



## OPEN ACCESS

## EDITED BY

Jacob Raber,  
Oregon Health and Science University,  
United States

## REVIEWED BY

Antonino Casile,  
Italian Institute of Technology (IIT), Italy  
Tim Verbelen,  
Ghent University, Belgium

## \*CORRESPONDENCE

Ivlin Peev Stoianov  
✉ ivlinpeev.stoianov@cnr.it

RECEIVED 23 December 2022

ACCEPTED 03 March 2023

PUBLISHED 20 March 2023

## CITATION

Priorelli M and Stoianov IP (2023) Flexible intentions: An Active Inference theory. *Front. Comput. Neurosci.* 17:1128694. doi: 10.3389/fncom.2023.1128694

## COPYRIGHT

© 2023 Priorelli and Stoianov. This is an open-access article distributed under the terms of the [Creative Commons Attribution License \(CC BY\)](https://creativecommons.org/licenses/by/4.0/). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

# Flexible intentions: An Active Inference theory

Matteo Priorelli and Ivlin Peev Stoianov\*

Institute of Cognitive Sciences and Technologies (ISTC), National Research Council of Italy (CNR), Padua, Italy

We present a normative computational theory of how the brain may support visually-guided goal-directed actions in dynamically changing environments. It extends the Active Inference theory of cortical processing according to which the brain maintains beliefs over the environmental state, and motor control signals try to fulfill the corresponding sensory predictions. We propose that the neural circuitry in the Posterior Parietal Cortex (PPC) compute flexible intentions—or motor plans from a belief over targets—to dynamically generate goal-directed actions, and we develop a computational formalization of this process. A proof-of-concept agent embodying visual and proprioceptive sensors and an actuated upper limb was tested on target-reaching tasks. The agent behaved correctly under various conditions, including static and dynamic targets, different sensory feedbacks, sensory precisions, intention gains, and movement policies; limit conditions were individuated, too. Active Inference driven by dynamic and flexible intentions can thus support goal-directed behavior in constantly changing environments, and the PPC might putatively host its core intention mechanism. More broadly, the study provides a normative computational basis for research on goal-directed behavior in end-to-end settings and further advances mechanistic theories of active biological systems.

## KEYWORDS

Active Inference, sensorimotor control, Posterior Parietal Cortex, intentions, Predictive Coding

## 1. Introduction

Traditionally, sensorimotor control in goal-directed actions like object-reaching is viewed as a sensory-response mapping involving several steps, starting with perception, movement planning in the body posture domain, translation of this plan in muscle commands, and finally movement execution (Erlhagen and Schöner, 2002). However, each of these steps is hindered by noise and delays, which make the approach unfeasible to operate in changing environments (Franklin and Wolpert, 2011). Instead, Predictive Coding or “Bayesian Brain” theories propose that prior knowledge and expectations over the environmental and bodily contexts provide crucial anticipatory information (Rao and Ballard, 1999). Under this perspective, motor control begins with target anticipation and motor planning even before obtaining sensory evidence. Here, we take on this view and extend an increasingly popular Predictive Coding based theory of action, Active Inference (Friston et al., 2010), with the formalization of flexible target-dependent motor plans. Moreover, based on extensive neural evidence for the role of the PPC in goal coding and motor planning (Snyder et al., 2000; Galletti et al., 2022), we propose that this cortical structure is the most likely neural correlate of the core intention manipulation process.

In primates, the dorsomedial visual stream provides critical support for continuously monitoring the body posture and the spatial location of objects to specify and guide actions, and for performing visuomotor transformations in the course of the evolving movement

(Cisek and Kalaska, 2010; Fattori et al., 2017; Galletti and Fattori, 2018). The PPC, located at the apex of the dorsal stream, is also bidirectionally connected to frontal areas, motor and somatosensory cortex, placing it in a privileged position to set goal-directed actions and continuously adjust motor plans by tracking moving targets and posture (Andersen, 1995; Gamberini et al., 2021) in a common reference frame (Cohen and Andersen, 2002). Undoubtedly, the PPC plays a crucial role in visually-guided motor control (Desmurget et al., 1999; Filippini et al., 2018; Gamberini et al., 2021)—with the specific subregion V6A involved in the control of reach-to-grasp actions (Galletti et al., 2022)—but its peculiar role is still disputed. The most consistent view is that the PPC estimates the states of both body and environment and optimizes their interactions (Medendorp and Heed, 2019). Others see the PPC as a task estimator (Haar and Donchin, 2020) or as being involved in endogenous attention and task setting (Corbetta and Shulman, 2002). Its underlying computational mechanism is not fully understood, especially as regards the definition of goals in motor planning and their integration within the control process (Shadmehr and Krakauer, 2008). For example, the prevailing Optimal Feedback Control theory defines motor goals through task-specific cost functions (Todorov, 2004). Neural-level details of motor goal coding are becoming increasingly important in light of the growing demand for neural interfaces that provide information about motor intents (Gallego et al., 2022) in support of intelligent assistive devices (Velliste et al., 2008; Srinivasan et al., 2021).

