specification corresponds to choosing a γ that prioritizes either the policy driven by habits (**E**), *drive forward*, or the policy driven by deliberate control (**G**), *dodge to the left*. This is determined by the direction of change of γ: when γ decreases, the habitual policy (*drive forward*) is prioritized; whereas when γ increases, the deliberate policy (*dodge to the left*) is prioritized. Signal intensity corresponds to the specific value of γ that results from the minimization of the expected free energy G(γ) (intended as a control cost): the higher the signal, the greater the engagement of deliberate control and the requested mental effort.

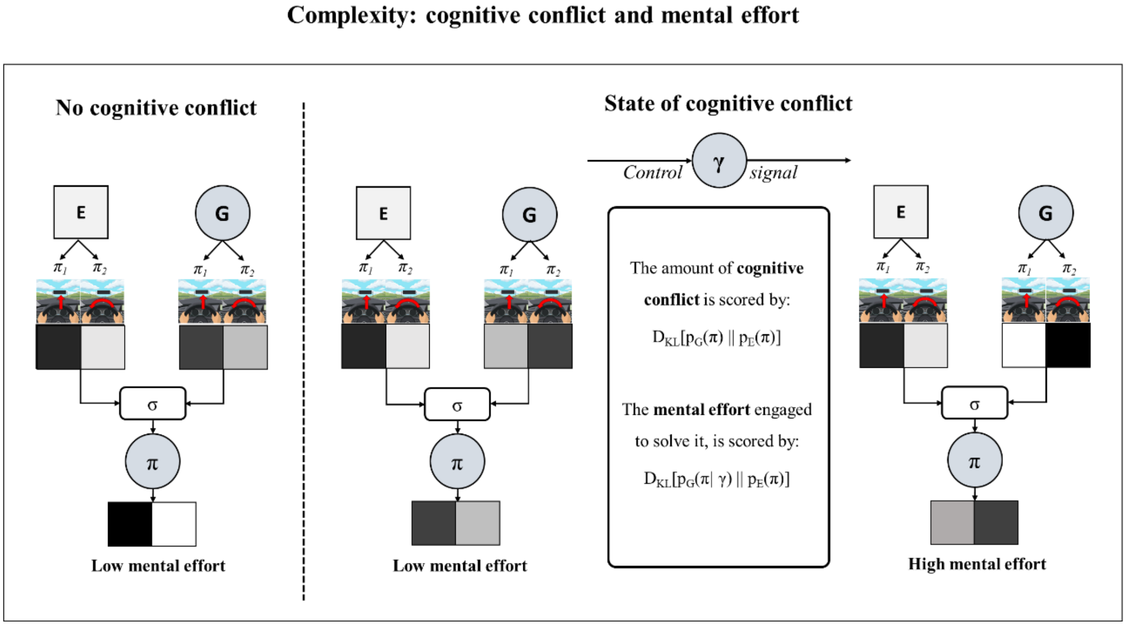

Figure 6.3: Clarification of the differences between cognitive conflict and mental effort. Cognitive conflict corresponds to the KL divergence between beliefs about the policy to pursue under deliberate control and the policy to pursue under habitual control. It is high when the KL between the habitual and deliberate policy distributions is high; conversely, if the two policies were the same, there would be no cognitive conflict to resolve. Mental effort is a complexity cost required to resolve cognitive conflict: the KL divergence between beliefs about the deliberate policy, given the control signal γ and the habitual policy. In case of a cognitive conflict, the agent can either invest low mental effort, hence prioritizing habitual policies, or invest high mental effort, hence prioritizing deliberate policies.

Monitoring changes in the state of the environment. In active inference, decisions are made on the basis of the agent's current estimate of the state of the environment and its beliefs about the relations between actions and outcomes. Monitoring significant *changes* in the state of the environment is key for accurate control, as it could signal that the existing action-outcome contingencies are no longer valid. For example, the unexpected appearance of a warning signal might indicate that the usual *driving forward* policy would no longer afford a safe journey, and the agent needs to learn the novel task contingencies to plan effectively.

We assume that the active inference agent constantly monitors the changes in the state of the environment, which are formalized here as Bayesian surprise, or the KL divergence between the agent's beliefs about the state of the environment at the current and the previous time steps. When Bayesian surprise exceeds a threshold (set arbitrarily to 1), two model parameters are modified. First, the entropy of the *action monitoring* matrix (i.e., the portion of the **a** *matrix* that links actions and outcomes) is increased; see Figure 6.4. The rationale is that when the agent experiences high Bayesian surprise, it becomes less certain about the links between its behaviour (*driving forward* - *dodging to the left*) and its outcomes (*injured deer* – *safe travel*). Second, the learning rate (η parameter) of the same *action monitoring* matrix is increased (note that the other parts of **a** are not updated during the simulations). This implies that the agent updates its model faster in the presence of novel observations and hence can re-learn efficiently the relations between her actions and their subsequent outcomes. Taken together, these two changes imply that in (volatile) situations characterized by high Bayesian surprise, the agent is less confident about how to act to achieve desired outcomes but learns the novel task contingencies faster (Behrens et al., 2007; Sales et al., 2019).

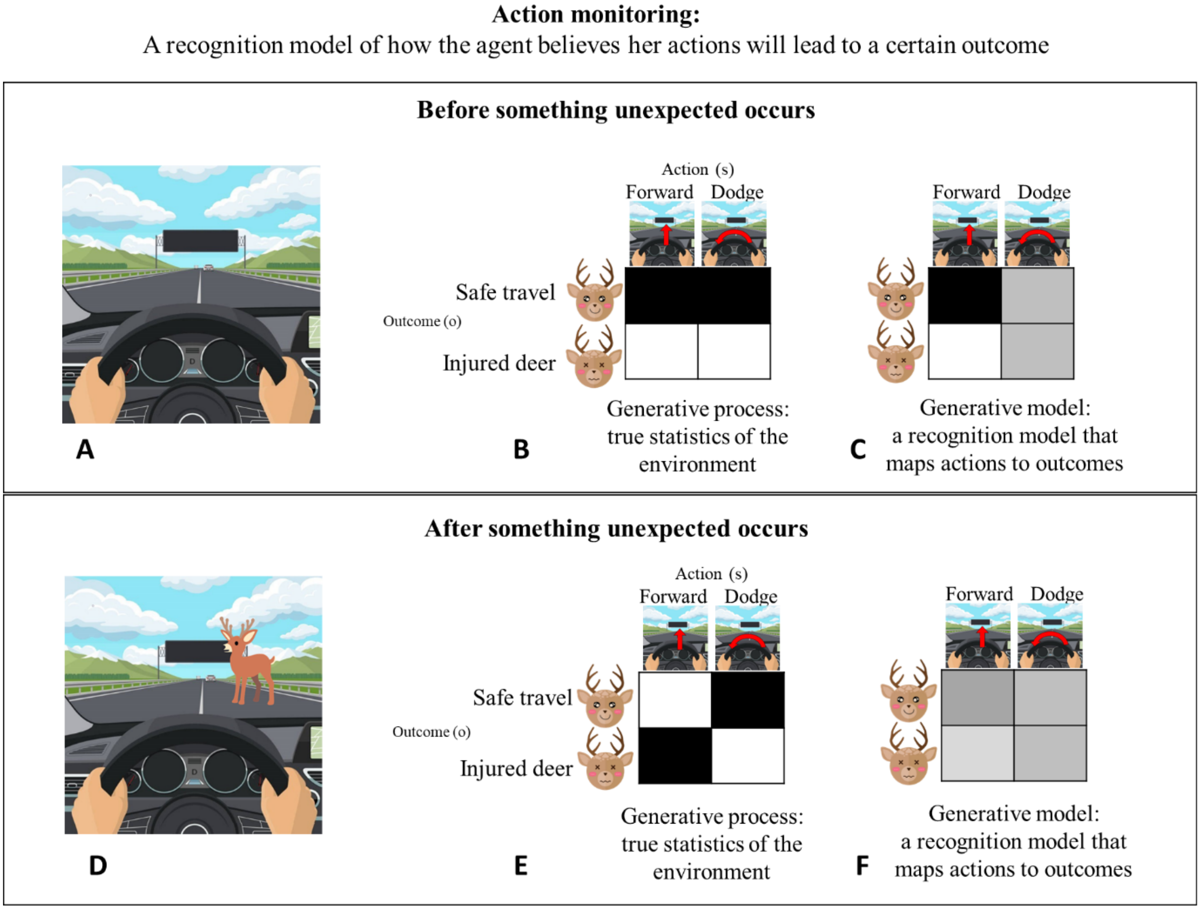
Action monitoring:
A recognition model of how the agent believes her actions will lead to a certain outcome

Before something unexpected occurs

After something unexpected occurs

Figure 6.4: How the action monitoring matrix (i.e., the portion of the ***a*** *matrix that links actions and outcomes) changes after significant surprises. (A-C) When the road is free and no unexpected events are observed, the driver's beliefs about her current policy (drive forward) reflect the true statistics of the environment, compare the left columns shown in B and C, which show action-outcome contingencies in the generative process and generative model, respectively. (D-F) However, unexpected events like the presence of a deer might signal that*

action-outcome contingencies in the generative process might have changed (i.e., volatility), as shown in (E). In the presence of unexpected events, the entropy of the action monitoring matrix (F) is increased, which implies that the agent is no more confident that she will obtain her preferred outcomes by following the drive forward policy. In parallel, the learning rate (η parameter) of the same matrix is increased, too, which implies that the matrix can learn novel contingencies faster.

6.3 Neuronal dynamics of cognitive control

Active inference can be used to simulate putative neuronal dynamics associated with the free energy minimization and the optimization of model parameters (Bastos et al., 2012; Da Costa et al., 2021; Friston, 2005; Friston et al., 2006; Friston et al., 2017; Isomura, 2021; Parr et al., 2022). Below we focus on how we simulate neuronal dynamics in two brain structures that play a crucial role in cognitive control: the dACC and the dopaminergic (DA) system.

6.3.1 Neuronal dynamics in the dACC and the specification of a control signal

In this setting, we associate neuronal dynamics in the dACC to the specification of a cognitive control signal, represented by the precision parameter γ. The model finds the optimal value for γ via a gradient descent of expected free energy $G(\gamma)$, by computing the derivative of equation 4[3].

In our subsequent simulations, the final (optimized) value of the control signal γ for each trial, which we associated with dACC activity, will be illustrated in panels D of Figures 6.6-6.8. Please note that the gradient of the expected free energy $G(\gamma)$ and the optimal control signal γ change at each simulation trial, because the parameters of the generative model are continuously updated and the agent collects more observations that change the expected free energy of her policies and accumulates more evidence in the **E** matrix.

6.3.2 Dopaminergic activity

Previous implementations of active inference (without cognitive control) assumed that dopaminergic activity corresponds to the precision of policies γ, which in turn indexes how much an agent is confident in her choices (Friston et al., 2014; Langdon et al., 2018; Schwartenbeck, FitzGerald, Mathys, Dolan, & Friston, 2015). In this setting, policy precision updates after receiving new observations reflect phasic dopaminergic responses in the brain. Positive spikes reflect the fact that novel observations increase the agent's confidence

---

[3] For ease of implementation, we simplified the model described in the main text in two ways. First, we used the values of **G**, **e** and **a** from the previous trial, not the current trial (note that numerically these values are very similar). Second, instead of computing the gradient of expected free energy for each component of Equation 4, we considered the estimated probability of each outcome weighted by the value of the $\mathbf{G(\pi)}$ vector associated to that outcome.

about the policy she is pursuing, whereas negative spikes reflect the fact that the novel observations decrease her confidence:

$$\beta_0 = \frac{1}{\gamma} \qquad (6)$$

$$\beta_{updated} = \beta - \beta_0 + \underbrace{(\pi_{posterior} - \pi_{prior}) * G_\pi}_{G\ error} \qquad (7)$$

with the posterior belief about the policies being:

$$\pi = \sigma\left(\ln E - F - \frac{1}{\beta}G\right) \qquad (8)$$

Dopamine, therefore, plays a central role in belief updating, as it encodes the meaningful information content of observations (FitzGerald et al., 2015; Friston, Shiner, et al., 2012; Friston et al., 2014; Schwartenbeck, FitzGerald, Mathys, Dolan, Kronbichler, et al., 2015) which under certain conditions resemble reward prediction errors (Cohen et al., 2012; D'Ardenne et al., 2008; Flagel et al., 2011; Schultz, 1998; Schultz et al., 1997). Importantly, at difference with previous implementations of active inference (without cognitive control), in the current model (with cognitive control) the precision parameter γ is optimized twice: first, during the specification of the control signal (Equation 5, discussed in the previous section) and second, when new observations are gathered that change the confidence in the policies (Equation 7). The latter update, described in Equation 7, depends on a "G error" that scores the difference between posterior and prior policies – and has been called an "affective charge" in previous active inference studies (Hesp et al., 2021). To the extent that dopaminergic activity corresponds to the precision parameter γ, it combines the contributions of both Equation 5 and 7, respectively. In our subsequent simulations, these two contributions are illustrated separately (in panels D and F, respectively, of Figures 6.6-6.8), and they are subsequently summed up (in panels G of Figures 6.6-6.8) to model dopaminergic activity.

6.4 Simulations of cognitive control in the driving task

Below we illustrate three simulations of cognitive control in the driving scenario. The three simulations use the same active inference model and parameters but present the driver with different challenges. In Simulation 1, the driver does not encounter any unexpected danger and can safely drive towards her office. This simulation illustrates habitization and the diminishing costs of cognitive control in habitual settings. In the other two simulations, the driver encounters an unexpected danger (Simulation 2) or a danger signalled by a traffic sign (Simulation 3). These two simulations illustrate how the demands of cognitive control increase under changing environmental conditions - and in which situations exerting cognitive control can be effective or not effective.

6.4.1 Simulation 1: Habitization and the decrease of “attention to action” in habitual settings

The first simulation of the driving scenario exemplifies a case in which the driver can safely proceed from home to the office because no unexpected events occur along the way. The driver can thus safely select the policy *drive forward* at each trial and achieve the desired outcome (safe travel) at the last trial (Figure 6.5A). However, the process that leads to the selection of the *drive forward* policy differs during the first and subsequent trials, respectively. During the first trials, the agent relies on deliberate control (driven by **G**) and invests cognitive effort (Figure 6.5E), hence eliciting activation of dACC (Figure 6.5D). During the task, however, the agent experiences a stable context, with no unexpected events (Bayesian surprise decreases significantly, Figure 6.5B). Consequently, the policy to drive forward slowly becomes a habit, as evidence for the policy is accumulated in **E** (Friston et al., 2016; Maisto et al., 2019). This habitual learning decreases the demands for cognitive control and attentional resources and permits selecting "chunks" of actions (Anderson et al., 2019; Botvinick, Huffstetler, et al., 2009; Dezfouli et al., 2014; Dezfouli & Balleine, 2013; Miller et al., 1960; Taatgen & Lee, 2003; Tessari et al., 2006, 2021, 2022).

This habitization of behavior is reflected in various ways in the model. The decreased cognitive demands become evident by considering that expected free energy G(γ) (Figure 6.5C) and the dACC activity that represents the control signal (Figure 6.5D) decrease with time, which implies that desired outcomes (safe travel) can be obtained at a smaller cognitive cost. The decrease in mental effort shown in Figure 6.5E is another index of habitization, and the fact the agent experiences no cognitive conflict between her habitual policy (to *drive forward*) and alternative policies (Figure 6.5H). Furthermore, the increased precision of the information in the "monitoring matrix" (Figure 6.5I) shows that the agent becomes increasingly confident in the ways her actions produce desired outcomes, which in turn causes G(γ) to decrease. Finally, the decreased cognitive demands are reflected in dopaminergic activity (Figure 6.5G). Please note the decrease in dopaminergic activity over time reflects the decrease of both the control signal (Figure 6.5D) and the **G** error (Figure 6.5F).