Intentions encode motor goals—or plans—set before the beginning of motor acts themselves and could be therefore viewed as memory holders of voluntary actions (Andersen, 1995; Snyder et al., 1997; Lau et al., 2004; Fogassi et al., 2005). Several cortical areas handle different aspects of this process: the Premotor cortex (PM) encodes structuring while the Supplementary Motor Area (SMA) controls phasing (Gallego et al., 2022). In turn, the PPC plays a role in building motor plans and their dynamic tuning, as different PPC neurons are sensitive to different intentions (Snyder et al., 2000). Notably, intention neurons respond not only when performing a given action but also during its observation, allowing observers to *predict the goal of the observed action and, thus, to “read” the intention of the acting individual* (Fogassi et al., 2005). Motor goals have also been observed down the motor hierarchy, which is an expression of Hierarchical Predictive Coding in the motor domain (Friston et al., 2011).

To investigate how neural circuitry in the PPC supports sensory-guided actions through motor intentions from a computational point of view, we adopted the Active Inference theory of cognitive and motor control, which provides fundamental insights of increasing appeal about the computational role and principles of the nervous system, especially about the perception-action loop (Friston and Kiebel, 2009; Friston et al., 2010; Bogacz, 2017; Parr et al., 2022). Indeed, Active Inference provides a formalization of these two cortical tasks, both of which are viewed as aiming to resolve the critical goal of all organisms: to survive in uncertain environments by operating within preferred states (e.g., maintaining a constant temperature). Accordingly, both tasks are implemented by dynamic minimization of a quantity called *free energy*, whose process generally corresponds to the minimization of high- and low-level prediction errors, that is, the satisfaction of prior and sensory expectations. There are two branches of Active

Inference appropriate to tackle two different levels of control. The discrete framework can explain high-level cognitive control processes such as planning and decision-making, i.e., it evaluates expected outcomes to select actions in discrete entities (Pezzulo et al., 2018). In turn, dynamic adjustment of action plans in the PPC matches by functionality the Active Inference framework in continuous state space (Friston et al., 2011, 2012). In short, this theory departs from classical views of perception, motor planning (Erlhagen and Schöner, 2002), and motor control (Todorov, 2004), unifying and considering them as a dynamic probabilistic inference problem (Toussaint and Storkey, 2006; Kaplan and Friston, 2018; Levine, 2018; Millidge et al., 2020). The biologically implausible cost functions typical of Optimal Control theories are replaced by high-level priors defined in the extrinsic state space, allowing complex movements such as walking or handwriting (Friston, 2011; Adams et al., 2013).

In the following, we first outline the background computational framework and then elaborate on movement planning and intentionality in continuous Active Inference. Our most critical contributions regard the formalization of goal-directed behavior and the processes linking dynamic goals (e.g., moving visual targets) with motor plans through the definition of flexible intentions. We also investigate a more parsimonious approach to motor control based solely on proprioceptive predictions. We then provide implementation details and a practical demonstration of the theoretical contribution in terms of a simulated Active Inference agent, which we show is capable of detecting and reaching static visual goals and tracking moving targets. We also provide detailed performance statistics and investigate the effects of system parameters whose balance is critical to movement stability. Additionally, gradient analysis provides crucial insights into the causes of the movements performed. Finally, we discuss how intentions could be selected to perform a series of goal-directed steps, e.g., a multi-phase action, and illustrate conditions for neurological disorders.

## 2. Computational background

We first outline the computational principles of the underlying probabilistic and Predictive Coding approach and provide background on variational inference, free energy minimization, Active Inference, and variational autoencoders necessary to comprehend the following main contribution.

### 2.1. The Bayesian brain hypothesis

An interesting visual phenomenon, called *binocular rivalry*, happens when two different images are presented simultaneously to each eye: the perception does not conform to the visual input but alternates between the two images. How and why does this happen? It is well-known that priors play a fundamental role in driving the dynamics of perceptual experience, but dominant views of the brain as a feature detector that passively receives sensory signals and computes motor commands have so far failed to explain how such illusions could arise.