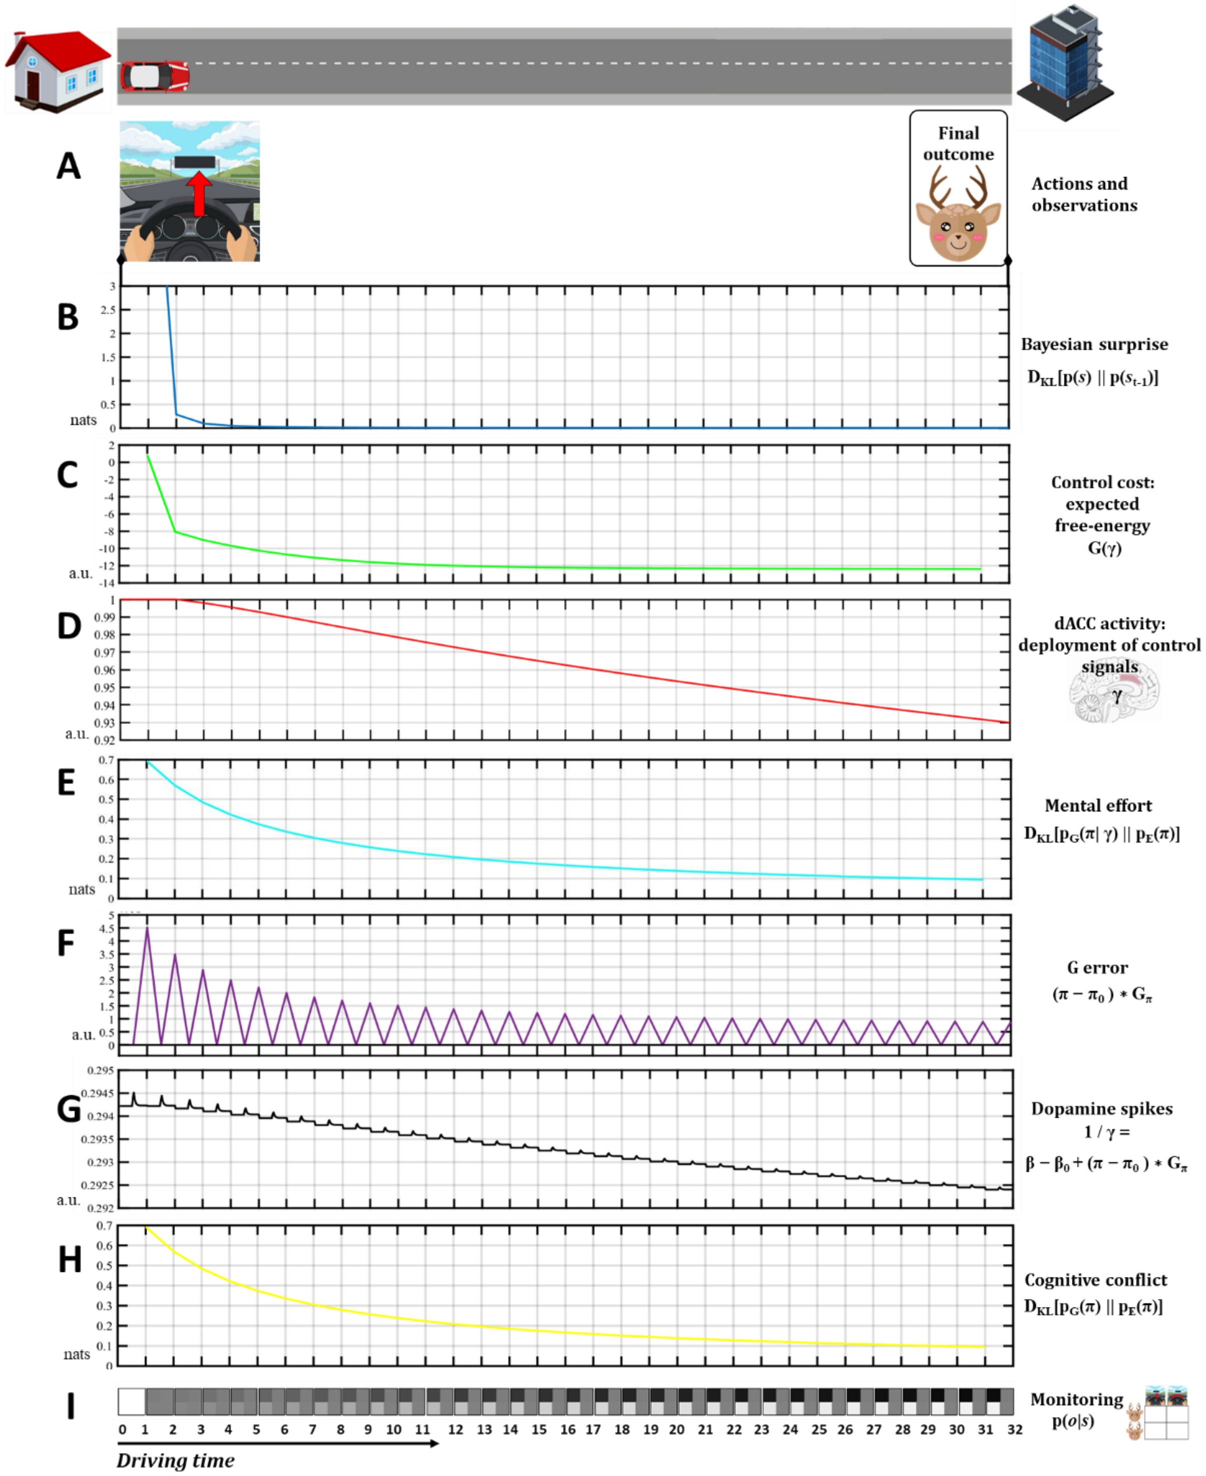

Figure 6.5: Results of Simulation 1. (A) This simulation illustrates the case of a driver who drives safely over 32 trials (or road-steps), without unexpected events. The simulation illustrates several key variables of cognitive control. These include (B) Bayesian surprise; (C) expected free energy G(γ); (D) the precision parameter γ, which we associated with a control signal and dACC activity (please note that γ starts from 1 and it can increase or decrease, but it only decreases in this simulation); (E) mental effort; (F) "G error", which indexes the confidence in expected free energy; (G) simulated dopaminergic responses; (H) cognitive conflict; and (I) the monitoring matrix. In this simulation, all the variables shown in panels B-H decrease over time,

whereas the part of the monitoring matrix associated with "driving forward" (shown in panel I) becomes increasingly more precise. Please see the main text for explanations.

6.4.2 Discussion of simulation 1

This simulation exemplifies the habitization of behaviour that arises in stable environmental conditions when the driver does not experience significant surprises. Developing habits in stable conditions is an efficient solution to the trade-off humans and other animals have to solve – about selecting the best course of action without wasting excessive resources. In reinforcement learning, it is common to cast this trade-off in terms of a competition between a model-free controller, which is cheap but rigid, and a model-based controller, which is more cognitively expensive but more flexible, and which might have distinct neuronal substrates (S. W. Lee et al., 2014; Abrahamse et al., 2013; Dezfouli & Balleine, 2013; Pezzulo et al., 2015). After (over)training, the action becomes habitual, that is, guided by model-free methods based on cached action values, or by stimulus-response mechanisms, which can be engaged automatically and with minimal processing resources enabling fast and skilled performance (Dezfouli et al., 2014; Dezfouli & Balleine, 2013; Keramati et al., 2011; Maisto, et al., 2019). It is argued that habitization depends on the fact that after sufficient learning, the model-free controller becomes less uncertain than the model-based controller and is selected to control behavior (Daw et al., 2005).

Here, we offer a different (but perhaps complementary) perspective on the process that guides the formation of a habit. There are at least four important points that characterize our proposal. First, the passage from a more complex to a simpler controller (in our proposal, **G** and **E**) is guided by a control signal that is continuously optimized based on internal estimates of Bayesian surprise and cognitive conflict, not a comparison of the uncertainty of the two controllers. The second related point is that in our proposal, the transition to habits is smooth because the different (simpler and more complex) controllers do not "compete" but are weighted by the control signal. In our simulation, the control signal becomes increasingly smaller and this reflects a smooth process of habitization over time (Daw et al., 2011; Dolan & Dayan, 2013; D. Lee et al., 2012; Otto et al., 2013; Pezzulo, Rigoli, et al., 2013). Third, in our proposal, habitization reflects not only the passage from a more complex to a simpler controller but also a series of other changes at the cognitive monitoring and control variables level. As explained above, several variables such as mental effort, cognitive conflict, the precision of the monitoring matrix and learning rate all change in a coordinated manner to pass from a mode of "high attention to action" to a mode of "low attention to action". This is in keeping with broad theoretical perspectives that emphasize the importance of higher-level cognitive control mechanisms in mediating the trade-off between simpler and more complex behavioural controllers (Behrens et al., 2007; Norman & Shallice, 1986; Shenhav et al., 2013). Finally, and importantly, we show that the solution to the above problems of cognitive control and the balance of simpler and more complex controllers can be characterized within a biologically plausible (free energy minimization) scheme that optimizes both performance and cognitive resources.

In sum, simulation 1 illustrated a situation where the demands for cognitive control gradually decrease and behaviour can be executed with low “attention to action” – passing from deliberate to habitual behaviour. The next two simulations illustrate the opposite case: a situation in which the demands for cognitive control (that are initially low) suddenly increase and the system has to suddenly put “attention to action” – hence passing from habitual to deliberate behaviour.

6.4.3 Simulation 2: the failure to increase “attention to action” in time after surprising events

Simulation 2 illustrates a more challenging scenario, in which the driver initially drives safely towards the office but she subsequently sees an unexpected deer along the road (at trial 19), which she hits shortly after (at trial 27); see Figure 6.6. The first part of the simulation, before trial 19, is analogous to Simulation 1. For simplicity, we reused exactly the *e* parameters learned during Simulation 1; and this is why the other parameters (Bayesian surprise, conflict and so on) settle fast during the first trials. During this period, the driver pursues the *drive forward* policy, which shortly becomes a habitual behaviour, requiring little cognitive control (see the low values of the variables shown in Figures 6.7B-E). However, when the driver sees the unexpected deer at trial 19, Bayesian surprise suddenly increases (Figure 6.6B), signalling changed environmental conditions (or volatility) and producing a cascade of effects. The driver is not confident anymore that her current plan will bring the desired outcomes, meaning that her monitoring matrix becomes more entropic (Figure 6.6I). She experiences a situation of cognitive conflict (Figure 6.6H) and has to assess whether, to avoid the danger, she must override her habitual behaviour and engage the deliberate policy to dodge instead. Managing this conflict increases mental effort (Figure 6.6E) and resources for cognitive control (Figures 6.7C-D) to devaluate habitual control in favour of deliberate control.

However, a crucial thing to notice is that deciding on the right policy to avoid danger requires time. In this simple example, we assume that the driver is inexperienced and hence firstly needs to re-build some precision in the monitoring matrix (Figure 6.2I). In the meantime, the driver has little confidence in her current action policy, as evident in the dopaminergic responses (Figure 6.6G), which show increasingly larger negative peaks, as the driver approaches the deer. In this specific simulation, the re-planning and the increase in cognitive control (Figure 6.2D) are not fast enough; hence, the driver hits the deer at trial 27, producing a dis-preferred (*injured deer*) outcome. This simulation, therefore, illustrates the case of an inexperienced driver who drives habitually, with low “attention to action” – and is not fast enough to adapt to changed environmental contingencies.

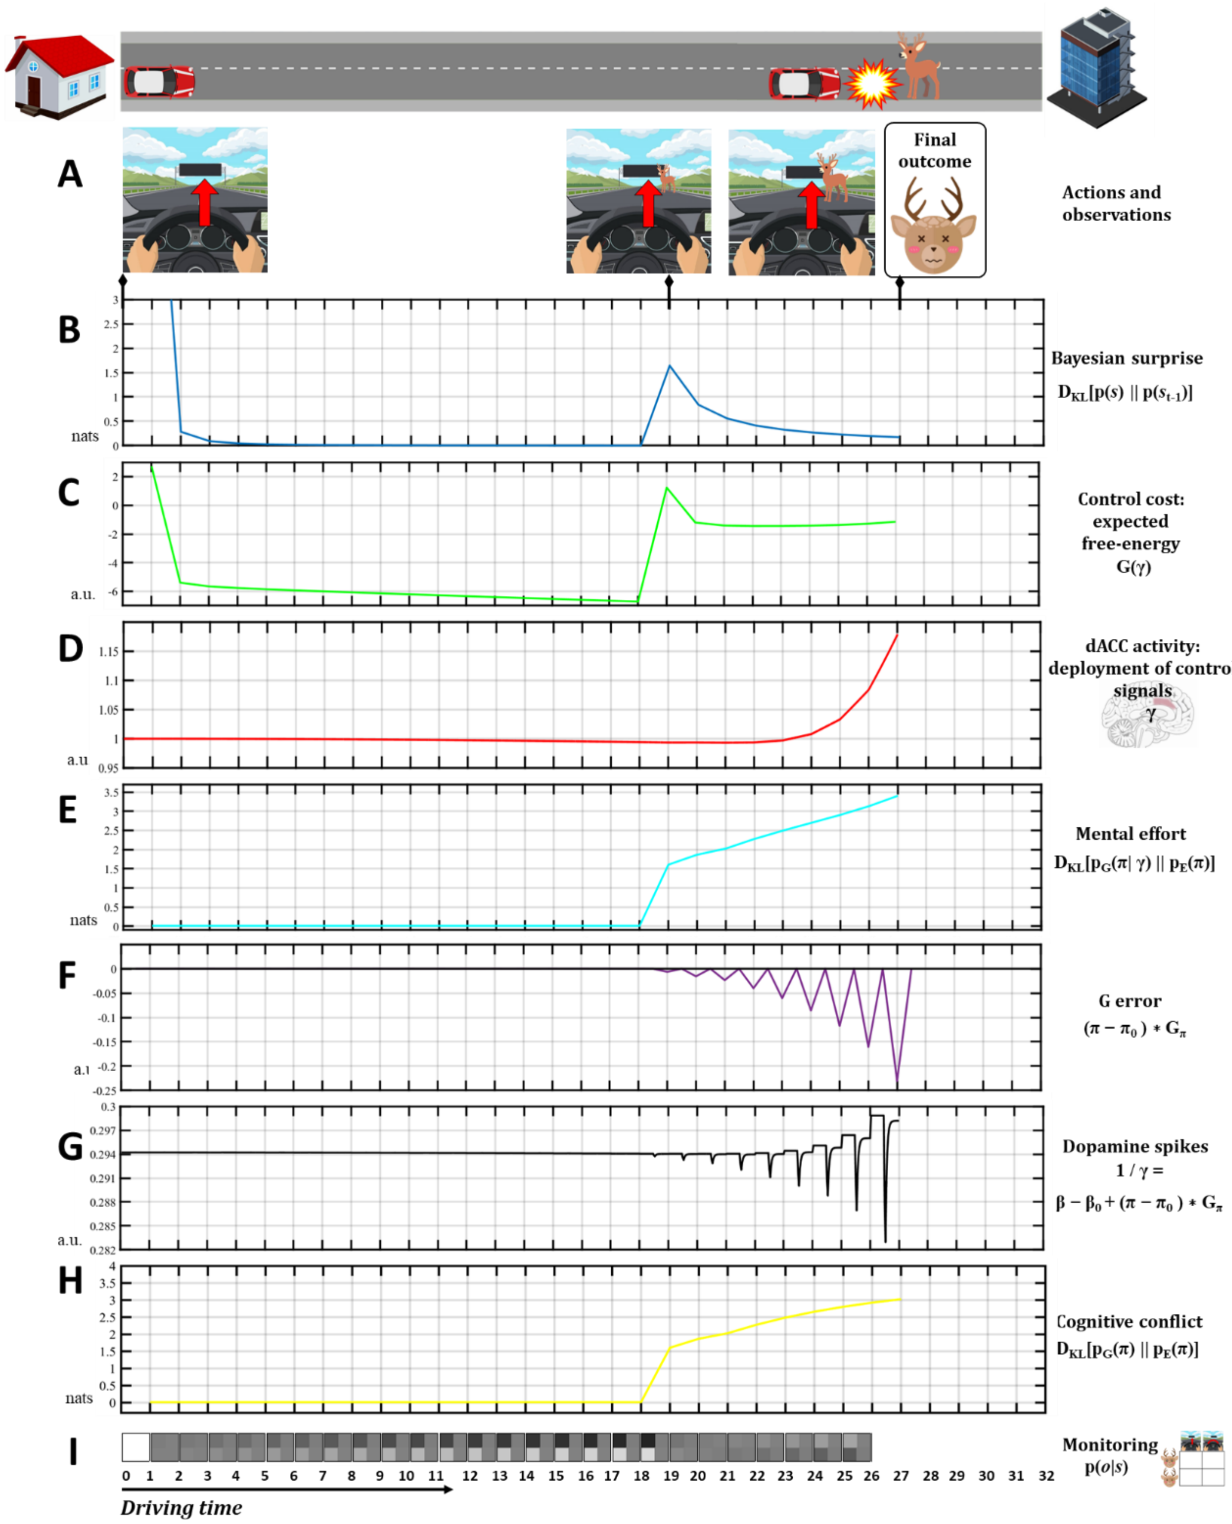

Figure 6.6: Results of Simulation 2. A driver encounters an unexpected deer on the road (trial 19) and hits it (trial 27), as it is not fast enough to re-adapt to the novel situation. Panels A-I are the same as in Figure 6.6. See the main text for an explanation.

Please note that we could have provided the driver's generative model with additional pre-configured monitoring matrices to deal with dangerous situations, like a deer on the road. In this way, the driver could have solved the problem of avoiding the deer by inferring the most contextually appropriate monitoring matrix, and then "loading" it. However, we decided to simulate an inexperienced driver with no pre-configured monitoring matrices (whose learning rate is not unrealistically high). This helps us illustrate that significant environmental changes (as signalled by high Bayesian surprise) promote model learning – which in our model, depends on the fact that the monitoring matrix becomes more entropic and the learning rate increases – but this, in turn, requires some time.

6.4.4 Discussion of simulation 2

This simulation exemplifies a failure to engage in deliberative behaviour (and to devaluate habitual behaviour) because cognitive control is engaged too late or too weakly. In the simulation, the driver experiences a significant *cognitive conflict* between the habitual policy to *drive forward* and the deliberate policy to *dodge to the left.* This kind of cognitive conflict has been extensively linked to cognitive control dynamics. One early proposal is that conflict can lead to compensatory adjustments in perceptual and attention processing (Berlyne, 1960). More recently, it has been proposed that conflict directly influences the regulation of cognitive control (Botvinik et al., 2001). This research identified several types of events that may trigger cognitive conflict – namely, explicit feedbacks about one's performance, the anticipation of a negative outcome or the observation of unexpected environmental events. All these events indicate a potential need to increase the intensity (and specify the identity) of the control signal.

In turn, context monitoring is essential to determine the identity and intensity of the control signal. Formally, this emerges from the fact that the expected accuracy of the expected free energy $G(\gamma)$ multiplies the value of an outcome by the probability of achieving it. When Bayesian surprise is high, the monitoring matrix loses precision and $G(\gamma)$ increases, because the driver is not sure anymore that she will obtain her preferred outcome by pursuing the current course of action. In our simulation, between trials 19 and 27, the driver accumulates evidence indicating that the *drive forward* policy will lead to an undesired outcome (*injured deer*). This permits minimizing $G(\gamma)$ by changing plan and engaging in the deliberate policy *dodge to the left*, which predicts a desired outcome (*safe travel*).

Unfortunately, in this specific simulation, cognitive control is engaged too late (but see the next simulation for a more positive outcome). This simulation, therefore, exemplifies a case in which, despite the cognitive conflict has been noticed, an insufficient level of control was allocated to the task and then habitual control was insufficiently devaluated (Botvinick et al., 2001, 2004). This example is representative of a broader class of conditions, described as *default override* (Shenav et al., 2013). Several studies reported dACC activity in tasks that require overcoming an automatic but task-inappropriate response (Silton et al., 2010). In these situations, the desired behaviour is less automatic than the default behaviour that occurs automatically – and represents the source of the conflict (Shenav et al., 2013).

Another relevant example is reversal learning in the context of motor control when overtrained actions can become resistant to devaluation. This manifests as an after-effect on motor adaptation (Izawa et al., 2008; Wei & Körding, 2009). For example, for a pilot, driving a car is an overlearned procedural task that uses a habitual behavioural controller and low cognitive resources. When a perturbation is introduced in the control task (for example, the steering direction of the wheel is inverted, such that when turning right, the car steers to the left), there is a strong "overriding" of habitual control that can lead to the wrong outcomes. Devaluating the motor habit requires a very high intensity of the control signal (Shadmehr et al., 2010; M. A. Smith et al., 2006). From an active inference perspective, if deliberate and habitual controllers are arranged in a hierarchy (Pezzulo, et al., 2015), overtraining produces very strong beliefs in the lower (habitual) hierarchical level and therefore, the descending predictions generated by the higher (deliberate) levels cannot contextualize them and override the default behaviour.

There are two ways to solve (or at least alleviate) this problem. First, the agent can engage a strong control signal. Indeed, a large body of evidence indicates that when a cognitive conflict occurs, a control signal must have an adequate intensity to override in time the default response (Cavanagh et al., 2013; De Martino et al., 2006; Paus et al., 1993; Shenhav et al., 2013). Second, the agent can be endowed with a precise monitoring matrix, which increases the confidence that the deliberate model will produce desired outcomes and prioritizes it. The next simulation will show that when these conditions hold, the active inference agent is able to override habitual behaviour and select the most adaptive course of actions to avoid danger.