In recent years, there has been increasing attention to a radically new theory of the mind called the *Bayesian brain*, according to which our brain is a sophisticated machine that constantly makes use of Bayesian reasoning to capture causal relationships in the world and deliver optimal behavior in an uncertain environment (Doya, 2007; Hohwy, 2013; Pezzulo et al., 2017). At the core of the theory is the Bayes theorem, whose application here implies that posterior beliefs about the world are updated according to the product of prior beliefs and the likelihood of observing sensory input. In this view, perception is more than a simple bottom-up feedforward mechanism that detects features and objects from the current sensorium; rather, it comprises a predictive top-down generative model which continuously anticipates the sensory input to test hypotheses and explain away ambiguities.

According to the Bayesian brain hypothesis, this complex task is accomplished by *Predictive Coding*, implemented through message passing of top-down predictions and bottom-up prediction errors between adjacent cortical layers (Rao and Ballard, 1999). The former are generated from latent states maintained at the highest levels, representing beliefs about the causes of the environment, while the latter are computed by comparing sensory-level predictions with the actual observations. Each prediction will then act as a cause for the layer below, while the prediction error will convey information to the layer above. It is thanks to this hierarchical organization and through error minimization at every layer that the cortex is supposed to be able to mimic and capture the inherently hierarchical relationships that model the world. In this view, sensations are only needed in that they provide, through the computation of prediction errors, a measure of how good the model is and a cue to correct future predictions. Thus, ascending projections do not encode the features of a stimulus, but rather how much the brain is *surprised* about it, considering the strict correlation between surprise and model uncertainty.

## 2.2. Variational bayes

Organisms are supposed to implement model fit or error minimization by some form of variational inference, a broad family of techniques based on the *calculus of variations* and used to approximate intractable posteriors that would otherwise be infeasible to compute analytically or even with classical sampling methods like Monte Carlo (Bishop, 2006). Under the Bayesian brain hypothesis, we can assume that the nervous system maintains latent variables  $\mathbf{z}$  about both the unknown state of the external world and the internal state of the organism. By exploiting a prior knowledge  $p(\mathbf{z})$  and the partial evidence  $p(\mathbf{s})$  of the environment provided by its sensors, it can apply Bayesian inference to improve its knowledge (Ma et al., 2006). To do so, given the observation  $\mathbf{s}$ , the nervous system needs to evaluate the posterior  $p(\mathbf{z}|\mathbf{s})$ :

$$p(\mathbf{z}|\mathbf{s}) = \frac{p(\mathbf{z}, \mathbf{s})}{p(\mathbf{s})} \quad (1)$$

However, directly computing such quantity is infeasible due to the intractability of the marginal  $p(\mathbf{s}) = \int p(\mathbf{z}, \mathbf{s}) d\mathbf{z}$ , which involves integration over the joint density  $p(\mathbf{z}, \mathbf{s})$ . What does the variational approach is approximating the posterior with a simpler to compute

*recognition* distribution  $q(\mathbf{z}) \approx p(\mathbf{z}|\mathbf{s})$  through minimization of the Kullback-Leibler (KL) divergence between them:

$$D_{KL}[q(\mathbf{z})||p(\mathbf{z}|\mathbf{s})] = \int_{\mathbf{z}} q(\mathbf{z}) \ln \frac{q(\mathbf{z})}{p(\mathbf{z}|\mathbf{s})} d\mathbf{z} \quad (2)$$

The KL divergence can be rewritten as the difference between log evidence  $\ln p(\mathbf{s})$  and a quantity  $\mathcal{L}(q)$  known as *evidence lower bound*, or ELBO (Bishop, 2006):

$$D_{KL}[q(\mathbf{z})||p(\mathbf{z}|\mathbf{s})] = \ln p(\mathbf{s}) - \int_{\mathbf{z}} q(\mathbf{z}) \ln \frac{p(\mathbf{z}, \mathbf{s})}{q(\mathbf{z})} d\mathbf{z} = \ln p(\mathbf{s}) - \mathcal{L}(q) \quad (3)$$

Since the KL divergence is always nonnegative, the ELBO provides a lower bound on log evidence, i.e.,  $\mathcal{L}(q) \leq \ln p(\mathbf{s})$ . Therefore, minimizing the KL divergence with respect to  $q(\mathbf{z})$  is equivalent to maximizing  $\mathcal{L}(q)$ , which at its maximum corresponds to an approximate density that is closest the most to the real posterior, depending on the particular choice of the form of  $q(\mathbf{z})$ . In general, few assumptions are made about the form of this distribution—a multivariate Gaussian is a typical choice—with a trade-off between having a tractable optimization process and still leading to a good approximate posterior.

## 2.3. Free energy and prediction errors

How can Bayesian inference be implemented through a simple message passing of prediction errors? Friston (2002, 2005) proposed an elegant solution based on the so-called *free energy*, a concept borrowed from thermodynamics and defined as the negative ELBO. Accordingly, Equation (3) can be rewritten as:

$$\mathcal{F}(\mathbf{z}, \mathbf{s}) = -\mathcal{L}(q) = D_{KL}[q(\mathbf{z})||p(\mathbf{z}|\mathbf{s})] - \ln p(\mathbf{s}) = \int_{\mathbf{z}} q(\mathbf{z}) \ln \frac{q(\mathbf{z})}{p(\mathbf{z}, \mathbf{s})} d\mathbf{z} \quad (4)$$

Minimizing the free energy with respect to the latent states  $\mathbf{z}$ —a process called *perceptual inference*—is then equivalent to ELBO maximization and provides an upper bound on surprise:

$$\mathbf{z} = \arg \min_{\mathbf{z}} \mathcal{F}(\mathbf{z}, \mathbf{s}) \quad (5)$$

In this way, the organism indirectly minimizes model uncertainty and is able to learn the causal relations between unknown states and sensory input, and to generate predictions based on its current representation of the environment. Free energy minimization is simpler than dealing with the KL divergence between the approximate and true posteriors as the former depends on quantities that the organism has access to, namely the approximate posterior and the *generative model*.

To this concern, it is necessary to distinguish between the latter and the real distribution producing sensory data, called *generative process*, which can be modeled with the following non-linear stochastic equations:

$$\begin{aligned} \mathbf{s} &= \mathbf{g}(\mathbf{z}) + \mathbf{w}_s \\ \dot{\mathbf{z}} &= \mathbf{f}(\mathbf{z}) + \mathbf{w}_z \end{aligned} \quad (6)$$

Where the function  $\mathbf{g}$  maps latent states or causes  $\mathbf{z}$  to observed states or sensations  $\mathbf{s}$ , the function  $\mathbf{f}$  encodes the dynamics of the

system, i.e., the evolution of  $\mathbf{z}$  over time, while  $\mathbf{w}_s$  and  $\mathbf{w}_z$  are noise terms that describe system uncertainty.

Nervous systems are supposed to approximate the generative process by making a few assumptions: that (i) under the mean-field approximation the recognition density can be partitioned into independent distributions:  $q(\mathbf{z}) = \prod_i q(\mathbf{z}_i)$ , and that (ii) under the Laplace approximation each of these partitions is Gaussian:  $q(\mathbf{z}_i) = \mathcal{N}(\boldsymbol{\mu}_i, \boldsymbol{\Pi}_i^{-1})$ , where  $\boldsymbol{\mu}_i$  represents the most plausible hypothesis—also called *belief* about the hidden state  $\mathbf{z}_i$ —and  $\boldsymbol{\Pi}_i$  is its precision matrix (Friston et al., 2007). In this way, the free energy does not depend on  $\mathbf{z}$  and simplifies as follows:

$$\mathcal{F}(\boldsymbol{\mu}, \mathbf{s}) = -\ln p(\boldsymbol{\mu}, \mathbf{s}) + C = -\ln p(\mathbf{s}|\boldsymbol{\mu}) - \ln p(\boldsymbol{\mu}) + C \quad (7)$$

Where  $C$  is a constant term. A more precise description of the unknown environmental dynamics can be achieved by considering not only the 1st order of Equation 6 but also higher temporal orders of the corresponding approximations:  $\tilde{\boldsymbol{\mu}} = \{\boldsymbol{\mu}, \boldsymbol{\mu}', \boldsymbol{\mu}'', \dots\}$ —called *generalized coordinates* (Friston, 2008; Friston et al., 2008). This allows us to better represent the environment with the following generalized model:

$$\begin{aligned} \tilde{\mathbf{s}} &= \tilde{\mathbf{g}}(\tilde{\boldsymbol{\mu}}) + \mathbf{w}_s \\ \mathcal{D}\tilde{\boldsymbol{\mu}} &= \tilde{\mathbf{f}}(\tilde{\boldsymbol{\mu}}) + \mathbf{w}_\mu \end{aligned} \quad (8)$$