6.4.5 Simulation 3: increasing "attention to action" after a warning signal permits avoiding dangers

Simulation 3 illustrates the case of a driver who initially drives safely towards her office, but then (at trial 10) encounters an unexpected warning signal that informs her that the environment is volatile, hence decreasing her confidence in the current policy and increasing her learning rate. As a consequence of these preparatory processes, the driver has sufficient time and resources to subsequently avoid the deer, which she encounters at trial 19; see Figure 6.7.

As in the two previous simulations, the driver initially pursues the habitual policy to *drive forward*. At trial 10 (Figure 6.7A) the driver encounters a warning signal that is unexpected and therefore signals some volatility (Figure 6.7B). To make the simulation more challenging, we treat the warning signal as a generic advice that does not directly increase the belief about encountering wild animals along the road. Rather, the warning signal makes the driver less confident that her current course of actions will lead to the desired outcome, rendering her monitoring matrix less precise (Figure 6.7I) and increasing her learning rate – which are exactly the preconditions that promote model revision. The imprecise mapping between actions and outcomes increases the level of G(γ) (Figure 6.7C) and the mental effort (Figure 6.7E). At trial 19, the driver encounters a wild animal on the road. Different from Simulation 2, this surprising event (Figure 6.7B) comes while the agent had already estimated the environment to be volatile and was paying more "attention to action" – as evident by the high level of G(γ) (Figure 6.7C) and the high learning rate. This is why the driver is able to deploy an appropriate control signal in time (Figure 6.7D) and ultimately select the policy to *dodge to the* *left* that avoids the deer at trial 24 (Figure 6.7A). Interestingly, the time interval between trial 19 (when the deer is first seen) and trial 24 (when the deer is avoided) is characterized by a rapid switch in the valence assigned to the situation, with a rapid passage from the prediction of a negative outcome (negative peaks in Figure 6.7F, between trials 20-22) to the prediction of a positive outcome (positive peaks in Figure 6.7F, between trials 23-25) when the policy to *dodge to the left* is selected. The same dynamics is reflected in the simulated dopaminergic activity shown in Panel 8G, which sums up the contributions of the control signal (Panel 8D) and the G error (Panel 8F).

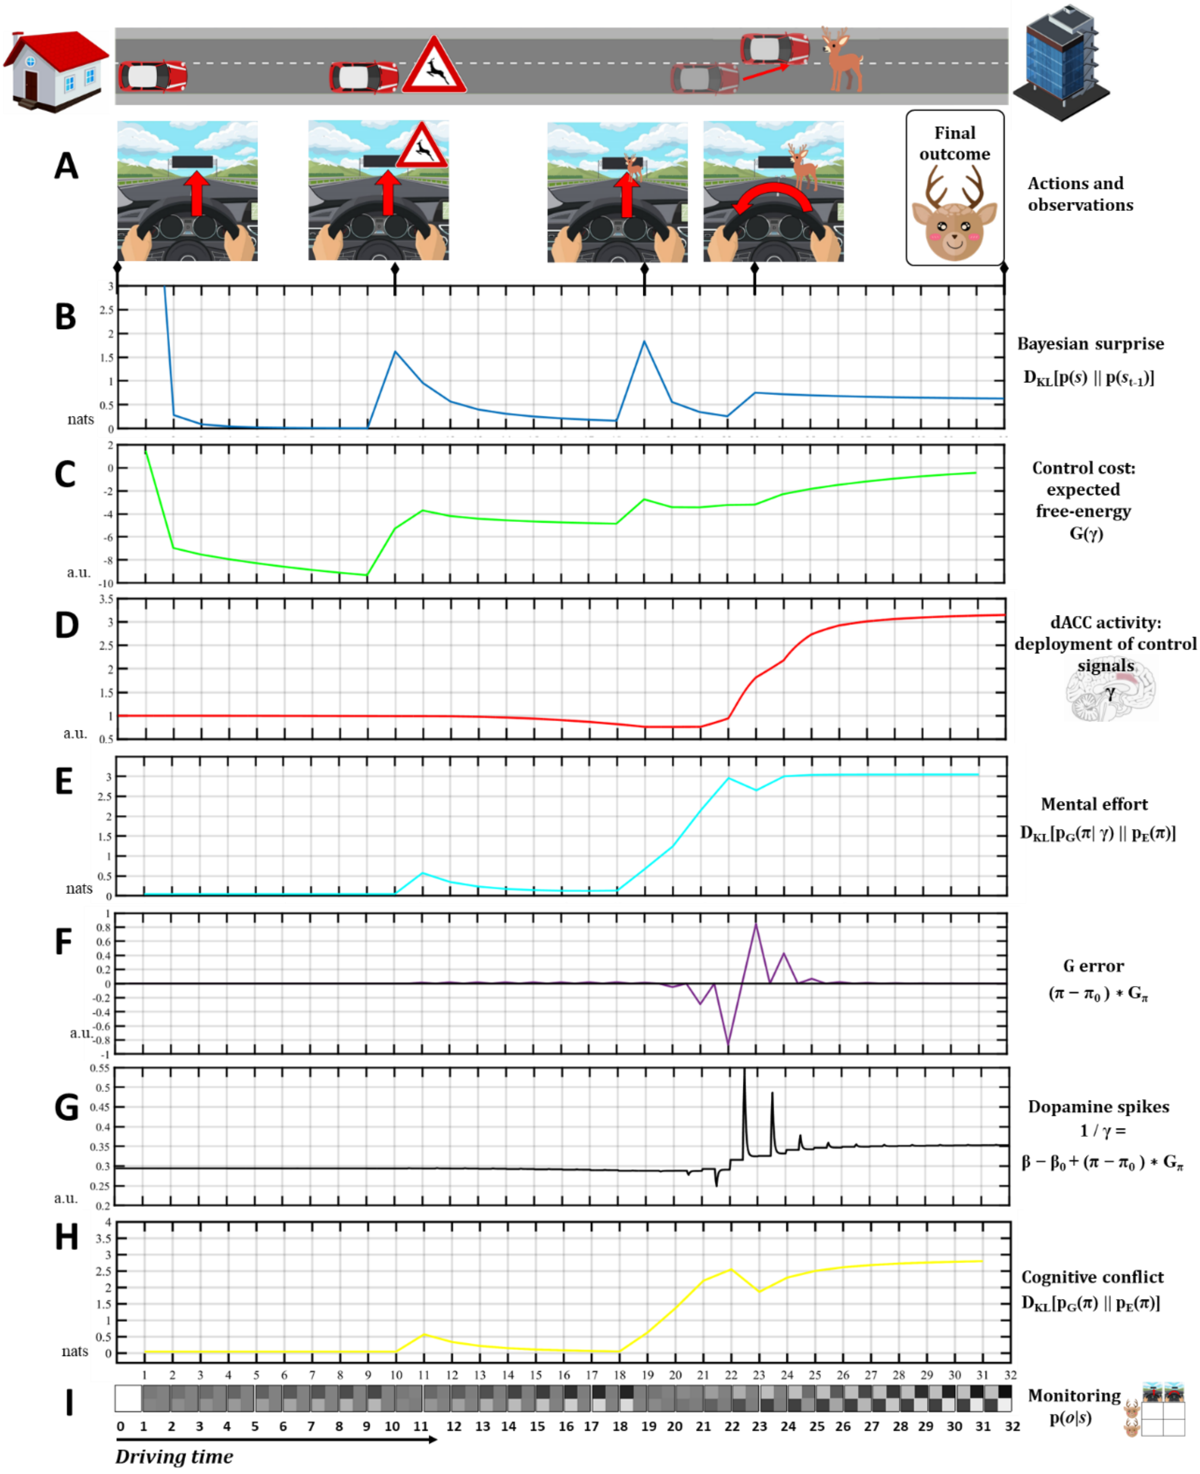

Figure 6.7: Results of Simulation 3. A driver encounters an unexpected warning signal (at trial 10) that triggers a series of preparatory processes, which help her avoid a deer that she subsequently encounters on the road (trial 19). Panels A-I are the same as in Figure 6.6. *See the main text for an explanation.*

6.4.6 Discussion of simulation 3

This simulation illustrated the importance of warning signals to pre-empt unexpected and aversive outcomes (Gabriel & Orona, 1982). At the behavioural level, the warning signal makes the driver more ready to change her course of action. This is because, in our simulation, the warning signal is a cue that the environment is volatile, determining an increased “attention to action” and learning rate. The fact that higher levels of attention and learning rates are applied in more volatile environments is in line with a large amount of evidence that links prediction errors to task accuracy and effort (Piray & Daw, 2020; Sales et al., 2019). For example, participants solving the Stroop task show less interference in incongruent trials that are frequent compared to the case in which they are rare (Lindsay & Jacoby, 1994; Logan & Zbrodoff, 1979). Similar results can be observed in imitation studies that manipulated the volatility of stimuli familiarity. Participants show a decrease in accuracy and an increase in effort and deliberation when they process familiar stimuli that are intermingled with unfamiliar stimuli, compared to the case in which the same stimuli are more predictable (Tessari et al., 2007, 2022; Tessari & Rumiati, 2004). Moreover, trying to exert cognitive control when a cue was given in both predictable and unpredictable conditions was further detrimental in RTs (i.e. processing time; Tessari et al., 2022).

Several theories in cognitive science emphasize that expectations and prediction errors play a key role in learning, too (Daw et al., 2011; Rescorla & Wagner, 1972). For example, according to the Pearce-Hall theory of learning, surprising events lead to increased attention levels that facilitate learning (Pearce & Hall, 1980; Roesch et al., 2012). In keeping, it has been repeatedly shown that the violation of expectations promotes learning (Behrens et al., 2007; Courville et al., 2006; Jacquet et al., 2012; Mathys et al., 2011; Soltani & Izquierdo, 2019; Swan & Pearce, 1988). According to another perspective, increasing the learning rate as an effect of high prediction errors could be considered a form of control signal that shapes subsequent responses (Shenhav et al., 2013).

At the neurophysiological level, the dopaminergic system, including substantia nigra, the ventral tegmental area, striatum, orbitofrontal cortex, amygdala and other associated structures (Schultz, 2007), has long been linked to learning by reward prediction error (Gershman, 2017; Glimcher, 2011; Schultz et al., 1997). Our framework does not directly link dopaminergic responses to learning by reward prediction error. However, this link is indirect (FitzGerald et al., 2015), as the conditions that promote high dopaminergic responses (e.g., increased control signals) are the same that promote learning (e.g., increased learning rates). In addition to the putative role played by the dopaminergic system, the mechanisms of learning and decision-making also involve the prefrontal and cingulate cortexes (Botvinick et al., 2001; Rushworth & Behrens, 2008; van Veen et al., 2001). For example, changes in context volatility that are associated with improved learning elicit dACC responses (Behrens et al., 2007; Brown & Braver, 2005; Carter et al., 1998); see below for a more detailed discussion of the relations between dopamine and dACC.

6.5 General discussion

We presented a novel theory of cognitive control within the framework of active inference. We showed that the necessary components of cognitive control identified in previous studies – namely, regulation, specification and monitoring (Shenhav et al., 2013) – can be straightforwardly described under the free energy minimization scheme of active inference. We proposed that cognitive control implies the optimization of a cognitive control signal (which corresponds to the precision parameter that balances deliberate and habitual components of control in active inference), via free energy minimization. This (meta-level or hierarchical) free energy minimization loop complements (or is *nested within*) the standard process of free energy minimization that optimizes perception and action in active inference (Pezzulo et al., 2018). Furthermore, we proposed that optimising the control signal (γ) that balances deliberate and habitual control components might reflect neuronal dynamics in the dorsal anterior cingulate cortex (dACC). In contrast, dopaminergic responses might reflect the further optimization of γ in relation to the current observations (i.e., policy scoring), in accordance with the *precision hypothesis* (FitzGerald et al., 2015).

To illustrate the novel theory, we discussed three simulations that exemplify the decrease of cognitive demands (or "attention to action") that leads to habitization of beahvior (Simulation 1) and the increase of cognitive demands that engage stronger cognitive control, showing under which conditions this could be insufficient (Simulation 2) or sufficient (Simulation 3) to select the correct policy to avoid dangers. We decided to illustrate the theory with a "driving task" rather than with classical setups to study cognitive control, such as the Stroop, the Posner and the Eriksen flanker task (MacLeod, 1991; Nee et al., 2007), because the driving task shows more clearly the importance of pursuing a goal and integrating information over time (i.e., across multiple trials), during which environmental conditions may change. Having said this, our simulations could also be taken as representative of *sequences of trials* in the Stroop, Postner, flanker, or similar tasks. Indeed, it has been extensively shown that the learned statistics (and the errors) across trials influence the deployment of cognitive resources and neuronal dynamics in subsequent trials (Gómez et al., 2019).

Throughout our simulations, we illustrated the interplay of the various cognitive variables and mechanisms that have been identified as crucial for cognitive control, which includes, for example, surprise and volatility, context monitoring, cognitive conflict, confidence in one's course of actions and the specification of a control signal. While these variables and mechanisms are sometimes studied independently, casting cognitive control in terms of active inference shows their relationships. For example, in our proposal, Bayesian surprise increases the level of expected free energy $G(\gamma)$, which in case of cognitive conflict may lead to an increase of the control signal γ. While these quantities are distinct, they clearly depend on one another, as illustrated in our simulations. Furthermore, it has been argued that the identity and intensity of the control signal are jointly specified in the brain; this is naturally reflected in our model since signal intensity and identity correspond to the value and direction of change of the same (γ) parameter. Our proposal also distinguishes between cognitive conflict and mental effort concepts that may sometimes be conflated. The former corresponds to the KL divergence between (probabilistic beliefs about) the habitual and the deliberate policies, whereas the latter is a complexity cost that can be invested to resolve the cognitive conflict: the KL divergence between beliefs about the deliberate policy, given the control signal γ and the habitual policy (see Figure 6.3).

Two prominent theories of cognitive control related to our proposal are *attention to action* (Norman & Shallice, 1986) and *expected value of control* (Shenhav et al., 2013). The former theory assumes that a (higher level) supervisory attentional system deploys cognitive control to bias (lower level) action selection, whereas the latter assumes that cognitive control is based on a cost-benefit evaluation that balances the payoff that one can obtain by engaging a controlled process against its cognitive costs. Our proposal is conceptually related to both theories but casts the computations of cognitive control within the overarching framework of active inference (Parr et al., 2022). Crucially, in our approach, cognitive control emerges as a straightforward extension of the (hierarchical) action-perception loop of active inference and does not require heterogeneous mechanisms for optimization, as in previous proposals. In other words, the amount of cognitive control is just another thing that is optimized by minimizing expected free energy. Another distinguishing feature of our proposal is that the expected free energy functional G(γ) is richer than the notion of expected value of control: it considers both reward-related factors that are key in the theory of expected value of control (Shenhav et al., 2013) and epistemic factors, such as environmental uncertainty and volatility, which are highlighted in other studies (Behrens et al., 2007). Furthermore, our approach is based on Bayesian (belief-based) inference, which provides a rich formal language to express the core concepts of cognitive control in rigorous ways; for example, Bayesian surprise can naturally score unexpected events, and KL divergences can naturally score cognitive conflict and mental effort (see above and Figure 6.3).

6.5.1 Putative neuronal underpinnings of the proposed model

Finally, the free energy minimization dynamics have straightforward correspondences with neuronal responses in various brain areas (Da Costa et al., 2021; Friston et al., 2017). Below, we discuss the putative neuronal underpinnings of the proposed model by focusing on three brain structures that are key to cognitive control: the dorsal anterior cingulate cortex (dACC), the dopaminergic (DA) system and the locus coeruleus (LC).

6.5.2 The dorsal anterior cingulate cortex (dACC)

Various studies implied the dorsal anterior cingulate cortex (dACC) in cognitive control and the modulation of adaptive behaviour. An influential proposal reviewed above is that the dACC integrates information about costs and rewards, to compute the net value associated with the allocation of control to a given task, to determine whether and how much control should be invested – and ultimately, to deploy the optimized control signal (Shenhav et al., 2013). It has been proposed that the dACC might be directly involved in two of the three key processes of cognitive control: the *monitoring* of conflict, volatility and the probability of undesired outcomes (Botvinik et al., 2001); and the *specification* of a control signal (Shenhav et al., 2013). Furthermore, the dACC could be indirectly involved in the third key process of cognitive control: it could deploy information about conflicts to other brain areas, such as the lateral prefrontal cortex (lPFC), which are responsible for *regulation*. A large body of evidence indicates that dACC is active in conditions requiring control intensity adjustments and influences brain structures responsible for regulation. For example, various neuroimaging studies suggest a relation between dACC activity during conflicts and the subsequent increment of activity in areas associated with attentional regulation (Cavanagh & Frank, 2014; Kerns, 2006; Kerns et al., 2004; King et al., 2010; MacDonald et al., 2000). Similar evidence comes from EEG studies, where electrophysiological indicators of dACC responses have been detected during attentional adjustments after conflict and errors (Aarts et al., 2008; Carter et al., 1998; Crottaz-Herbette & Menon, 2006; Forster et al., 2011; Shenhav et al., 2013; Sohn et al., 2007).

Following these findings, our proposal associates dACC activity with the optimization of cognitive control signals, by minimizing the expected free energy G(γ). This perspective assigns dACC a role in prioritizing (precision-weighting) simpler or more complex behavioural controllers, which are present across the dorsolateral prefrontal hierarchy (Pezzulo et al., 2015, 2018). Given its key role in monitoring, the dACC could also be implied in calculating cognitive conflict (Botvinik et al., 2001), which in our model directly influences (and is largely reflected in) the dynamics of expected free energy G(γ).

6.5.3 The dopaminergic (DA) system

Dopaminergic (DA) activity plays a crucial role in cognitive control. The high-level modulations in prefrontal areas are highly influenced by dopamine activity, particularly from midbrain dopaminergic neurons (Brozoski et al., 1979; Cools et al., 2019; Sawaguchi & Goldman-Rakic, 1991). An influential view of dopamine's role in cognitive control is that it acts as a motivational modulator that influences value-related computations, for the selection of adaptive behaviour (Cools, 2016). In other words, it modulates the value of decisions, influencing the willingness for costly control to pursue alternative ones (Aarts et al., 2008; M. Botvinick & Braver, 2015; Padmala & Pessoa, 2011; Westbrook & Braver, 2016). Along similar lines, it has been proposed that dopamine may influence multiple mechanisms for motivation (Braver et al., 2014): from stimulus-rewards Pavlovian associations (Swart et al., 2017) to cognitive cost-benefit analysis and reinforcement learning (Salamone et al., 2016; Schultz et al., 1997; Tobler et al., 2005), and putatively by promoting learning based on reward prediction errors (Montague et al., 1996).