Where  $\mathcal{D}$  is the differential (shift) operator matrix such that  $\mathcal{D}\tilde{\boldsymbol{\mu}} = \{\boldsymbol{\mu}', \boldsymbol{\mu}'', \dots\}$ ,  $\tilde{\mathbf{s}}$  denotes the generalized sensors, while  $\tilde{\mathbf{g}}$  and  $\tilde{\mathbf{f}}$  denote the generalized model functions of all temporal orders. Note that in this system, the sensory data at a particular dynamical order  $s^{[d]}$ —where  $[d]$  is the order—engage only with the same order of belief  $\boldsymbol{\mu}^{[d]}$ , while the generalized equation of motion, or system dynamics, specifies the coupling between adjacent orders. Such equations are generated from the generalized likelihood and prior distributions, which can be expanded as follows:

$$\begin{aligned} p(\tilde{\mathbf{s}}|\tilde{\boldsymbol{\mu}}) &= \prod_d p(s^{[d]}|\boldsymbol{\mu}^{[d]}) \\ p(\tilde{\boldsymbol{\mu}}) &= \prod_d p(\boldsymbol{\mu}^{[d+1]}|\boldsymbol{\mu}^{[d]}) \end{aligned} \quad (9)$$

As defined above, these variational probability distributions are assumed to be Gaussian:

$$\begin{aligned} p(s^{[d]}|\boldsymbol{\mu}^{[d]}) &= \frac{\boldsymbol{\Pi}_s}{\sqrt{(2\pi)^L}} \exp\left(-\frac{1}{2}\mathbf{e}_s^{[d]T}\boldsymbol{\Pi}_s\mathbf{e}_s^{[d]}\right) \\ p(\boldsymbol{\mu}^{[d+1]}|\boldsymbol{\mu}^{[d]}) &= \frac{\boldsymbol{\Pi}_\mu}{\sqrt{(2\pi)^M}} \exp\left(-\frac{1}{2}\boldsymbol{\epsilon}_\mu^{[d]T}\boldsymbol{\Pi}_\mu\boldsymbol{\epsilon}_\mu^{[d]}\right) \end{aligned} \quad (10)$$

Where  $L$  and  $M$  are the dimensions of sensations and internal beliefs, respectively with precisions  $\boldsymbol{\Pi}_s$  and  $\boldsymbol{\Pi}_\mu$ . Note that the probability distributions are expressed in terms of sensory and dynamics prediction errors:

$$\boldsymbol{\epsilon}_s^{[d]} = \mathbf{s}^{[d]} - \mathbf{g}^{[d]}(\boldsymbol{\mu}^{[d]}) \quad (11)$$

$$\boldsymbol{\epsilon}_\mu^{[d]} = \boldsymbol{\mu}^{[d+1]} - \mathbf{f}^{[d]}(\boldsymbol{\mu}^{[d]}) \quad (12)$$

The factorized probabilistic approximation of the dynamic model allows easy state estimation performed by iterative gradient

descent over the generalized coordinates, that is, by changing the belief  $\tilde{\boldsymbol{\mu}}$  over the hidden states at every temporal order:

$$\dot{\tilde{\boldsymbol{\mu}}} - \mathcal{D}\tilde{\boldsymbol{\mu}} = -\partial_{\tilde{\boldsymbol{\mu}}}\mathcal{F}(\tilde{\boldsymbol{\mu}}, \tilde{\mathbf{s}}) \quad (13)$$

Gradient descent is tractable because the Gaussian variational functions are smooth and differentiable and the derivatives are easily computed in terms of generalized prediction errors, since the logarithm of Equation (7) vanishes the exponent of the Gaussian. The belief update thus turns to:

$$\dot{\tilde{\boldsymbol{\mu}}} = \mathcal{D}\tilde{\boldsymbol{\mu}} + \frac{\partial \tilde{\mathbf{g}}^T}{\partial \tilde{\boldsymbol{\mu}}} \tilde{\boldsymbol{\Pi}}_s \tilde{\boldsymbol{\epsilon}}_s + \frac{\partial \tilde{\mathbf{f}}^T}{\partial \tilde{\boldsymbol{\mu}}} \tilde{\boldsymbol{\Pi}}_\mu \tilde{\boldsymbol{\epsilon}}_\mu - \mathcal{D}^T \tilde{\boldsymbol{\Pi}}_\mu \tilde{\boldsymbol{\epsilon}}_\mu \quad (14)$$

It is crucial to keep in mind the nature of the three components that compose this update equation: a likelihood error computed at the sensory level, a backward error arising from the next temporal order, and a forward error coming from the previous order. These terms represent the free energy gradients relative to the belief  $\boldsymbol{\mu}^{[d]}$  of Equation (11) for the likelihood, and  $\boldsymbol{\mu}^{[d+1]}$  and  $\boldsymbol{\mu}^{[d]}$  of Equation (12) for the dynamics errors.

In short, by making a few plausible simplifying assumptions, the complexity of free energy minimization reduces to the generation of predictions, which are constantly compared with sensory observations to determine a prediction error signal. This error then flows back through the cortical hierarchy to adjust the distribution parameters accordingly and minimize sensory surprise—or maximize evidence—in the long run.