In keeping with previous active inference works, we linked dopaminergic responses to the precision parameter γ (Friston et al., 2014; Langdon et al., 2018; Schwartenbeck, FitzGerald, Mathys, Dolan, & Friston, 2015a). Crucially, however, in this model, the precision parameter γ is optimized twice: during the specification of the control signal (Equation 5) and when new observations are gathered that change the confidence in the policies (Equation 7). Therefore, dopaminergic activity might correspond to the combination of these two components. This is illustrated in Figures 6.6-6.8, in which we plot dopaminergic activity (panels

G) as the sum of the two components (illustrated in panels D and F, respectively). These two optimizations of the precision parameter γ are not unrelated: the former mediates the selection of a "general modality" (e.g., habitual or deliberative) of control, whereas the latter mediates the selection of a specific policy and the assessment of the value of its outcomes. This perspective could help contextualize the aforementioned theories, which assign a motivational role to dopaminergic activity. This role is evident by considering that in Figures 6.6-6.8, dopaminergic activity tends to increase with control demands, suggesting that it might act as a motivational modulator that influences value-related computations (Cools, 2016). At the same time, positive and negative peaks of dopaminergic activity track at a fine-grained level the changing confidence in the current policies, when novel observations are gathered that predict their success or failure.

6.5.4 Locus coeruleus (LC) and noradrenaline

The locus coeruleus (LC) is one of the main sources of noradrenaline (NA) in the brain, and it could be a key area to process information about statistical regularities, such as surprise, uncertainty, and volatility of the environment (Behrens et al., 2007; Sales et al., 2019). In our proposal, the LC could be responsible for calculating Bayesian surprise, with two main effects. First, the LC could deploy information about Bayesian surprise to the dACC, via cortico-LC connections – as this information is essential to initiate adjustments to the level of control in the dACC. For example, high levels of Bayesian surprise lead to changes to $\mathbf{G}$ and expected free energy $G(\gamma)$. Second, a high level of Bayesian surprise could promote the release of noradrenaline, therefore increasing the learning rate and fostering faster model updating. Note that here the model learning is not a reward-based process, but a process that learns the statistical regularities of actions and outcomes (in the monitoring matrix). A complementary perspective is that noradrenaline promotes reward-based learning, too. In this perspective, there would be a hierarchy of neuromodulators: noradrenaline (which lies higher in the hierarchy) would promote the release of dopamine (which lies lower in the hierarchy), which in turn would mediate faster learning via reward prediction error (Silvetti et al., 2018).

In sum, in this section we have provided possible links between the computations of the active inference scheme – and in particular, those required to optimize a control signal – and neuronal dynamics in the dorsal anterior cingulate cortex (dACC), the dopaminergic (DA) system and the locus coeruleus (LC). Our proposals identify a coherent brain circuit that supports the various facets of cognitive control. For example, a putative process of "increase of attention to action" could start when conditions of high Bayesian surprise and cognitive conflict are detected (via interactions between LC and dACC). This, in turn, would increase the learning rate (via noradrenaline) and the control signal (in the dACC), which in turn would change the weight assigned to different behavioural controllers (in the lateral prefrontal hierarchy) – possibly, prioritizing a more deliberate mode of behaviour that engages higher hierarchical levels. Active inference provides a coherent framework to formalize all these subprocesses parsimoniously (i.e., in terms of the optimization of a single, free energy functional) and permits generating of crisp quantitative predictions, such as the ones shown in Figures 6.6-6.8. Nevertheless, various aspects of this proposal are speculative and remain to be tested in future studies.

Chapter 7

Discussion

In the introduction, I have shown how the free energy principle aims to explain what special properties things that exist must have. The principle describes things defined in a particular way and then displays the dynamics that satisfy that definition. By doing this, it can formalize the rules that govern the self-organization of any dynamic system. Self-organization implies moving through a form of itinerancy within the organism's attracting states by maintaining a low-entropy distribution over their set (Friston 2012a; Friston and Ao 2012). The explanatory power of the FEP is that all complex behaviour emerges from the simple physics described by the Fokker-Plank equation. More practically, the FEP can be seen as a mathematical rationale that can be used to develop methods and mechanics for specific applications. Indeed, while the mechanics explain what a system is doing, the principle explains why (Ramstead et al., 2022).

This translates into process theories that function as corollaries for how the free energy principle is realized. Active inference offers an account of how the process of free energy minimization is carried on in biological and artificial systems. Perception minimizes free energy by updating the Bayesian belief of a generative model in order to make it a more accurate representation of the environment, given the upcoming observations. Action minimizes free energy by changing the sensory data to make the world more compatible with our preferred prior beliefs and goals (Parr et al., 2022). This unique mathematical solution allows for the representation of cognitive processing and the underlying neural dynamics. I have shown three implementations of this in a hierarchical setting.

The first study proposed a unified computational account of the cognitive processing underlying action observation, understanding, learning and imitation abilities. W provided a computational analysis of hierarchical action representations by considering four levels of skilled movements. We have shown that the higher the level of the hidden variable, the more the representation is associated with abstract and generalized action meanings by providing four simulations of a tennis learner who observes a teacher performing tennis shots, forms hierarchical representations of the observed actions, and imitates them. These simulations show 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. Finally, we illustrate the neurobiological implications and the possible neural implementation of our proposed model, including the multiple routes for action understanding in the dorsal and ventral streams and mirror mechanisms (di Pellegrino et al., 1992; Rizzolatti & Craighero, 2004; Tessari et al., 2022).

The second study provided a unified perspective on the visuomotor system and its deficits by jointly describing them in neuropsychological, computational and neuroanatomical terms. By adopting the active inference formalization, we provided a computational account of apraxia by showing how the generative model from study 1 is able to suit the neurofunctional evidence of previous cognitive and neuroanatomical models.

We associated the generative model with the structure of a neurocognitive model and addressed how the behavioural impairment may be produced at the neurocomputational level. We provided simulations to show how the impairment can be observed both at the behavioural level, with a simulated patient that can or cannot perform the task correctly, and at the level of probabilistic beliefs. Although we acknowledge that our simple discrete models constrain the possibility for more elaborated theoretical accounts (Lanillos et al., 2020), our approach is inspired by computational phenotyping, which originated in psychiatric research (Schwartenbeck & Friston, 2016) where the computations performed by the brain are associated with measurable behaviours (Krakauer & Shadmehr, 2007; Mirza et al., 2016; Testolin & Zorzi, 2016).

In the third study, we formalize cognitive control as the optimization of a single control signal: a (precision) parameter that balances habitual and deliberate components of action selection in hierarchical active inference (Pezzulo et al., 2015, 2018). In the simulations, the agent is required to engage an optimal level of cognitive control to complete the task successfully, without wasting cognitive resources. This is achieved by weighting deliberate and habitual action components depending on the parameter that constitutes a precision estimate of beliefs about expected free energy. This precision parameter plays the same role as attentional resources (Cooper & Shallice, 2000; Shallice & Burgess, 1993) and the control signal (Shenhav et al., 2013) in cognitive theories of cognitive control: it determines the weighting of habitual and deliberate components of action selection. Our simulations illustrate the interplay of the various cognitive variables and mechanisms that have been identified as crucial for cognitive control, such as surprise and volatility, context monitoring, cognitive conflict, confidence in one’s course of actions and the specification of a control signal. All these emerge as a straightforward extension of the (hierarchical) action-perception loop of active inference and do not require heterogeneous optimisation mechanisms, as in previous proposals. In other words, the amount of cognitive control is just another thing that is optimized by minimizing expected free energy.

With these studies, I have shown how cognitive processing can be regarded as an emergent property of the process of minimizing variational free energy at different temporal scales. Setting the appropriate generative model is the only necessary step to model any cognitive task, and I showed how even pathological function can be regarded as a form of optimal inference with suboptimal prior. This hints that brain organization and neural dynamics just recapitulate the shapes of the generative models. An important future direction is to set behavioural tasks that can effectively represent cognitive processing and translate them into generative models that allow effective model fitting and the recovery of computational parameters (Mirza et al., 2016, 2019). This, in line which the idea of computational phenotyping (Schwartenbeck & Friston, 2016), allows for a more quantitive approach to otherwise more narrative aspects of behaviour by casting them in terms of, e.g., probabilistic belief (Smith, Ramstead, et al., 2022). Furthermore, given the biological plausibility of the neural process theories that accompany active inference, the same models can be accompanied by neuroimaging investigations, as shown by Schwartenbeck and colleagues (2015). In conclusion, active inference offers a way of understanding sentient behavior starting from the dynamics of self-organization.

Casting each component of cognition, from perception to planning and action, as an interrelated form of inference offers a unified perspective on adaptive systems and their biological substrates.

Acknowledgements

I'm very glad that, for the first time in something school- or work-related, the content of this thesis is mostly the result of something I actually wanted to do and was interested in. I got genuinely passionate about the topics and felt motivated to explore them, so this work feels truly "mine." This may also be due to the way the thesis was created: it is the result of three years of effort to understand a theoretical framework that I randomly discovered during the first month of my literature exploration. I then decided to make it the main topic of my project. Strangely enough, this process was reinforced by the fact that I spent two of the three years of my PhD working from home during the Covid-19 lockdowns. This opportunity allowed me to do something simple yet effective: sit and read. I started with very little understanding of everything related to active inference, but the little I learned was enough to keep me motivated to delve deeper into the topic. In doing so, I rediscovered the pleasure of learning something on my own, just by reading and trying. Then, I had the chance to spend the last year abroad in Nijmegen, where I had a real and wonderful taste of what doing science and research truly means. For this, I want to express my gratitude. First and foremost, I extend a huge thank you to my supervisor, Alessia Tessari, who gave me two very important things: autonomy and freedom. She allowed me to learn by doing and explore, with the assurance that things wouldn't lead to a dead end. I also want to thank Giovanni Pezzulo for providing me with the first exchange on the topic of active inference and for offering continuous collaboration, hints, and ideas that helped me build my knowledge and the ability to structure it properly. Floris de Lange also deserves a thank you for allowing me to spend a year at the Donders Institute in Nijmegen, which confirmed my belief that science can feel like not working when done in a certain way. Lastly, I thank my friends and colleagues at the University of Bologna, my friends at the Donders Institute, and all the friends who were close to me in Nijmegen.

References

Aarts, E., Roelofs, A., & Turennout, M. van. (2008). Anticipatory Activity in Anterior Cingulate Cortex Can Be Independent of Conflict and Error Likelihood. Journal of Neuroscience, 28(18), 4671–4678. https://doi.org/10.1523/JNEUROSCI.4400-07.2008

Abrahamse, E., Ruitenberg, M., De Kleine, E., & Verwey, W. (2013). Control of automated behavior: Insights from the discrete sequence production task. Frontiers in Human Neuroscience, 7. https://www.frontiersin.org/article/10.3389/fnhum.2013.00082

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 : Clinical, 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

Ackley, D. H., Hinton, G. E., & Sejnowski, T. J. (1985). A learning algorithm for boltzmann machines. Cognitive Science, 9(1), 147–169. https://doi.org/10.1016/S0364-0213(85)80012-4

Adams, R. A., Shipp, S., & Friston, K. J. (2013). Predictions not commands: Active inference in the motor system. Brain Structure & Function, 218(3), 611–643. https://doi.org/10.1007/s00429-012-0475-5

Aglioti, S. M., Cesari, P., Romani, M., & Urgesi, C. (2008). Action anticipation and motor resonance in elite basketball players. Nature Neuroscience, 11(9), 1109–1116. https://doi.org/10.1038/nn.2182

Anderson, J. R. (1982). Acquisition of cognitive skill. Psychological Review, 89(4), 369–406. https://doi.org/10.1037/0033-295X.89.4.369

Anderson, J. R., Betts, S., Bothell, D., Hope, R., & Lebiere, C. (2019). Learning rapid and precise skills. Psychological Review, 126(5), 727–760. https://doi.org/10.1037/rev0000152

Andrews, M. (2021). The math is not the territory: Navigating the free energy principle. Biology & Philosophy, 36(3), 30. https://doi.org/10.1007/s10539-021-09807-0

Arnold, L. (1995). Random dynamical systems. In L. Arnold, C. K. R. T. Jones, K. Mischaikow, G. Raugel, & R. Johnson (Eds.), Dynamical Systems: Lectures Given at the 2nd Session of the Centro Internazionale Matematico Estivo (C.I.M.E.) held in Montecatini Terme, Italy, June 13–22, 1994 (pp. 1–43). Springer. https://doi.org/10.1007/BFb0095238

Ashby, W. R. (1957). An introduction to cybernetics. Chapman & Hall Ltd. http://dspace.utalca.cl/handle/1950/6344

Aziz-Zadeh, L., & Damasio, A. (2008). Embodied semantics for actions: Findings from functional brain imaging. Journal of Physiology-Paris, 102(1), 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, 113(3), 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. Advances in Neural Information Processing Systems 18, 99–106.

Balleine, B. W., & Dickinson, A. (1998). Goal-directed instrumental action: Contingency and incentive learning and their cortical substrates. Neuropharmacology, 37(4), 407–419. https://doi.org/10.1016/S0028-3908(98)00033-1

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

Bardakan, M. M., Schmidt, C. C., Hesse, M. D., Fink, G. R., & Weiss, P. H. (2022). Neuropsychological Evidence for a Motor Working Memory Subsystem Related to Apraxia. Journal of Cognitive Neuroscience, 34(11), 2016–2027. https://doi.org/10.1162/jocn_a_01893

Barsalou, L. W. (2008). Grounded Cognition. Annual Review of Psychology, 59(1), 617–645. https://doi.org/10.1146/annurev.psych.59.103006.093639

Barsalou, L. W. (2014). Cognitive Psychology: An Overview for Cognitive Scientists. Psychology Press. https://doi.org/10.4324/9781315807485

Barto, A., Mirolli, M., & Baldassarre, G. (2013). Novelty or surprise? Frontiers in Psychology, 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; a Journal Devoted to the Study of the Nervous System and Behavior, 37(5), 696–699. https://doi.org/10.1016/s0010-9452(08)70617-8

Bartolo, A., & Ham, H. S. (2016). A Cognitive Overview of Limb Apraxia. Current Neurology and Neuroscience Reports, 16(8), 75. https://doi.org/10.1007/s11910-016-0675-0

Bastos, A. M., Usrey, W. M., Adams, R. A., Mangun, G. R., Fries, P., & Friston, K. J. (2012). Canonical microcircuits for predictive coding. Neuron, 76(4), 695–711. https://doi.org/10.1016/j.neuron.2012.10.038

Bayles, K. A., Kaszniak, A. W., & Tomoeda, C. K. (1987). Communication and cognition in normal aging and dementia (pp. xvi, 400). College-Hill Press/Little, Brown & Co.

Beal, M. (2003). Variational algorithms for approximate Bayesian inference /. PhD Thesis.

Behrens, T. E. J., Woolrich, M. W., Walton, M. E., & Rushworth, M. F. S. (2007). Learning the value of information in an uncertain world. Nature Neuroscience, 10(9), Article 9. https://doi.org/10.1038/nn1954

Bengio, Y. (2011). Deep learning of representations for unsupervised and transfer learning. Proceedings of the 2011 International Conference on Unsupervised and Transfer Learning Workshop - Volume 27, 17–37.

Berlyne, D. E. (1954). A Theory of Human Curiosity. British Journal of Psychology. General Section, 45(3), 180–191. https://doi.org/10.1111/j.2044-8295.1954.tb01243.x

Berlyne, D. E. (1960). Conflict, arousal, and curiosity (pp. xii, 350). McGraw-Hill Book Company. https://doi.org/10.1037/11164-000

Bi, Y., Wang, X., & Caramazza, A. (2016). Object Domain and Modality in the Ventral Visual Pathway. Trends in Cognitive Sciences, 20(4), 282–290. https://doi.org/10.1016/j.tics.2016.02.002

Binkofski, F., & Buxbaum, L. J. (2013). Two action systems in the human brain. Brain and Language, 127(2), 222–229. https://doi.org/10.1016/j.bandl.2012.07.007

Blake, A., & Yuille, A. L. (Eds.). (1992). Active Vision. MIT Press.

Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent dirichlet allocation. The Journal of Machine Learning Research, 3(null), 993–1022.

Bonini, L. (2016). The Extended Mirror Neuron Network Anatomy, Origin, and Functions. The 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 A. M. Borghi & F. Binkofski (Eds.), Words as Social Tools: An Embodied View on Abstract Concepts (pp. 19–37). Springer. 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(3), 763–773. https://doi.org/10.1016/j.neuropsychologia.2009.10.029

Botvinick, M., & Braver, T. (2015). Motivation and cognitive control: From behavior to neural mechanism. Annual Review of Psychology, 66, 83–113. https://doi.org/10.1146/annurev-psych-010814-015044

Botvinick, M. M. (2008). Hierarchical models of behavior and prefrontal function. Trends in Cognitive Sciences, 12(5), 201–208. https://doi.org/10.1016/j.tics.2008.02.009

Botvinick, M. M., Braver, T. S., Barch, D. M., Carter, C. S., & Cohen, J. D. (2001). Conflict monitoring and cognitive control. Psychological Review, 108(3), 624–652. https://doi.org/10.1037/0033-295x.108.3.624

Botvinick, M. M., Cohen, J. D., & Carter, C. S. (2004). Conflict monitoring and anterior cingulate cortex: An update. Trends in Cognitive Sciences, 8(12), 539–546. https://doi.org/10.1016/j.tics.2004.10.003

Botvinick, M. M., Huffstetler, S., & McGuire, J. T. (2009). Effort discounting in human nucleus accumbens. Cognitive, Affective & Behavioral Neuroscience, 9(1), 16–27. https://doi.org/10.3758/CABN.9.1.16

Botvinick, M. M., Niv, Y., & Barto, A. G. (2009). Hierarchically organized behavior and its neural foundations: A reinforcement learning perspective. Cognition, 113(3), 262–280. https://doi.org/10.1016/j.cognition.2008.08.011

Botvinick, M., Ritter, S., Wang, J. X., Kurth-Nelson, Z., Blundell, C., & Hassabis, D. (2019). Reinforcement Learning, Fast and Slow. Trends in Cognitive Sciences, 23(5), 408–422. https://doi.org/10.1016/j.tics.2019.02.006

Botvinick, M., & Toussaint, M. (2012). Planning as inference. Trends in Cognitive Sciences, 16(10), 485–488. https://doi.org/10.1016/j.tics.2012.08.006

Bracci, S., Daniels, N., & Op de Beeck, H. (2017). Task Context Overrules Object- and Category-Related Representational Content in the Human Parietal Cortex. Cerebral Cortex, cercor;bhw419v1. https://doi.org/10.1093/cercor/bhw419

Brass, M., Bekkering, H., & Prinz, W. (2001). Movement observation affects movement execution in a simple response task. Acta Psychologica, 106(1–2), 3–22. https://doi.org/10.1016/s0001-6918(00)00024-x

Brass, M., & Heyes, C. (2005). Imitation: Is cognitive neuroscience solving the correspondence problem? Trends in Cognitive Sciences, 9(10), 489–495. https://doi.org/10.1016/j.tics.2005.08.007