## 2.4. Active Inference

Describing the relationship between Predictive Coding and Bayesian inference still does not explain why has the cortex evolved in such a peculiar way. The answer comes from the so-called *free energy principle* (FEP), regard to which the Bayesian brain hypothesis is supposed to be a corollary. Indeed, learning the causal relationships of some observed data (e.g., what causes an increase in body temperature) is insufficient to keep organisms alive (e.g., maintaining the temperature in a vital range).

The FEP states that, for an organism to maintain a state of homeostasis and survive, it must constantly and actively restrict the set of latent states in which it lives to a narrow range of life-compatible possibilities, counteracting the natural tendency for disorder (Friston, 2012)—hence the relationship with thermodynamics. If these states are defined by the organism's phenotype, from the point of view of its internal model they are exactly the states that it expects to be less surprising. Thus, while perceptual inference tries to optimize the belief about hidden causes to explain away sensations, if on the other hand the assumptions defined by the phenotype are considered to be the true causes of the world, interacting with the external environment means that the agent will try to sample those sensations that make the assumptions true, fulfilling its needs and beliefs. *Active inference* becomes a self-fulfilling prophecy. In this view, there is no difference between a desire and a belief: we simply seek the states in which we expect to find ourselves (Friston et al., 2010; Buckley et al., 2017).

For achieving a goal-directed behavior, it is then sufficient to minimize the free energy also with respect to the action (see Equation 7):

$$\mathbf{a} = \arg \min_a \mathcal{F}(\boldsymbol{\mu}, \mathbf{s}) \quad (15)$$

Given that motor control signals only depend on sensory information, we obtain:

$$\dot{\mathbf{a}} = -\partial_a \mathcal{F}(\tilde{\mathbf{s}}, \tilde{\boldsymbol{\mu}}) = -\frac{\partial \mathcal{F}}{\partial \tilde{\mathbf{s}}} \frac{\partial \tilde{\mathbf{s}}}{\partial \mathbf{a}} = -\frac{\partial \tilde{\mathbf{s}}^T}{\partial \mathbf{a}} \tilde{\boldsymbol{\Pi}}_s \tilde{\boldsymbol{\epsilon}}_s \quad (16)$$

Minimizing the free energy of all sensory signals is certainly useful, as every likelihood contribution will drive the belief update; however, it requires the knowledge of an inverse mapping from exteroceptive sensations to actions (Baltieri and Buckley, 2019), which is considered a "hard problem" being in general highly non-linear and not univocal (Friston et al., 2010). In a more realistic scenario, only proprioception drives the minimization of free energy with respect to the motor signals; this process is easier to realize since the corresponding sensory prediction is already in the intrinsic domain. Control signals sent from the motor cortex are then not motor commands as in classical views of Optimal Control theories; rather, they consist of predictions that define the desired trajectory. Under this perspective, proprioceptive prediction errors computed locally at the spinal cord serve two purposes that only differ in how these signals are conveyed. They drive the current belief toward sensory observations—happening to realize perception—like for exteroceptive signals. But they also drive sensory observations toward the current belief by suppression in simple reflex arcs that activate the corresponding muscles—thus happening to realize movement (Adams et al., 2013; Parr and Friston, 2018; Versteeg et al., 2021).

In conclusion, perception and action can be seen as two sides of the same coin implementing the common vital goal of minimizing entropy or average surprise. In this view, what we perceive never tries to perfectly match the real state of affairs of the world, but is constantly biased toward our preferred states. This means that action only indirectly fulfills future goals; instead, it continuously tries to fill the gap between sensations and predictions generated from our already biased beliefs.

## 2.5. Variational autoencoders

Variational Autoencoders (VAEs) belong to the family of generative models, since they learn the joint distribution  $p(\mathbf{z}, \mathbf{s})$  and can generate synthetic data similar to the input, given a prior distribution  $p(\mathbf{z})$  over the latent space. VAEs use the variational Bayes approach to capture the posterior distribution  $p(\mathbf{z}|\mathbf{s})$  of the latent representation of the inputs when the computation of the marginal is intractable (Goodfellow et al., 2016). A VAE is composed of two probability distributions, both of which are assumed to be Gaussian: a probabilistic *encoder* corresponding to the recognition distribution  $q(\mathbf{z}|\mathbf{s})$ , and a generative function  $p(\mathbf{s}|\mathbf{z})$  called probabilistic *decoder* computing a distribution over the input

space given a latent representation  $\mathbf{z}$  (Figure 3C):

$$\begin{aligned} q(\mathbf{z}|\mathbf{s}) &= \mathcal{N}(\mathbf{z}|\boldsymbol{\mu}_\phi, \boldsymbol{\Sigma}_\phi) \\ p(\mathbf{s}|\mathbf{z}) &= \mathcal{N}(\mathbf{s}|\boldsymbol{\mu}_\theta, \boldsymbol{\Sigma}_\theta) \end{aligned} \quad (17)$$

Although VAEs have many similarities with traditional autoencoders, they are actually a derivation of the AEVB algorithm when a neural network is used for the recognition distribution (Kingma and Welling, 2014). Unlike other variational techniques, the approximate posterior is generally not assumed to be factorial, but since the calculation of the ELBO gradient  $\nabla_\phi \mathcal{L}_{\theta, \phi}(\mathbf{s})$  is biased, a method called *reparametrization trick* is used so that it is independent of the parameters  $\phi$ . This method works by expressing the latent variable  $\mathbf{z}$  by a function:

$$\mathbf{z} = r(\boldsymbol{\epsilon}, \phi, \mathbf{s}) \quad (18)$$

Where  $\boldsymbol{\epsilon}$  is an auxiliary variable independent of  $\phi$  and  $\mathbf{s}$ . The ELBO  $\tilde{\mathcal{L}}_{\theta, \phi}(\mathbf{s})$  for a single data point can thus be expressed as:

$$\tilde{\mathcal{L}}_{\theta, \phi}(\mathbf{s}) = -D_{KL}[q(\mathbf{z}|\mathbf{s})||p(\mathbf{z})] + \frac{1}{M} \sum_m \log p(\mathbf{s}|\mathbf{z}_m) \quad (19)$$

Which can be minimized through backpropagation. Here, the KL divergence can be seen as a regularizer, while the second RHS term is an expected negative reconstruction term that depends on all the  $m$ th components of the latent variable  $\mathbf{z}$ .

## 3. A framework for flexible intentions

In what follows, we develop a computational theory of the circuitry controlling goal-directed actions in a dynamically changing environment through flexible intentions and discuss its putative neural basis in the PPC and related areas. We first elaborate on intentionality in Active Inference, then provide a proof-of-concept agent endowed with visual input. The theory is exemplified and assessed in the following sections through simulations of visually-guided behaviors. The theoretical work is motivated by basic research showing the critical role of the PPC in goal-directed sensorimotor control through intention coding (Andersen, 1995; Desmurget et al., 1999; Galletti and Fattori, 2018) and extends previous theoretical and applied research on Active Inference (Friston et al., 2009; Pio-Lopez et al., 2016; Lanillos and Cheng, 2018; Limanowski and Friston, 2020) and VAE-based vision support (Rood et al., 2020; Sancaktar et al., 2020). The simulations are inspired by a classical monkey reaching task (Breviglieri et al., 2014).

### 3.1. Flexible intentions

State-of-the-art implementations of continuous Active Inference have proven to successfully tackle a wide range of tasks, from oculomotion dynamics (Adams et al., 2015) to the well-known mountain car problem (Friston et al., 2009). Most simulations involve reaching movements in robotic experiments, where several strategies have been tried for designing goal states,