Braver, T. S., Krug, M. K., Chiew, K. S., Kool, W., Westbrook, J. A., Clement, N. J., Adcock, R. A., Barch, D. M., Botvinick, M. M., Carver, C. S., Cools, R., Custers, R., Dickinson, A., Dweck, C. S., Fishbach, A., Gollwitzer, P. M., Hess, T. M., Isaacowitz, D. M., Mather, M., … MOMCAI group. (2014). Mechanisms of motivation-cognition interaction: Challenges and opportunities. Cognitive, Affective & Behavioral Neuroscience, 14(2), 443–472. https://doi.org/10.3758/s13415-014-0300-0

Brown, J. W., & Braver, T. S. (2005). Learned predictions of error likelihood in the anterior cingulate cortex. Science (New York, N.Y.), 307(5712), 1118–1121. https://doi.org/10.1126/science.1105783

Brozoski, T. J., Brown, R. M., Rosvold, H. E., & Goldman, P. S. (1979). Cognitive deficit caused by regional depletion of dopamine in prefrontal cortex of rhesus monkey. Science, 205(4409), 929–931. https://doi.org/10.1126/science.112679

Buxbaum, L. J. (2001). Ideomotor apraxia: A call to action. Neurocase, 7(6), 445–458. https://doi.org/10.1093/neucas/7.6.445

Buxbaum, L. J., Kyle, K. M., & Menon, R. (2005). On beyond mirror neurons: Internal representations subserving imitation and recognition of skilled object-related actions in humans. Cognitive Brain Research, 25(1), 226–239. https://doi.org/10.1016/j.cogbrainres.2005.05.014

Buxbaum, L. J., & Randerath, J. (2018). Limb apraxia and the left parietal lobe. Handbook of Clinical Neurology, 151, 349–363. https://doi.org/10.1016/B978-0-444-63622-5.00017-6

Buxbaum, L. J., Schwartz, M. F., & Carew, T. G. (1997). The role of semantic memory in object use. Cognitive Neuropsychology, 14, 219–254. https://doi.org/10.1080/026432997381565

Buxbaum, L. J., Veramonti, T., & Schwartz, M. F. (2000). Function and manipulation tool knowledge in apraxia: Knowing “what for” but not “how.” Neurocase, 6, 83–97. https://doi.org/10.1093/neucas/6.2.83

Buxbaum, L., & Kalénine, S. (2010). Action knowledge, visuomotor activation, and embodiment in the two action systems. Annals of the New York Academy of Sciences, 1191, 201–218. https://doi.org/10.1111/j.1749-6632.2010.05447.x

Byrne, R. W., & Russon, A. E. (1998). Learning by imitation: A hierarchical approach. Behavioral and Brain Sciences, 21(5), 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. Human Brain Mapping, 39(10), 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. Cerebral Cortex (New York, N.Y.: 1991), 15(8), 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. Current Biology: CB, 16(19), 1905–1910. https://doi.org/10.1016/j.cub.2006.07.065

Carter, C. S., Braver, T. S., Barch, D. M., Botvinick, M. M., Noll, D., & Cohen, J. D. (1998). Anterior cingulate cortex, error detection, and the online monitoring of performance. Science (New York, N.Y.), 280(5364), 747–749. https://doi.org/10.1126/science.280.5364.747

Catani, M., & ffytche, D. H. (2005). The rises and falls of disconnection syndromes. Brain, 128(10), 2224–2239. https://doi.org/10.1093/brain/awh622

Catani, M., Howard, R. J., Pajevic, S., & Jones, D. K. (2002). Virtual in vivo interactive dissection of white matter fasciculi in the human brain. NeuroImage, 17(1), 77–94. https://doi.org/10.1006/nimg.2002.1136

Catani, M., Jones, D. K., & ffytche, D. H. (2005). Perisylvian language networks of the human brain. Annals of Neurology, 57(1), 8–16. https://doi.org/10.1002/ana.20319

Catmur, C., Walsh, V., & Heyes, C. (2009). Associative sequence learning: The role of experience in the development of imitation and the mirror system. Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences, 364(1528), 2369–2380. https://doi.org/10.1098/rstb.2009.0048

Cavanagh, J. F., Eisenberg, I., Guitart-Masip, M., Huys, Q., & Frank, M. J. (2013). Frontal Theta Overrides Pavlovian Learning Biases. Journal of Neuroscience, 33(19), 8541–8548. https://doi.org/10.1523/JNEUROSCI.5754-12.2013

Cavanagh, J. F., & Frank, M. J. (2014). Frontal theta as a mechanism for cognitive control. Trends in Cognitive Sciences, 18(8), 414–421. https://doi.org/10.1016/j.tics.2014.04.012

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

Clark, A. (2013b). Whatever next? Predictive brains, situated agents, and the future of cognitive science. Behavioral and Brain Sciences, 36(3), 181–204. https://doi.org/10.1017/S0140525X12000477

Cohen, J. Y., Haesler, S., Vong, L., Lowell, B. B., & Uchida, N. (2012). Neuron-type-specific signals for reward and punishment in the ventral tegmental area. Nature, 482(7383), 85–88. https://doi.org/10.1038/nature10754

Collins, R. P., Litman, J. A., & Spielberger, C. D. (2004). The measurement of perceptual curiosity. Personality and Individual Differences, 36(5), 1127–1141. https://doi.org/10.1016/S0191-8869(03)00205-8

Coltheart, M., Masterson, J., Byng, S., Prior, M., & Riddoch, J. (1983). Surface dyslexia. The Quarterly Journal of Experimental Psychology Section A, 35(3), 469–495. https://doi.org/10.1080/14640748308402483

Cools, R. (2016). The costs and benefits of brain dopamine for cognitive control. WIREs Cognitive Science, 7(5), 317–329. https://doi.org/10.1002/wcs.1401

Cools, R., Froböse, M., Aarts, E., & Hofmans, L. (2019). Dopamine and the motivation of cognitive control. Handbook of Clinical Neurology, 163, 123–143. https://doi.org/10.1016/B978-0-12-804281-6.00007-0

Cooper, R., & Shallice, T. (2000). Contention scheduling and the control of routine activities. Cognitive Neuropsychology, 17(4), 297–338. https://doi.org/10.1080/026432900380427

Corcoran, A. W., Pezzulo, G., & Hohwy, J. (2020). From Allostatic Agents to Counterfactual Cognisers: Active Inference, Biological Regulation, and the Origins of Cognition. Biology and Philosophy, 35(3), 1–45. https://doi.org/10.1007/s10539-020-09746-2

Coslett, H. B. (1991). Read but not write “Idea”: Evidence for a third reading mechanism. Brain and Language, 40(4), 425–443. https://doi.org/10.1016/0093-934X(91)90141-M

Courville, A. C., Daw, N. D., & Touretzky, D. S. (2006). Bayesian theories of conditioning in a changing world. Trends in Cognitive Sciences, 10(7), 294–300. https://doi.org/10.1016/j.tics.2006.05.004

Crauel, H., & Flandoli, F. (1994). Attractors for random dynamical systems. Probability Theory and Related Fields, 100(3), 365–393. https://doi.org/10.1007/BF01193705

Creem-Regehr, S. H., Dilda, V., Vicchrilli, A. E., Federer, F., & Lee, J. N. (2007). The influence of complex action knowledge on representations of novel graspable objects: Evidence from functional magnetic resonance imaging. Journal of the International Neuropsychological Society, 13(6), 1009–1020. https://doi.org/10.1017/S1355617707071093

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(3), 1257–1267. https://doi.org/10.1016/j.neuroimage.2006.01.033

Crottaz-Herbette, S., & Menon, V. (2006). Where and when the anterior cingulate cortex modulates attentional response: Combined fMRI and ERP evidence. Journal of Cognitive Neuroscience, 18(5), 766–780. https://doi.org/10.1162/jocn.2006.18.5.766

Csibra, G., & Gergely, G. (2007). ‘Obsessed with goals’: Functions and mechanisms of teleological interpretation of actions in humans. Acta Psychologica, 124(1), 60–78. https://doi.org/10.1016/j.actpsy.2006.09.007

Cubelli, R., Bartolo, A., Nichelli, P., & Della Sala, S. (2006). List effect in apraxia assessment. Neuroscience Letters, 407(2), 118–120. https://doi.org/10.1016/j.neulet.2006.08.019

Cubelli, R., Marchetti, C., Boscolo, G., & Della Sala, S. (2000). Cognition in action: Testing a model of limb apraxia. Brain and Cognition, 44(2), 144–165. https://doi.org/10.1006/brcg.2000.1226

Da Costa, L., Parr, T., Sajid, N., Veselic, S., Neacsu, V., & Friston, K. (2020). Active inference on discrete state-spaces: A synthesis. Journal of Mathematical Psychology, 99, 102447. https://doi.org/10.1016/j.jmp.2020.102447

Da Costa, L., Parr, T., Sengupta, B., & Friston, K. (2021). Neural Dynamics under Active Inference: Plausibility and Efficiency of Information Processing. Entropy (Basel, Switzerland), 23(4), 454. https://doi.org/10.3390/e23040454

D’Ardenne, K., McClure, S. M., Nystrom, L. E., & Cohen, J. D. (2008). BOLD responses reflecting dopaminergic signals in the human ventral tegmental area. Science (New York, N.Y.), 319(5867), 1264–1267. https://doi.org/10.1126/science.1150605

Davis, C. P., & Yee, E. (2021). Building semantic memory from embodied and distributional language experience. WIREs Cognitive Science, 12(5), e1555. https://doi.org/10.1002/wcs.1555

Daw, N. D., Gershman, S. J., Seymour, B., Dayan, P., & Dolan, R. J. (2011). Model-based influences on humans’ choices and striatal prediction errors. Neuron, 69(6), 1204–1215. https://doi.org/10.1016/j.neuron.2011.02.027

Daw, N. D., Niv, Y., & Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. Nature Neuroscience, 8(12), 1704–1711. https://doi.org/10.1038/nn1560

Dayan, P., Hinton, G. E., Neal, R. M., & Zemel, R. S. (1995). The Helmholtz machine. Neural Computation, 7(5), 889–904. https://doi.org/10.1162/neco.1995.7.5.889

De Martino, B., Kumaran, D., Seymour, B., & Dolan, R. J. (2006). Frames, Biases, and Rational Decision-Making in the Human Brain. Science (New York, N.Y.), 313(5787), 684–687. https://doi.org/10.1126/science.1128356

De Renzi, E., Faglioni, P., Scarpa, M., & Crisi, G. (1986). Limb apraxia in patients with damage confined to the left basal ganglia and thalamus. Journal of Neurology, Neurosurgery, and Psychiatry, 49(9), 1030–1038. https://doi.org/10.1136/jnnp.49.9.1030

De Renzi, E., Motti, F., & Nichelli, P. (1980). Imitating gestures. A quantitative approach to ideomotor apraxia. Archives of Neurology, 37(1), 6–10. https://doi.org/10.1001/archneur.1980.00500500036003

Dezfouli, A., & Balleine, B. W. (2012). Habits, action sequences and reinforcement learning. European Journal of Neuroscience, 35(7), 1036–1051. https://doi.org/10.1111/j.1460-9568.2012.08050.x

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(12), 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. Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences, 369(1655), 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. Experimental Brain Research, 91(1), 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. Proceedings of IJCAI 2011, 2113–2119.

Dolan, R. J., & Dayan, P. (2013). Goals and habits in the brain. Neuron, 80(2), 312–325. https://doi.org/10.1016/j.neuron.2013.09.007

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

Doya, K. (2002). Metalearning and neuromodulation. Neural Netw, 15(4–6), 495–506.

Dressing, A., Nitschke, K., Kümmerer, D., Bormann, T., Beume, L., Schmidt, C. S. M., Ludwig, V. M., Mader, I., Willmes, K., Rijntjes, M., Kaller, C. P., Weiller, C., & Martin, M. (2018). Distinct Contributions of Dorsal and Ventral Streams to Imitation of Tool-Use and Communicative Gestures. Cerebral Cortex, 28(2), 474–492. https://doi.org/10.1093/cercor/bhw383

Edelman, G. M., & Gally, J. A. (2001). Degeneracy and complexity in biological systems. Proceedings of the National Academy of Sciences of the United States of America, 98(24), 13763–13768. https://doi.org/10.1073/pnas.231499798

Engel, A. K., Friston, K. J., & Kragic, D. (2016). The Pragmatic Turn: Toward Action-Oriented Views in Cognitive Science. MIT Press.

Engel, A. K., Maye, A., Kurthen, M., & König, P. (2013). Where's the action? The pragmatic turn in cognitive science. Trends in Cognitive Sciences, 17(5), 202–209. https://doi.org/10.1016/j.tics.2013.03.006

Engelkamp, J., & Zimmer, H. D. (1984). Motor programme information as a separable memory unit. Psychological Research, 46(3), 283–299. https://doi.org/10.1007/BF00308889

Engelkamp, J., & Zimmer, H. D. (1985). Motor programs and their relation to semantic memory. German Journal of Psychology, 9(3), 239–254.

Engelkamp, J., & Zimmer, H. D. (1989). Memory for action events: A new field of research. Psychological Research, 51(4), 153–157. https://doi.org/10.1007/BF00309142

Evans, D. J., & Searles, D. J. (2002). The Fluctuation Theorem. Advances in Physics, 51(7), 1529–1585. https://doi.org/10.1080/00018730210155133

Fasoli, D. (2013). Attacking the brain with neuroscience: Mean-field theory, finite size effects and encoding capability of stochastic neural networks [Phdthesis, Université Nice Sophia Antipolis]. https://tel.archives-ouvertes.fr/tel-00850289

Feldman, H., & Friston, K. (2010). Attention, Uncertainty, and Free-Energy. Frontiers in Human Neuroscience, 4, 215. https://doi.org/10.3389/fnhum.2010.00215

Feynman, R. P. (1972). Statistical mechanics: A set of lectures. Benjamin/Cummings.

Finkel, L., Hogrefe, K., Frey, S. H., Goldenberg, G., & Randerath, J. (2018). It takes two to pantomime: Communication meets motor cognition. NeuroImage: Clinical, 19, 1008–1017. https://doi.org/10.1016/j.nicl.2018.06.019

FitzGerald, T. H. B., Dolan, R. J., & Friston, K. (2015). Dopamine, reward learning, and active inference. Frontiers in Computational Neuroscience, 9, 136. https://doi.org/10.3389/fncom.2015.00136

FitzGerald, T. H., Dolan, R. J., & Friston, K. (2015). Dopamine, reward learning, and active inference. Frontiers in Computational Neuroscience, 9. http://www.ncbi.nlm.nih.gov/pmc/articles/PMC4631836/

Flagel, S. B., Clark, J. J., Robinson, T. E., Mayo, L., Czuj, A., Willuhn, I., Akers, C. A., Clinton, S. M., Phillips, P. E. M., & Akil, H. (2011). A selective role for dopamine in reward learning. Nature, 469(7328), 53–57. https://doi.org/10.1038/nature09588

Flanagan, J. R., & Johansson, R. S. (2003). Action plans used in action observation. Nature, 424(6950), 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(5722), 662–667. https://doi.org/10.1126/science.1106138

Forster, S. E., Carter, C. S., Cohen, J. D., & Cho, R. Y. (2011). Parametric manipulation of the conflict signal and control-state adaptation. Journal of Cognitive Neuroscience, 23(4), 923–935. https://doi.org/10.1162/jocn.2010.21458

Friston, K. (2005). A theory of cortical responses. Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences, 360(1456), 815–836. https://doi.org/10.1098/rstb.2005.1622

Friston, K. (2008). Hierarchical Models in the Brain. PLOS Computational Biology, 4(11), e1000211. https://doi.org/10.1371/journal.pcbi.1000211

Friston, K. (2010). The free-energy principle: A unified brain theory? Nature Reviews. Neuroscience, 11(2), 127–138. https://doi.org/10.1038/nrn2787

Friston, K. (2019). A free energy principle for a particular physics (arXiv:1906.10184). arXiv. https://doi.org/10.48550/arXiv.1906.10184

Friston, K. (2021). World model learning and inference. Neural Networks, 144, 573–590.

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., & Ao, P. (2012). Free energy, value, and attractors. Computational and Mathematical Methods in Medicine, 2012, 937860. https://doi.org/10.1155/2012/937860

Friston, K., & Buzsáki, G. (2016). The Functional Anatomy of Time: What and When in the Brain. Trends in Cognitive Sciences, 20(7), 500–511. https://doi.org/10.1016/j.tics.2016.05.001

Friston, K., Da Costa, L., Sajid, N., Heins, C., Ueltzhöffer, K., Pavliotis, G. A., & Parr, T. (2022). The free energy principle made simpler but not too simple (arXiv:2201.06387). arXiv. https://doi.org/10.48550/arXiv.2201.06387

Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., O'Doherty, J. P., & Pezzulo, G. (2016a). 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., O'Doherty, J. P., & Pezzulo, G. (2016b). 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. (2017a). Active Inference: A Process Theory. Neural Computation, 29(1), 1–49. https://doi.org/10.1162/NECO_a_00912

Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., & Pezzulo, G. (2017b). Active Inference: A Process Theory. Neural Computation, 29(1), 1–49. https://doi.org/10.1162/NECO_a_00912

Friston, K., & Frith, C. (2015). A Duet for one. Consciousness and Cognition, 36, 390–405. https://doi.org/10.1016/j.concog.2014.12.003

Friston, K., Heins, C., Ueltzhöffer, K., Da Costa, L., & Parr, T. (2021). Stochastic Chaos and Markov Blankets. Entropy, 23(9), Article 9. https://doi.org/10.3390/e23091220

Friston, K. J., Lin, M., Frith, C. D., Pezzulo, G., Hobson, J. A., & Ondobaka, S. (2017). Active Inference, Curiosity and Insight. Neural Computation, 29(10), 2633–2683. https://doi.org/10.1162/neco_a_00999

Friston, K. J., & Price, C. J. (2003). Degeneracy and redundancy in cognitive anatomy. Trends in Cognitive Sciences, 7(4), 151–152. https://doi.org/10.1016/S1364-6613(03)00054-8

Friston, K., & Kiebel, S. (2009). Predictive coding under the free-energy principle. Philosophical Transactions of the Royal Society B: Biological Sciences, 364(1521), 1211–1221. https://doi.org/10.1098/rstb.2008.0300

Friston, K., Kilner, J., & Harrison, L. (2006). A free energy principle for the brain. Journal of Physiology, Paris, 100(1–3), 70–87. https://doi.org/10.1016/j.jphysparis.2006.10.001

Friston, K., Mattout, J., & Kilner, J. (2011). Action understanding and active inference. Biological Cybernetics, 104(1–2), 137–160. https://doi.org/10.1007/s00422-011-0424-z

Friston, K., Rigoli, F., Ognibene, D., Mathys, C., Fitzgerald, T., & Pezzulo, G. (2015). Active inference and epistemic value. Cognitive Neuroscience, 6(4), 187–214. https://doi.org/10.1080/17588928.2015.1020053

Friston, K., Rosch, R., Parr, T., Price, C., & Bowman, H. (2017). Deep temporal models and active inference. Neuroscience and Biobehavioral Reviews, 77, 388–402. https://doi.org/10.1016/j.neubiorev.2017.04.009

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., Samothrakis, S., & Montague, R. (2012). Active inference and agency: Optimal control without cost functions. Biological Cybernetics, 106(8–9), 523–541. https://doi.org/10.1007/s00422-012-0512-8

Friston, K., Schwartenbeck, P., Fitzgerald, T., Moutoussis, M., Behrens, T., & Dolan, R. (2013). The anatomy of choice: Active inference and agency. Frontiers in Human Neuroscience, 7. https://www.frontiersin.org/article/10.3389/fnhum.2013.00598

Friston, K., Schwartenbeck, P., FitzGerald, T., Moutoussis, M., Behrens, T., & Dolan, R. J. (2014). The anatomy of choice: Dopamine and decision-making. Philosophical Transactions of the Royal Society B: Biological Sciences, 369(1655), 20130481. https://doi.org/10.1098/rstb.2013.0481

Friston, K., Shiner, T., FitzGerald, T., Galea, J. M., Adams, R., Brown, H., Dolan, R. J., Moran, R., Stephan, K. E., & Bestmann, S. (2012). Dopamine, Affordance and Active Inference. PLoS Computational Biology, 8(1), e1002327. https://doi.org/10.1371/journal.pcbi.1002327

Gabriel, M., & Orona, E. (1982). Parallel and serial processes of the prefrontal and cingulate cortical systems during behavioral learning. Brain Research Bulletin, 8(6), 781–785. https://doi.org/10.1016/0361-9230(82)90107-1

Galetzka, C. (2017). The Story So Far: How Embodied Cognition Advances Our Understanding of Meaning-Making. Frontiers in Psychology, 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(12), 493–501. https://doi.org/10.1016/S1364-6613(98)01262-5

Galvez-Pol, A., Forster, B., & Calvo-Merino, B. (2020). Beyond action observation: Neurobehavioral mechanisms of memory for visually perceived bodies and actions. Neuroscience and Biobehavioral Reviews, 116, 508–518. https://doi.org/10.1016/j.neubiorev.2020.06.014

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(6), 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(7), 287–292. https://doi.org/10.1016/S1364-6613(03)00128-1

Gershman, S. J. (2017). Dopamine, Inference, and Uncertainty. Neural Computation, 29(12), 3311–3326. https://doi.org/10.1162/neco_a_01023

Glimcher, P. W. (2011). Understanding dopamine and reinforcement learning: The dopamine reward prediction error hypothesis. Proceedings of the National Academy of Sciences, 108(supplement_3), 15647–15654. https://doi.org/10.1073/pnas.1014269108

Goldenberg, G. (1995). Imitating gestures and manipulating a mannikin: The representations of the human body in ideomotor apraxia. Neuropsychologia, 33, 63–72. https://doi.org/10.1016/0028-3932(94)00104-W

Goldenberg, G. (2014). Apraxia – The cognitive side of motor control. Cortex, 57, 270–274. https://doi.org/10.1016/j.cortex.2013.07.016

Goldenberg, G., & Hagmann, S. (1997). The meaning of meaningless gestures: A study of visuo-imitative apraxia. Neuropsychologia, 35(3), 333–341. https://doi.org/10.1016/S0028-3932(96)00085-1

Goldenberg, G., & Hagmann, S. (1998). Tool use and mechanical problem solving in apraxia. Neuropsychologia, 36(7), 581–589. https://doi.org/10.1016/S0028-3932(97)00165-6

Goldenberg, G., Hermsdörfer, J., & Spatt, J. (1996). Ideomotor apraxia and cerebral dominance for motor control. Cognitive Brain Research, 3(2), 95–100. https://doi.org/10.1016/0926-6410(95)00034-8

Gómez, C. M., Arjona, A., Donnarumma, F., Maisto, D., Rodríguez-Martínez, E. I., & Pezzulo, G. (2019). Tracking the Time Course of Bayesian Inference With Event-Related Potentials:A Study Using the Central Cue Posner Paradigm. Frontiers in Psychology, 10. https://www.frontiersin.org/articles/10.3389/fpsyg.2019.01424

Goodale, M. A., & Milner, A. D. (1992). Separate visual pathways for perception and action. Trends in Neurosciences, 15(1), 20–25. https://doi.org/10.1016/0166-2236(92)90344-8

Gottlieb, J., & Oudeyer, P.-Y. (2018). Towards a neuroscience of active sampling and curiosity. Nature Reviews Neuroscience, 19(12), 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. Annals of the New York Academy of Sciences, 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, 26(4), 590–616. https://doi.org/10.1016/j.humov.2007.05.009

Gregory, R. L., Longuet-Higgins, H. C., & Sutherland, N. S. (1980). Perceptions as hypotheses. Philosophical Transactions of the Royal Society of London. B, Biological Sciences, 290(1038), 181–197. https://doi.org/10.1098/rstb.1980.0090

Hagura, N., Takei, T., Hirose, S., Aramaki, Y., Matsumura, M., Sadato, N., & Naito, E. (2007). Activity in the Posterior Parietal Cortex Mediates Visual Dominance over Kinesthesia. Journal of Neuroscience, 27(26), 7047–7053. https://doi.org/10.1523/JNEUROSCI.0970-07.2007

Haken, H. (2004). An Introduction. In H. Haken (Ed.), Synergetics: Introduction and Advanced Topics (pp. 1–387). Springer. https://doi.org/10.1007/978-3-662-10184-1_1

Hamzei, F., Vry, M.-S., Saur, D., Glauche, V., Hoeren, M., Mader, I., Weiller, C., & Rijntjes, M. (2016). The Dual-Loop Model and the Human Mirror Neuron System: An Exploratory Combined fMRI and DTI Study of the Inferior Frontal Gyrus. Cerebral Cortex, 26(5), 2215–2224. https://doi.org/10.1093/cercor/bhv066

Hasson, U., Chen, J., & Honey, C. J. (2015). Hierarchical process memory: Memory as an integral component of information processing. Trends in Cognitive Sciences, 19(6), 304–313. https://doi.org/10.1016/j.tics.2015.04.006

Hauk, O., Johnsrude, I., & Pulvermüller, F. (2004). Somatotopic Representation of Action Words in Human Motor and Premotor Cortex. Neuron, 41(2), 301–307. https://doi.org/10.1016/S0896-6273(03)00838-9

Heilman, K. M., Maher, L. M., Greenwald, M. L., & Rothi, L. J. G. (1997). Conceptual apraxia from lateralized lesions. Neurology, 49(2), 457–464. https://doi.org/10.1212/WNL.49.2.457

Heilman, K. M., & Rothi, L. J. G. (1993). Apraxia. In Clinical neuropsychology, 3rd ed (pp. 141–163). Oxford University Press.

Heilman, K. M., & Rothi, L. J. G. (2003). Apraxia. In Clinical neuropsychology, 4th ed (pp. 215–235). Oxford University Press.

Heilman, K. M., Rothi, L. J., & Valenstein, E. (1982). Two forms of ideomotor apraxia. Neurology, 32(4), 342–342. https://doi.org/10.1212/WNL.32.4.342

Hermsdörfer, J., Li, Y., Randerath, J., Goldenberg, G., & Johannsen, L. (2012). Tool use without a tool: Kinematic characteristics of pantomiming as compared to actual use and the effect of brain damage. Experimental Brain Research, 218(2), 201–214. https://doi.org/10.1007/s00221-012-3021-z

Hermsdörfer, J., Li, Y., Randerath, J., Roby-Brami, A., & Goldenberg, G. (2013). Tool use kinematics across different modes of execution. Implications for action representation and apraxia. Cortex; a Journal Devoted to the Study of the Nervous System and Behavior, 49(1), 184–199. https://doi.org/10.1016/j.cortex.2011.10.010

Hesp, C., Ramstead, M., Constant, A., Badcock, P., Kirchhoff, M., & Friston, K. (2019). A Multi-scale View of the Emergent Complexity of Life: A Free-Energy Proposal. In G. Y. Georgiev, J. M. Smart, C. L. Flores Martinez, & M. E. Price (Eds.), Evolution, Development and Complexity (pp. 195–227). Springer International Publishing. https://doi.org/10.1007/978-3-030-00075-2_7

Hesp, C., Smith, R., Parr, T., Allen, M., Friston, K. J., & Ramstead, M. J. D. (2021). Deeply Felt Affect: The Emergence of Valence in Deep Active Inference. Neural Computation, 33(2), 398–446. https://doi.org/10.1162/neco_a_01341

Heugten, C. M. van. (1998). Apraxia in stroke patients: Assessment and treatment [Thesis fully internal (DIV)]. s.n.

Heyes, C. (2011). Automatic imitation. Psychological Bulletin, 137(3), 463–483. https://doi.org/10.1037/a0022288

Hobson, J. A., & Friston, K. J. (2012). Waking and dreaming consciousness: Neurobiological and functional considerations. Progress in Neurobiology, 98(1), 82–98. https://doi.org/10.1016/j.pneurobio.2012.05.003

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: A Journal of Neurology, 137(Pt 10), 2796–2810. https://doi.org/10.1093/brain/awu203

Hohwy, J. (2016). The Self-Evidencing Brain. Noûs, 50(2), 259–285. https://doi.org/10.1111/nous.12062

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(5), 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. A Bradford Book.

Iacoboni, M., Woods, R. P., Brass, M., Bekkering, H., Mazziotta, J. C., & Rizzolatti, G. (1999). Cortical mechanisms of human imitation. Science (New York, N.Y.), 286(5449), 2526–2528. https://doi.org/10.1126/science.286.5449.2526

Isomura, T. (2021). Active inference leads to Bayesian neurophysiology. Neuroscience Research. https://doi.org/10.1016/j.neures.2021.12.003

Itti, L., & Baldi, P. (2009). Bayesian surprise attracts human attention. Vision Research, 49(10), 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(10), 1489–1506. 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(11), 1254–1259. https://doi.org/10.1109/34.730558

Izawa, J., Rane, T., Donchin, O., & Shadmehr, R. (2008). Motor Adaptation as a Process of Reoptimization. Journal of Neuroscience, 28(11), 2883–2891. https://doi.org/10.1523/JNEUROSCI.5359-07.2008

Jacob, P., & Jeannerod, M. (2005). The motor theory of social cognition: A critique. Trends in Cognitive Sciences, 9(1), 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(6). https://doi.org/10.1371/journal.pone.0039629

James, W. (1890). The principles of psychology, Vol I. (pp. xii, 697). Henry Holt and Co. 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(2), 350–355. https://doi.org/10.1016/j.cognition.2010.01.004

Jeannerod, M. (1994). The representing brain: Neural correlates of motor intention and imagery. Behavioral and Brain Sciences, 17(2), 187–245. https://doi.org/10.1017/S0140525X00034026

Jeannerod, M. (2001). Neural simulation of action: A unifying mechanism for motor cognition. NeuroImage, 14(1 Pt 2), S103-109. https://doi.org/10.1006/nimg.2001.0832

Jeannerod, M. (2006). Motor cognition: What actions tell the self (pp. x, 209). Oxford University Press. https://doi.org/10.1093/acprof:oso/9780198569657.001.0001

Jeannerod, M., Arbib, M. A., Rizzolatti, G., & Sakata, H. (1995). Grasping objects: The cortical mechanisms of visuomotor transformation. Trends in Neurosciences, 18(7), 314–320.

Jimura, K., Locke, H. S., & Braver, T. S. (2010). Prefrontal cortex mediation of cognitive enhancement in rewarding motivational contexts. Proceedings of the National Academy of Sciences of the United States of America, 107(19), 8871–8876. https://doi.org/10.1073/pnas.1002007107

Kahneman, D. (2011). Thinking, Fast and Slow. Penguin UK.

Kahneman, D., Treisman, A., & Burkell, J. (1983). The cost of visual filtering. Journal of Experimental Psychology: Human Perception and Performance, 9(4), 510–522. https://doi.org/10.1037/0096-1523.9.4.510

Kalénine, S., Buxbaum, L. J., & Coslett, H. B. (2010). Critical brain regions for action recognition: Lesion symptom mapping in left hemisphere stroke. Brain: A Journal of Neurology, 133(11), 3269–3280. https://doi.org/10.1093/brain/awq210

Kauffman, S. A., & Kauffman, M. of the S. F. I. and P. of B. S. A. (1993). The Origins of Order: Self-organization and Selection in Evolution. Oxford University Press.

Keramati, M., Dezfouli, A., & Piray, P. (2011). Speed/Accuracy Trade-Off between the Habitual and the Goal-Directed Processes. PLOS Computational Biology, 7(5), e1002055. https://doi.org/10.1371/journal.pcbi.1002055

Kerns, J. G. (2006). Anterior cingulate and prefrontal cortex activity in an FMRI study of trial-to-trial adjustments on the Simon task. NeuroImage, 33(1), 399–405. https://doi.org/10.1016/j.neuroimage.2006.06.012

Kerns, J. G., Cohen, J. D., MacDonald, A. W., Cho, R. Y., Stenger, V. A., & Carter, C. S. (2004). Anterior cingulate conflict monitoring and adjustments in control. Science (New York, N.Y.), 303(5660), 1023–1026. https://doi.org/10.1126/science.1089910

Kiebel, S., Daunizeau, J., & Friston, K. (2008). A Hierarchy of Time-Scales and the Brain. PLOS Computational Biology, 4(11), e1000209. https://doi.org/10.1371/journal.pcbi.1000209

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

Kilner, J. M. (2011). More than one pathway to action understanding. Trends in Cognitive Sciences, 15(8), 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. Cognitive Processing, 8(3), 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. Current Biology: CB, 13(6), 522–525. https://doi.org/10.1016/s0960-9822(03)00165-9

King, J. A., Korb, F. M., von Cramon, D. Y., & Ullsperger, M. (2010). Post-error behavioral adjustments are facilitated by activation and suppression of task-relevant and task-irrelevant information processing. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience, 30(38), 12759–12769. https://doi.org/10.1523/JNEUROSCI.3274-10.2010

Kirchhoff, M., Parr, T., Palacios, E., Friston, K., & Kiverstein, J. (2018). The Markov blankets of life: Autonomy, active inference and the free energy principle. Journal of The Royal Society Interface, 15(138), 20170792. https://doi.org/10.1098/rsif.2017.0792

Klein, E., Moeller, K., Glauche, V., Weiller, C., & Willmes, K. (2013). Processing Pathways in Mental Arithmetic—Evidence from Probabilistic Fiber Tracking. PLOS ONE, 8(1), e55455. https://doi.org/10.1371/journal.pone.0055455

Kleineberg, N. N., Dovern, A., Binder, E., Grefkes, C., Eickhoff, S. B., Fink, G. R., & Weiss, P. H. (2018). Action and semantic tool knowledge – Effective connectivity in the underlying neural networks. Human Brain Mapping, 39(9), 3473–3486. https://doi.org/10.1002/hbm.24188

Knill, D. C., & Pouget, A. (2004). The Bayesian brain: The role of uncertainty in neural coding and computation. Trends in Neurosciences, 27(12), 712–719. https://doi.org/10.1016/j.tins.2004.10.007

Kool, W., & Botvinick, M. (2014). A labor/leisure tradeoff in cognitive control. Journal of Experimental Psychology. General, 143(1), 131–141. https://doi.org/10.1037/a0031048

Kool, W., McGuire, J. T., Rosen, Z. B., & Botvinick, M. M. (2010). Decision making and the avoidance of cognitive demand. Journal of Experimental Psychology. General, 139(4), 665–682. https://doi.org/10.1037/a0020198

Krakauer, J. W., & Shadmehr, R. (2007). Towards a computational neuropsychology of action. In P. Cisek, T. Drew, & J. F. Kalaska (Eds.), Progress in Brain Research (Vol. 165, pp. 383–394). Elsevier. https://doi.org/10.1016/S0079-6123(06)65024-3

Kreher, B. W., Schnell, S., Mader, I., Il'yasov, K. A., Hennig, J., Kiselev, V. G., & Saur, D. (2008). Connecting and merging fibres: Pathway extraction by combining probability maps. NeuroImage, 43(1), 81–89. https://doi.org/10.1016/j.neuroimage.2008.06.023

Lacquaniti, F., Guigon, E., Bianchi, L., Ferraina, S., & Caminiti, R. (1995). Representing Spatial Information for Limb Movement: Role of Area 5 in the Monkey. Cerebral Cortex, 5(5), 391–409. https://doi.org/10.1093/cercor/5.5.391

Lambon Ralph, M. A. (2014). Neurocognitive insights on conceptual knowledge and its breakdown. Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences, 369(1634), 20120392. https://doi.org/10.1098/rstb.2012.0392

Laming, D. R. J. (1968). Information theory of choice-reaction times. Academic Press.

Langdon, A. J., Sharpe, M. J., Schoenbaum, G., & Niv, Y. (2018). Model-based predictions for dopamine. Current Opinion in Neurobiology, 49, 1–7. https://doi.org/10.1016/j.conb.2017.10.006

Lanillos, P., Oliva, D., Philippsen, A., Yamashita, Y., Nagai, Y., & Cheng, G. (2020). A review on neural network models of schizophrenia and autism spectrum disorder. Neural Networks, 122, 338–363. https://doi.org/10.1016/j.neunet.2019.10.014

Lauro-Grotto, R., Piccini, C., & Shallice, T. (1997). Modality-specific operations in semantic dementia. Cortex: A Journal Devoted to the Study of the Nervous System and Behavior, 33, 593–622. https://doi.org/10.1016/S0010-9452(08)70720-2

Law, S.-P., Wong, W., & Chiu, K. M. Y. (2005). Preserved reading aloud with semantic deficits: Evidence for a non-semantic lexical route for reading Chinese. Neurocase, 11(3), 167–175. https://doi.org/10.1080/13554790590944618

Lee, D., Seo, H., & Jung, M. W. (2012). Neural Basis of Reinforcement Learning and Decision Making. Annual Review of Neuroscience, 35(1), 287–308. https://doi.org/10.1146/annurev-neuro-062111-150512

Lee, S. W., Shimojo, S., & O'Doherty, J. P. (2014). Neural computations underlying arbitration between model-based and model-free learning. Neuron, 81(3), 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. Journal of the Optical Society of America. A, Optics, Image Science, and Vision, 20(7), 1434–1448. https://doi.org/10.1364/josaa.20.001434

Leshinskaya, A., & Caramazza, A. (2014). Nonmotor Aspects of Action Concepts. Journal of Cognitive Neuroscience, 26(12), 2863–2879. https://doi.org/10.1162/jocn_a_00679

Leshinskaya, A., & Caramazza, A. (2016). For a cognitive neuroscience of concepts: Moving beyond the grounding issue. Psychonomic Bulletin & Review, 23(4), 991–1001. https://doi.org/10.3758/s13423-015-0870-z

Lindsay, D. S., & Jacoby, L. L. (1994). Stroop process dissociations: The relationship between facilitation and interference. Journal of Experimental Psychology: Human Perception and Performance, 20(2), 219–234. https://doi.org/10.1037/0096-1523.20.2.219

Lingnau, A., & Downing, P. E. (2015). The lateral occipitotemporal cortex in action. Trends in Cognitive Sciences, 19(5), 268–277. https://doi.org/10.1016/j.tics.2015.03.006

Logan, G. D., & Zbrodoff, N. J. (1979). When it helps to be misled: Facilitative effects of increasing the frequency of conflicting stimuli in a Stroop-like task. Memory & Cognition, 7(3), 166–174. https://doi.org/10.3758/BF03197535

Logie, R. H. (1986). Visuo-spatial processing in working memory. The Quarterly Journal of Experimental Psychology A: Human Experimental Psychology, 38A(2), 229–247. https://doi.org/10.1080/14640748608401596

Lu, X., Huang, J., Yi, Y., Shen, M., Weng, X., & Gao, Z. (2016). Holding Biological Motion in Working Memory: An fMRI Study. Frontiers in Human Neuroscience, 10, 251. https://doi.org/10.3389/fnhum.2016.00251

Lungarella, M., Metta, G., Pfeifer, R., & Sandini, G. (2003). Developmental robotics: A survey. Connection Science, 15(4), 151–190. https://doi.org/10.1080/09540090310001655110

MacDonald, A. W., Cohen, J. D., Stenger, V. A., & Carter, C. S. (2000). Dissociating the role of the dorsolateral prefrontal and anterior cingulate cortex in cognitive control. Science (New York, N.Y.), 288(5472), 1835–1838. https://doi.org/10.1126/science.288.5472.1835

MacLeod, C. M. (1991). Half a century of research on the Stroop effect: An integrative review. Psychological Bulletin, 109(2), 163–203. https://doi.org/10.1037/0033-2909.109.2.163

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

Makris, N., & Pandya, D. N. (2009). The extreme capsule in humans and rethinking of the language circuitry. Brain Structure & Function, 213(3), 343–358. https://doi.org/10.1007/s00429-008-0199-8

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. Scientific Reports, 9(1), 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. Journal of Neuroscience, 34(45), 14827–14832. https://doi.org/10.1523/JNEUROSCI.2705-14.2014

Martin, M., Beume, L., Kümmerer, D., Schmidt, C. S. M., Bormann, T., Dressing, A., Ludwig, V. M., Umarova, R. M., Mader, I., Rijntjes, M., Kaller, C. P., & Weiller, C. (2016). Differential Roles of Ventral and Dorsal Streams for Conceptual and Production-Related Components of Tool Use in Acute Stroke Patients. Cerebral Cortex (New York, N.Y.: 1991), 26(9), 3754–3771. https://doi.org/10.1093/cercor/bhv179

Martin, M., Dressing, A., Bormann, T., Schmidt, C., Kümmerer, D., Beume, L., Saur, D., Mader, I., Rijntjes, M., Kaller, C., & Weiller, C. (2016). 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. (2016a). Brain activity underlying tool-related and imitative skills after major left hemisphere stroke. Brain: A Journal of Neurology, 139(Pt 5), 1497–1516. https://doi.org/10.1093/brain/aww035

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(5), 1497–1516. https://doi.org/10.1093/brain/aww035

Matheny, M. H., Emenheiser, J., Fon, W., Chapman, A., Salova, A., Rohden, M., Li, J., Hudoba de Badyn, M., Pósfai, M., & Duenas-Osorio, L. (2019). Exotic states in a simple network of nanoelectromechanical oscillators. Science, 363(6431), eaav7932.

Mathys, C., Daunizeau, J., Friston, K. J., & Stephan, K. E. (2011). A bayesian foundation for individual learning under uncertainty. Frontiers in Human Neuroscience, 5, 39. https://doi.org/10.3389/fnhum.2011.00039

Maturana, H. and Varela, F. (1980) Autopoiesis and Cognition The Realization of the Living. D. Reidel Publishing Co., Dordrecht. - References—Scientific Research Publishing. (n.d.). Retrieved December 26, 2022, from https://scirp.org/reference/referencespapers.aspx?referenceid=1546481

McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: I. An account of basic findings. Psychological Review, 88(5), 375–407. https://doi.org/10.1037/0033-295X.88.5.375

Mecklinger, A., Gruenewald, C., Besson, M., Magnié, M.-N., & Von Cramon, D. Y. (2002). Separable neuronal circuitries for manipulable and non-manipulable objects in working memory. Cerebral Cortex, 12, 1115–1123. https://doi.org/10.1093/cercor/12.11.1115

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(8), 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, 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, 48(7), 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 (pp. 95–102). Henry Holt and Co. https://doi.org/10.1037/10039-007

Miller, K. J., Ludvig, E. A., Pezzulo, G., & Shenhav, A. (2018). Chapter 18—Realigning Models of Habitual and Goal-Directed Decision-Making. In R. Morris, A. Bornstein, & A. Shenhav (Eds.), Goal-Directed Decision Making (pp. 407–428). Academic Press. https://doi.org/10.1016/B978-0-12-812098-9.00018-8

Milli, S., Lieder, F., & Griffiths, T. L. (2021). A rational reinterpretation of dual-process theories. Cognition, 217, 104881. https://doi.org/10.1016/j.cognition.2021.104881

Mirza, M. B., Adams, R. A., Friston, K., & Parr, T. (2019). Introducing a Bayesian model of selective attention based on active inference. Scientific Reports, 9(1), Article 1. https://doi.org/10.1038/s41598-019-50138-8

Mirza, M. B., Adams, R. A., Mathys, C. D., & Friston, K. (2016). Scene Construction, Visual Foraging, and Active Inference. Frontiers in Computational Neuroscience, 10(ARTN 56), Article ARTN 56. http://dx.doi.org/10.3389/fncom.2016.00056

Mishkin, M., Ungerleider, L. G., & Macko, K. A. (1983). Object vision and spatial vision: Two cortical pathways. Trends in Neurosciences, 6, 414–417. https://doi.org/10.1016/0166-2236(83)90190-X

Montague, P. R., Dayan, P., & Sejnowski, T. J. (1996). A framework for mesencephalic dopamine systems based on predictive Hebbian learning. Journal of Neuroscience, 16(5), 1936–1947. https://doi.org/10.1523/JNEUROSCI.16-05-01936.1996

Moors, A., & De Houwer, J. (2006). Automaticity: A Theoretical and Conceptual Analysis. Psychological Bulletin, 132(2), 297–326. https://doi.org/10.1037/0033-2909.132.2.297

Murray, J. D., Bernacchia, A., Freedman, D. J., Romo, R., Wallis, J. D., Cai, X., Padoa-Schioppa, C., Pasternak, T., Seo, H., Lee, D., & Wang, X.-J. (2014). A hierarchy of intrinsic timescales across primate cortex. Nature Neuroscience, 17(12), 1661–1663. https://doi.org/10.1038/nn.3862

Musso, M., Weiller, C., Horn, A., Glauche, V., Umarova, R., Hennig, J., Schneider, A., & Rijntjes, M. (2015). A single dual-stream framework for syntactic computations in music and language. NeuroImage, 117, 267–283. https://doi.org/10.1016/j.neuroimage.2015.05.020

Neal, R. M., & Hinton, G. E. (1998). A View of the Em Algorithm that Justifies Incremental, Sparse, and other Variants. In M. I. Jordan (Ed.), Learning in Graphical Models (pp. 355–368). Springer Netherlands. https://doi.org/10.1007/978-94-011-5014-9_12

Nee, D. E., Wager, T. D., & Jonides, J. (2007). Interference resolution: Insights from a meta-analysis of neuroimaging tasks. Cognitive, Affective, & Behavioral Neuroscience, 7(1), 1–17. https://doi.org/10.3758/CABN.7.1.1

Negri, G. A. L., Rumiati, R. I., Zadini, A., Ukmar, M., Mahon, B. Z., & Caramazza, A. (2007). What is the role of motor simulation in action and object recognition? Evidence from apraxia. Cognitive Neuropsychology, 24(8), 795–816. https://doi.org/10.1080/02643290701707412

Negri, G. A., Lunardelli, A., Reverberi, C., Gigli, G. L., & Rumiati, R. I. (2007). Degraded semantic knowledge and accurate object use. Cortex; a Journal Devoted to the Study of the Nervous System and Behavior, 43(3), 376–388. https://doi.org/10.1016/s0010-9452(08)70463-5

Newell, A., & Simon, H. A. (1972). Human problem solving (pp. xiv, 920). Prentice-Hall.

Niessen, E., Fink, G. R., & Weiss, P. H. (2014). Apraxia, pantomime and the parietal cortex. NeuroImage: Clinical, 5, 42–52. https://doi.org/10.1016/j.nicl.2014.05.017

Noppeney, U., Friston, K. J., & Price, C. J. (2004). Degenerate neuronal systems sustaining cognitive functions. Journal of Anatomy, 205(6), 433–442. https://doi.org/10.1111/j.0021-8782.2004.00343.x

Norman, D. A., & Shallice, T. (1986). Attention to Action. In R. J. Davidson, G. E. Schwartz, & D. Shapiro (Eds.), Consciousness and Self-Regulation: Advances in Research and Theory Volume 4 (pp. 1–18). Springer US. https://doi.org/10.1007/978-1-4757-0629-1_1

Ochipa, C., Rothi, L. J., & Heilman, K. M. (1994). Conduction apraxia. Journal of Neurology, Neurosurgery & Psychiatry, 57(10), 1241–1244. https://doi.org/10.1136/jnnp.57.10.1241

O'Keefe, J., & Recce, M. L. (1993). Phase relationship between hippocampal place units and the EEG theta rhythm. Hippocampus, 3(3), 317–330. https://doi.org/10.1002/hipo.450030307

O'Regan, J. K. (2011). Why Red Doesn't Sound Like a Bell: Understanding the feel of consciousness. Oxford University Press.

Otto, A. R., Gershman, S. J., Markman, A. B., & Daw, N. D. (2013). The curse of planning: Dissecting multiple reinforcement-learning systems by taxing the central executive. Psychological Science, 24(5), 751–761.

Ottoboni, G., Ceciliani, A., & Tessari, A. (2021). The Effect of Structured Exercise on Short-Term Memory Subsystems: New Insight on Training Activities. International Journal of Environmental Research and Public Health, 18(14), 7545. https://doi.org/10.3390/ijerph18147545

Pacherie, E. (2008). The phenomenology of action: A conceptual framework. Cognition, 107(1), 179. https://doi.org/10.1016/j.cognition.2007.09.003

Padmala, S., & Pessoa, L. (2011). Reward Reduces Conflict by Enhancing Attentional Control and Biasing Visual Cortical Processing. Journal of Cognitive Neuroscience, 23(11), 3419–3432. https://doi.org/10.1162/jocn_a_00011

Palacios, E. R., Razi, A., Parr, T., Kirchhoff, M., & Friston, K. (2020). On Markov blankets and hierarchical self-organisation. Journal of Theoretical Biology, 486, 110089. https://doi.org/10.1016/j.jtbi.2019.110089

Parr, T., & Friston, K. J. (2017). Working memory, attention, and salience in active inference. Scientific Reports, 7(1), 14678. https://doi.org/10.1038/s41598-017-15249-0

Parr, T., & Friston, K. J. (2018a). The Anatomy of Inference: Generative Models and Brain Structure. Frontiers in Computational Neuroscience, 12. https://www.frontiersin.org/articles/10.3389/fncom.2018.00090

Parr, T., & Friston, K. J. (2018b). The Discrete and Continuous Brain: From Decisions to Movement—And Back Again. Neural Computation, 30(9), 2319–2347. https://doi.org/10.1162/neco_a_01102

Parr, T., Limanowski, J., Rawji, V., & Friston, K. (2021). The computational neurology of movement under active inference. Brain, 144(6), 1799–1818. https://doi.org/10.1093/brain/awab085

Parr, T., Markovic, D., Kiebel, S. J., & Friston, K. J. (2019). Neuronal message passing using Mean-field, Bethe, and Marginal approximations. Scientific Reports, 9(1), Article 1. https://doi.org/10.1038/s41598-018-38246-3

Parr, T., Pezzulo, G., & Friston, K. (2022). Active Inference: The Free Energy Principle in Mind, Brain, and Behavior. MIT Press.

Parr, T., Rees, G., & Friston, K. J. (2018). Computational Neuropsychology and Bayesian Inference. Frontiers in Human Neuroscience, 12, 61. https://doi.org/10.3389/fnhum.2018.00061

Parr, T., Sajid, N., & Friston, K. J. (2020). Modules or Mean-Fields? Entropy, 22(5), Article 5. https://doi.org/10.3390/e22050552

Patterson, K., Nestor, P. J., & Rogers, T. T. (2007). Where do you know what you know? The representation of semantic knowledge in the human brain. Nature Reviews Neuroscience, 8(12), Article 12. https://doi.org/10.1038/nrn2277

Patterson, K., & Shewell, C. (1987). Speak and spell: Dissociations and word-class effects. In The cognitive neuropsychology of language (pp. 273–294). Lawrence Erlbaum Associates, Inc.

Paus, T., Petrides, M., Evans, A. C., & Meyer, E. (1993). Role of the human anterior cingulate cortex in the control of oculomotor, manual, and speech responses: A positron emission tomography study. Journal of Neurophysiology, 70(2), 453–469. https://doi.org/10.1152/jn.1993.70.2.453

Pearce, J. M., & Hall, G. (1980). A model for Pavlovian learning: Variations in the effectiveness of conditioned but not of unconditioned stimuli. Psychological Review, 87(6), 532–552. https://doi.org/10.1037/0033-295X.87.6.532

Pearl, J. (1988). Probabilistic reasoning in intelligent systems: Networks of plausible inference. San Mateo, Calif. : Morgan Kaufmann Publishers. http://archive.org/details/probabilisticrea00pear

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(1), 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, 119–142. https://doi.org/10.1002/hbm.10161

Petreska, B., Adriani, M., Blanke, O., & Billard, A. G. (2007). Apraxia: A review. In C. von Hofsten & K. Rosander (Eds.), Progress in Brain Research (Vol. 164, pp. 61–83). Elsevier. https://doi.org/10.1016/S0079-6123(07)64004-7

Pezzulo, G. (2013). Studying mirror mechanisms within generative and predictive architectures for joint action. Cortex, 49(10), 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. (2013). Computational Grounded Cognition: A new alliance between grounded cognition and computational modeling. Frontiers in Psychology, 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. Frontiers in Psychology, 2. https://doi.org/10.3389/fpsyg.2011.00005

Pezzulo, G., Rigoli, F., & Chersi, F. (2013). 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. (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., Rigoli, F., & Friston, K. (2018a). Hierarchical Active Inference: A Theory of Motivated Control. Trends in Cognitive Sciences, 22(4), 294–306. https://doi.org/10.1016/j.tics.2018.01.009

Pezzulo, G., Rigoli, F., & Friston, K. J. (2018b). Hierarchical Active Inference: A Theory of Motivated Control. Trends in Cognitive Sciences, 22(4), 294–306. https://doi.org/10.1016/j.tics.2018.01.009

Pezzulo, G., Rigoli, F., & Friston, K. J. (2018c). Hierarchical Active Inference: A Theory of Motivated Control. Trends in Cognitive Sciences, 22(4), 294–306. https://doi.org/10.1016/j.tics.2018.01.009

Pezzulo, G., Zorzi, M., & Corbetta, M. (2021). The secret life of predictive brains: What’s spontaneous activity for? Trends in Cognitive Sciences, 25(9), 730–743. https://doi.org/10.1016/j.tics.2021.05.007

Piray, P., & Daw, N. D. (2020). A simple model for learning in volatile environments. PLOS Computational Biology, 16(7), e1007963. https://doi.org/10.1371/journal.pcbi.1007963

Posner, M. I. (1978). Chronometric explorations of mind (pp. xiii, 271). Lawrence Erlbaum.

Posner, M. I., & Petersen, S. E. (1990). The Attention System of the Human Brain. Annual Review of Neuroscience, 13(1), 25–42. https://doi.org/10.1146/annurev.ne.13.030190.000325

Price, C. J., & Friston, K. J. (2002). Degeneracy and cognitive anatomy. Trends in Cognitive Sciences, 6(10), 416–421. https://doi.org/10.1016/S1364-6613(02)01976-9

Proietti, R., Pezzulo, G., & Tessari, A. (2021). An active inference model of hierarchical action understanding, learning and imitation. PsyArXiv. https://doi.org/10.31234/osf.io/ms95f

Rabbitt, P. M. (1966). Errors and error correction in choice-response tasks. Journal of Experimental Psychology, 71(2), 264–272. https://doi.org/10.1037/h0022853

Ramstead, M. J. D., Sakthivadivel, D. A. R., Heins, C., Koudahl, M., Millidge, B., Da Costa, L., Klein, B., & Friston, K. J. (2022). On Bayesian Mechanics: A Physics of and by Beliefs (arXiv:2205.11543). arXiv. https://doi.org/10.48550/arXiv.2205.11543

Rao, R. P. N., & Ballard, D. H. (1999). Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. Nature Neuroscience, 2(1), Article 1. https://doi.org/10.1038/4580

Renzi, E. D. (1985). Methods of Limb Apraxia Examination and Their Bearing On The Interpretation of The Disorder. In E. A. Roy (Ed.), Advances in Psychology (Vol. 23, pp. 45–64). North-Holland. https://doi.org/10.1016/S0166-4115(08)61135-8

Rescorla, R., & Wagner, A. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In Classical Conditioning II: Current Research and Theory: Vol. Vol. 2.

Rijntjes, M., Dettmers, C., Büchel, C., Kiebel, S., Frackowiak, R. S., & Weiller, C. (1999). A blueprint for movement: Functional and anatomical representations in the human motor system. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience, 19(18), 8043–8048.

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. Annual Review of Neuroscience, 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. Nature Reviews. Neuroscience, 2(9), 661–670. https://doi.org/10.1038/35090060

Rizzolatti, G., & Matelli, M. (2003). Two different streams form the dorsal visual system: Anatomy and functions. Experimental Brain Research, 153(2), 146–157. https://doi.org/10.1007/s00221-003-1588-0

Rizzolatti, G., & Sinigaglia, C. (2010). The functional role of the parieto-frontal mirror circuit: Interpretations and misinterpretations. Nature Reviews Neuroscience, 11(4), 264–274. https://doi.org/10.1038/nrn2805

Roesch, M. R., Esber, G. R., Li, J., Daw, N. D., & Schoenbaum, G. (2012). Surprise! Neural correlates of Pearce–Hall and Rescorla–Wagner coexist within the brain. European Journal of Neuroscience, 35(7), 1190–1200. https://doi.org/10.1111/j.1460-9568.2011.07986.x

Rojas-Carulla, M., Schölkopf, B., Turner, R., & Peters, J. (2018). Invariant models for causal transfer learning. The Journal of Machine Learning Research, 19(1), 1309–1342.

Rothi, L. J., Ochipa, C., & Heilman, K. M. (1991). A cognitive neuropsychological model of limb praxis. Cognitive Neuropsychology, 8(6), 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 E. A. Roy (Ed.), Advances in Psychology (Vol. 23, pp. 111–161). North-Holland. https://doi.org/10.1016/S0166-4115(08)61139-5

Rubin, J., Shamir, O., & Tishby, N. (2012). Trading value and information in MDPs. In Decision Making with Imperfect Decision Makers (pp. 57–74). Springer.

Rumelhart, D. E., & Norman, D. A. (1982). Simulating a Skilled Typist: A Study of Skilled Cognitive-Motor Performance. Cognitive Science, 6(1), 1–36. https://doi.org/10.1207/s15516709cog0601_1

Rumiati, R. I., Carmo, J. C., & Corradi-Dell’Acqua, C. (2009). Neuropsychological perspectives on the mechanisms of imitation. Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences, 364(1528), 2337–2347. https://doi.org/10.1098/rstb.2009.0063

Rumiati, R. I., Papeo, L., & Corradi-Dell’Acqua, C. (2010). Higher-level motor processes. Annals of the New York Academy of Sciences, 1191(1), 219–241. https://doi.org/10.1111/j.1749-6632.2010.05442.x

Rumiati, R. I., & Tessari, A. (2002a). Imitation of novel and well-known actions: The role of short-term memory. Experimental Brain Research, 142(3), 425–433. https://doi.org/10.1007/s00221-001-0956-x

Rumiati, R. I., & Tessari, A. (2002b). Imitation of novel and well-known actions: The role of short-term memory. Experimental Brain Research, 142(3), 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. Journal of Cognitive Neuroscience, 17(9), 1420–1431. https://doi.org/10.1162/0898929054985374

Rushworth, M. F. S., & Behrens, T. E. J. (2008). Choice, uncertainty and value in prefrontal and cingulate cortex. Nature Neuroscience, 11(4), 389–397. https://doi.org/10.1038/nn2066

Sajid, N., Parr, T., Hope, T. M., Price, C. J., & Friston, K. J. (2020). Degeneracy and redundancy in active inference. Cerebral Cortex, 30(11), 5750–5766. https://doi.org/10.1093/cercor/bhaa148

Salamone, J. D., Yohn, S. E., López-Cruz, L., San Miguel, N., & Correa, M. (2016). Activational and effort-related aspects of motivation: Neural mechanisms and implications for psychopathology. Brain, 139(5), 1325–1347. https://doi.org/10.1093/brain/aww050

Sales, A. C., Friston, K. J., Jones, M. W., Pickering, A. E., & Moran, R. J. (2019). Locus Coeruleus tracking of prediction errors optimises cognitive flexibility: An Active Inference model. PLOS Computational Biology, 15(1), e1006267. https://doi.org/10.1371/journal.pcbi.1006267

Saur, D., Kreher, B. W., Schnell, S., Kümmerer, D., Kellmeyer, P., Vry, M.-S., Umarova, R., Musso, M., Glauche, V., Abel, S., Huber, W., Rijntjes, M., Hennig, J., & Weiller, C. (2008). Ventral and dorsal pathways for language. Proceedings of the National Academy of Sciences of the United States of America, 105(46), 18035–18040. https://doi.org/10.1073/pnas.0805234105

Sawaguchi, T., & Goldman-Rakic, P. S. (1991). D1 dopamine receptors in prefrontal cortex: Involvement in working memory. Science (New York, N.Y.), 251(4996), 947–950. https://doi.org/10.1126/science.1825731

Schneider, W., & Chein, J. M. (2003). Controlled & automatic processing: Behavior, theory, and biological mechanisms. Cognitive Science, 27(3), 525–559. https://doi.org/10.1207/s15516709cog2703_8

Schultz, W. (1998). Predictive Reward Signal of Dopamine Neurons. Journal of Neurophysiology, 80(1), 1–27. https://doi.org/10.1152/jn.1998.80.1.1

Schultz, W. (2007). Multiple dopamine functions at different time courses. Annual Review of Neuroscience, 30, 259–288. https://doi.org/10.1146/annurev.neuro.28.061604.135722

Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. Science (New York, N.Y.), 275(5306), 1593–1599. https://doi.org/10.1126/science.275.5306.1593

Schwartenbeck, P., FitzGerald, T. H. B., Mathys, C., Dolan, R., & Friston, K. (2015a). The Dopaminergic Midbrain Encodes the Expected Certainty about Desired Outcomes. Cerebral Cortex, 25(10), 3434–3445. https://doi.org/10.1093/cercor/bhu159

Schwartenbeck, P., FitzGerald, T. H. B., Mathys, C., Dolan, R., & Friston, K. (2015b). The Dopaminergic Midbrain Encodes the Expected Certainty about Desired Outcomes. Cerebral Cortex (New York, N.Y.: 1991), 25(10), 3434–3445. https://doi.org/10.1093/cercor/bhu159

Schwartenbeck, P., FitzGerald, T. H. B., Mathys, C., Dolan, R., Kronbichler, M., & Friston, K. (2015). Evidence for surprise minimization over value maximization in choice behavior. Scientific Reports, 5(1), Article 1. https://doi.org/10.1038/srep16575

Schwartenbeck, P., & Friston, K. (2016). Computational Phenotyping in Psychiatry: A Worked Example. ENeuro, 3(4). 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

Seelke, A. M. H., Dooley, J. C., & Krubitzer, L. A. (2012). The Emergence of Somatotopic Maps of the Body in S1 in Rats: The Correspondence Between Functional and Anatomical Organization. PLOS ONE, 7(2), e32322. https://doi.org/10.1371/journal.pone.0032322

Shadmehr, R., Smith, M. A., & Krakauer, J. W. (2010). Error correction, sensory prediction, and adaptation in motor control. Annual Review of Neuroscience, 33, 89–108. https://doi.org/10.1146/annurev-neuro-060909-153135

Shallice, T., & Burgess, P. (1993). Supervisory control of action and thought selection. In Attention: Selection, awareness, and control: A tribute to Donald Broadbent (pp. 171–187). Clarendon Press/Oxford University Press.

Shannon, C. E. (1948). A mathematical theory of communication. The Bell System Technical Journal, 27(3), 379–423. https://doi.org/10.1002/j.1538-7305.1948.tb01338.x

Shelton, P. A., & Knopman, D. S. (1991). Ideomotor apraxia in Huntington's disease. Archives of Neurology, 48, 35–41. https://doi.org/10.1001/archneur.1991.00530130043018

Shenhav, A., Botvinick, M. M., & Cohen, J. D. (2013). The expected value of control: An integrative theory of anterior cingulate cortex function. Neuron, 79(2), 217–240. https://doi.org/10.1016/j.neuron.2013.07.007

Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending and a general theory. Psychological Review, 84(2), 127–190. https://doi.org/10.1037/0033-295X.84.2.127

Silton, R. L., Heller, W., Towers, D. N., Engels, A. S., Spielberg, J. M., Edgar, J. C., Sass, S. M., Stewart, J. L., Sutton, B. P., Banich, M. T., & Miller, G. A. (2010). The time course of activity in dorsolateral prefrontal cortex and anterior cingulate cortex during top-down attentional control. NeuroImage, 50(3), 1292–1302. https://doi.org/10.1016/j.neuroimage.2009.12.061

Silvetti, M., Vassena, E., Abrahamse, E., & Verguts, T. (2018). Dorsal anterior cingulate-brainstem ensemble as a reinforcement meta-learner. PLOS Computational Biology, 14(8), e1006370. https://doi.org/10.1371/journal.pcbi.1006370

Sirigu, A., Zalla, T., Pillon, B., Grafman, J., Agid, Y., & Dubois, B. (1995). Selective impairments in managerial knowledge following pre-frontal cortex damage. Cortex; a Journal Devoted to the Study of the Nervous System and Behavior, 31(2), 301–316. https://doi.org/10.1016/s0010-9452(13)80364-4

Smith, M. A., Ghazizadeh, A., & Shadmehr, R. (2006). Interacting Adaptive Processes with Different Timescales Underlie Short-Term Motor Learning. PLOS Biology, 4(6), e179. https://doi.org/10.1371/journal.pbio.0040179

Smith, R., Friston, K. J., & Whyte, C. J. (2022). A step-by-step tutorial on active inference and its application to empirical data. Journal of Mathematical Psychology, 107, 102632. https://doi.org/10.1016/j.jmp.2021.102632

Smith, R., Kirlic, N., Stewart, J. L., Touthang, J., Kuplicki, R., McDermott, T. J., Taylor, S., Khalsa, S. S., Paulus, M. P., & Aupperle, R. L. (2021). Long-term stability of computational parameters during approach-avoidance conflict in a transdiagnostic psychiatric patient sample. Scientific Reports, 11(1), Article 1. https://doi.org/10.1038/s41598-021-91308-x

Smith, R., Ramstead, M. J. D., & Kiefer, A. (2022). Active inference models do not contradict folk psychology. Synthese, 200(2), 81. https://doi.org/10.1007/s11229-022-03480-w

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

Sohn, M.-H., Albert, M. V., Jung, K., Carter, C. S., & Anderson, J. R. (2007). Anticipation of conflict monitoring in the anterior cingulate cortex and the prefrontal cortex. Proceedings of the National Academy of Sciences, 104(25), 10330–10334. https://doi.org/10.1073/pnas.0703225104

Sokolov, E. N. (1963). Higher Nervous Functions: The Orienting Reflex. Annual Review of Physiology, 25(1), 545–580. https://doi.org/10.1146/annurev.ph.25.030163.002553

Soltani, A., & Izquierdo, A. (2019). Adaptive learning under expected and unexpected uncertainty. Nature Reviews Neuroscience, 20(10), 635–644. https://doi.org/10.1038/s41583-019-0180-y

Stanovich, K. E., & West, R. F. (2000). Individual differences in reasoning: Implications for the rationality debate? Behavioral and Brain Sciences, 23(5), 645–665. https://doi.org/10.1017/S0140525X00003435

Strogatz, S. H. (2012). Sync: How Order Emerges from Chaos In the Universe, Nature, and Daily Life. Hachette UK.

Strogatz, S. H., & Stewart, I. (1993). Coupled Oscillators and Biological Synchronization. Scientific American, 269(6), 102–109.

Sutton, R. S., & Barto, A. G. (1998). Reinforcement Learning: An Introduction. MIT Press.

Swan, J. A., & Pearce, J. M. (1988). The orienting response as an index of stimulus associability in rats. Journal of Experimental Psychology: Animal Behavior Processes, 14(3), 292–301. https://doi.org/10.1037/0097-7403.14.3.292

Swart, J. C., Froböse, M. I., Cook, J. L., Geurts, D. E., Frank, M. J., Cools, R., & den Ouden, H. E. (2017). Catecholaminergic challenge uncovers distinct Pavlovian and instrumental mechanisms of motivated (in)action. ELife, 6, e22169. https://doi.org/10.7554/eLife.22169

Taatgen, N. A., & Lee, F. J. (2003). Production Compilation: A Simple Mechanism to Model Complex Skill Acquisition. Human Factors, 45(1), 61–76. https://doi.org/10.1518/hfes.45.1.61.27224

Tessari, A., Bosanac, D., & Rumiati, R. I. (2006). Effect of learning on imitation of new actions: Implications for a memory model. Experimental Brain Research, 173(3), 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(4), 1111–1126. https://doi.org/10.1093/brain/awm003

Tessari, A., Mengotti, P., Faccioli, L., Tuozzi, 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. Cognitive Neuropsychology, 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(6), 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(18), 2531–2535. https://doi.org/10.1097/00001756-200212200-00030

Testolin, A., & Zorzi, M. (2016). Probabilistic Models and Generative Neural Networks: Towards an Unified Framework for Modeling Normal and Impaired Neurocognitive Functions. Frontiers in Computational Neuroscience, 10. https://www.frontiersin.org/articles/10.3389/fncom.2016.00073

Thompson, E., & Varela, F. J. (2001). Radical embodiment: Neural dynamics and consciousness. Trends in Cognitive Sciences, 5, 418–425. https://doi.org/10.1016/S1364-6613(00)01750-2

Tobler, P. N., Fiorillo, C. D., & Schultz, W. (2005). Adaptive coding of reward value by dopamine neurons. Science (New York, N.Y.), 307(5715), 1642–1645. https://doi.org/10.1126/science.1105370

Todorov, E. (2009). Efficient computation of optimal actions. Proc Natl Acad Sci U S A, 106(28), 11478–11483. https://doi.org/10.1073/pnas.0710743106

Tononi, G., & Cirelli, C. (2006). Sleep function and synaptic homeostasis. Sleep Medicine Reviews, 10(1), 49–62. https://doi.org/10.1016/j.smrv.2005.05.002

Tononi, G., Sporns, O., & Edelman, G. M. (1994). A measure for brain complexity: Relating functional segregation and integration in the nervous system. Proceedings of the National Academy of Sciences, 91(11), 5033–5037. https://doi.org/10.1073/pnas.91.11.5033

Tononi, G., Sporns, O., & Edelman, G. M. (1999). Measures of degeneracy and redundancy in biological networks. Proceedings of the National Academy of Sciences, 96(6), 3257–3262. https://doi.org/10.1073/pnas.96.6.3257

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

van de Laar, T. W., & de Vries, B. (2019). Simulating Active Inference Processes by Message Passing. Frontiers in Robotics and AI, 6, 20. https://doi.org/10.3389/frobt.2019.00020

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(2), 220–250. https://doi.org/10.1016/j.plrev.2013.11.005

van Veen, V., Cohen, J. D., Botvinick, M. M., Stenger, V. A., & Carter, C. S. (2001). Anterior cingulate cortex, conflict monitoring, and levels of processing. NeuroImage, 14(6), 1302–1308. https://doi.org/10.1006/nimg.2001.0923

Varela, F., Lachaux, J.-P., Rodriguez, E., & Martinerie, J. (2001). The brainweb: Phase synchronization and large-scale integration. Nature Reviews Neuroscience, 2(4), Article 4. https://doi.org/10.1038/35067550

Vidaurre, D., Smith, S. M., & Woolrich, M. W. (2017). Brain network dynamics are hierarchically organized in time. Proceedings of the National Academy of Sciences of the United States of America, 114(48), 12827–12832. https://doi.org/10.1073/pnas.1705120114

Vingerhoets, G. (2014). Contribution of the posterior parietal cortex in reaching, grasping, and using objects and tools. Frontiers in Psychology, 5, 151. https://doi.org/10.3389/fpsyg.2014.00151

Von Hofsten, C. (2007). Action in development. Developmental Science, 10(1), 54–60. https://doi.org/10.1111/j.1467-7687.2007.00564.x

Von Neumann, J., & Burks, A. W. (1966). Theory of self-reproducing automata. IEEE Transactions on Neural Networks, 5(1), 3–14.

Vry, M.-S., Saur, D., Rijntjes, M., Umarova, R., Kellmeyer, P., Schnell, S., Glauche, V., Hamzei, F., & Weiller, C. (2012). Ventral and dorsal fiber systems for imagined and executed movement. Experimental Brain Research, 219(2), 203–216. https://doi.org/10.1007/s00221-012-3079-7

Vry, M.-S., Tritschler, L. C., Hamzei, F., Rijntjes, M., Kaller, C. P., Hoeren, M., Umarova, R., Glauche, V., Hermsdoerfer, J., Goldenberg, G., Hennig, J., & Weiller, C. (2015). The ventral fiber pathway for pantomime of object use. NeuroImage, 106, 252–263. https://doi.org/10.1016/j.neuroimage.2014.11.002

Watson, R. T., Fleet, W. S., Gonzalez-Rothi, L., & Heilman, K. M. (1986). Apraxia and the Supplementary Motor Area. Archives of Neurology, 43(8), 787–792. https://doi.org/10.1001/archneur.1986.00520080035016

Wei, K., & Körding, K. (2009). Relevance of Error: What Drives Motor Adaptation? Journal of Neurophysiology, 101(2), 655–664. https://doi.org/10.1152/jn.90545.2008

Weiller, C., Bormann, T., Saur, D., Musso, M., & Rijntjes, M. (2011a). How the ventral pathway got lost: And what its recovery might mean. Brain and Language, 118(1–2), 29–39. https://doi.org/10.1016/j.bandl.2011.01.005

Weiller, C., Bormann, T., Saur, D., Musso, M., & Rijntjes, M. (2011b). How the ventral pathway got lost – And what its recovery might mean. Brain and Language, 118(1), 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 in Cognitive Sciences, 13(9), 369–370; 370–371. https://doi.org/10.1016/j.tics.2009.06.007

Weiller, C., Reisert, M., Peto, I., Hennig, J., Makris, N., Petrides, M., Rijntjes, M., & Egger, K. (2021). The ventral pathway of the human brain: A continuous association tract system. NeuroImage, 234, 117977. https://doi.org/10.1016/j.neuroimage.2021.117977

Weisberg, J., van Turennout, M., & Martin, A. (2007). A Neural System for Learning about Object Function. Cerebral Cortex (New York, N.Y. : 1991), 17(3), 513–521. https://doi.org/10.1093/cercor/bhj176

Westbrook, A., & Braver, T. S. (2016). Dopamine Does Double Duty in Motivating Cognitive Effort. Neuron, 89(4), 695–710. https://doi.org/10.1016/j.neuron.2015.12.029

Wu, D. H., Martin, R. C., & Damian, M. F. (2002). A third route for reading? Implications from a case of phonological dyslexia. Neurocase, 8(4), 274–293. https://doi.org/10.1076/neur.8.3.274.16198

Wurm, M. F., & Caramazza, A. (2022). Two ‘what’ pathways for action and object recognition. Trends in Cognitive Sciences, 26(2), 103–116. https://doi.org/10.1016/j.tics.2021.10.003

Wurm, M. F., Caramazza, A., & Lingnau, A. (2017). Action Categories in Lateral Occipitotemporal Cortex Are Organized Along Sociality and Transitivity. Journal of Neuroscience, 37(3), 562–575. https://doi.org/10.1523/JNEUROSCI.1717-16.2016

Wurm, M. F., & Lingnau, A. (2015). Decoding Actions at Different Levels of Abstraction. Journal of Neuroscience, 35(20), 7727–7735. https://doi.org/10.1523/JNEUROSCI.0188-15.2015

Yuille, A., & Kersten, D. (2006). Vision as Bayesian inference: Analysis by synthesis? Trends in Cognitive Sciences, 10(7), 301–308. https://doi.org/10.1016/j.tics.2006.05.002

Zeki, S., & Shipp, S. (1988). The functional logic of cortical connections. Nature, 335(6188), 311–317. https://doi.org/10.1038/335311a0

Zénon, A., Solopchuk, O., & Pezzulo, G. (2019). An information-theoretic perspective on the costs of cognition. Neuropsychologia, 123, 5–18. https://doi.org/10.1016/j.neuropsychologia.2018.09.013