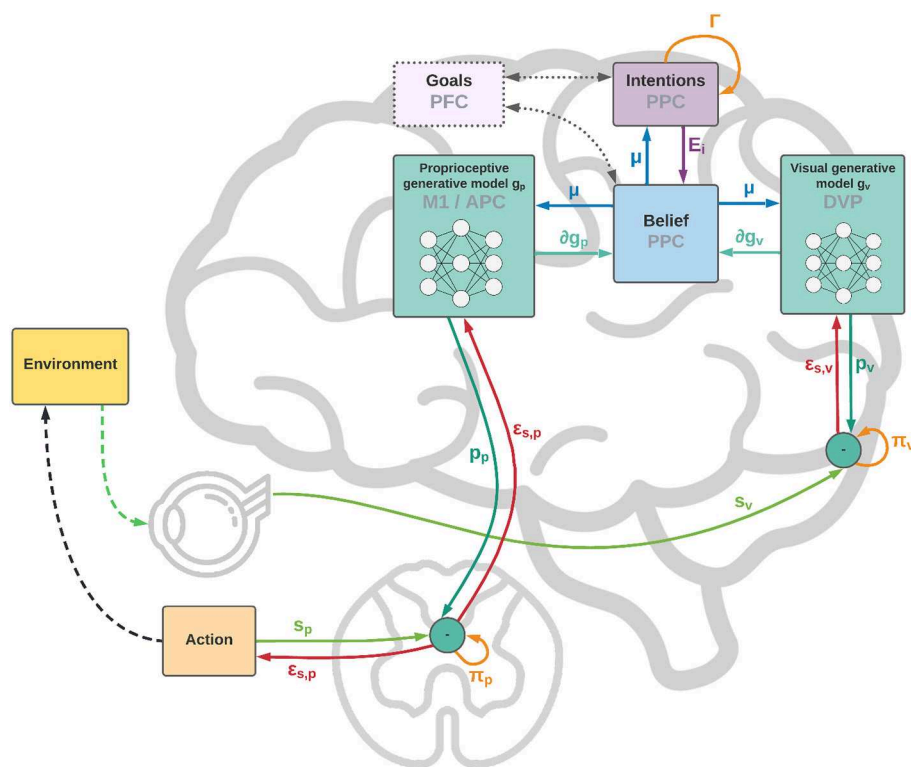


FIGURE 1

Functional architecture and cortical overlay. The process starts with the computation of future intentions  $h$  (not explicitly represented in the figure) in the PPC under the coordination of frontal and motor areas. In the middle of the sensorimotor hierarchy, the PPC maintains beliefs  $\mu$  over the latent causes of sensory observations  $s_p$  and  $s_v$ , and computes proprioceptive and visual predictions through the somatosensory and dorsal visual pathways (for simplicity, we have omitted the somatomotor pathway and considered a single mechanism for both motor control and belief inference). The lower layers of the hierarchy compute sensory prediction errors  $\epsilon_{s,p}$  and  $\epsilon_{s,v}$ , while the higher layers compute intention prediction errors  $E_i$ ; both are propagated back toward the PPC, which thus integrates information from multiple sensory modalities and intentions. Free energy is minimized throughout the cortical hierarchy by changing the belief about the causes of the current observation (perception) and by sending proprioceptive predictions from the motor cortex to the reflex arcs (action). An essential element of this process is the computation of gradients  $\partial g_p$  and  $\partial g_v$  by inverse mappings from the sensations toward the deepest latent states. In this process, intentions act as high-level attractors and the belief propagated down to compute sensorimotor predictions embeds a component directing the body state toward the goals.

which are expressed in terms of an attractor embedded in the system dynamics. However, there seem to be a few issues regarding biological plausibility. First, the goal state is usually static and the agent is not able to deal with continuously changing environments, expecting that the world will always evolve in the same way (Baïoumy et al., 2020). For dynamic goals, one has to use low-level information of sensory signals (e.g., a visual input about a moving target) directly into the high-level dynamics function (Friston, 2011). Second, when goals are specified in an exteroceptive domain, one uses sensory predictions to obtain a belief update direction through backpropagation of the corresponding error (Oliver et al., 2019; Sancaktar et al., 2020). In this case, the same generative model that produces predictions and compares them with the actual observations, has to be duplicated into the system dynamics to further compare the belief with the desired cue. In other words, two specular mechanisms are used for the same model, with additional concerns when the latter can be changed by learning.

A common question seems to be behind these two similar issues: how does dynamic sensory information get available for generating high-level dynamic goals? The same inference process

of environmental causes should be at work for the same signal flow, and a goal state should be computed locally without information passed inconsistently. How then to design a flexible exteroceptive attractor that avoids implausible scenarios?

Although the high-level latent state could be as simple as encoding body configurations only, an agent could also maintain a dynamically estimated belief over moving objects in the scene. An intention can then be computed by exploiting this new information to compute a future action goal in terms of body posture, so that the attractor—either defined in the belief domain or at the sensory level—is not fixed but depends on current perceptual and internal representations of the world (but also on past memorized experiences). This intention may also depend on priors generated from higher-level areas (Friston et al., 2011), so that the considered belief is located at an intermediate level between the generative models that produce sensory predictions, and the ones that define its evolution over time. In a non-trivial task, its dynamics may be generally composed of several contributions and not restricted to a single intention: we thus propose to decompose it into a set of functions, each one providing an independent expectation that the agent will find itself in a particular state. The belief is then

constantly subject to several forces of two natures: one from lower hierarchical levels—proportional to *sensory prediction errors*—that pulls it toward what the agent is currently perceiving, and one from lateral or higher connections—which we call *intention prediction errors*—that pulls it toward what the agent expects to perceive in the future.

As shown in [Figure 1](#), from a neural perspective the PPC is the ideal candidate for a cortical structure computing beliefs over bodily states and flexible intentions: on one hand, being at the apex of the Dorsal Visual Stream (DVS) and other sensory generative models, and on the other linked with motor and frontal areas that produce continuous trajectories and plans of discrete action chunks. The PPC is known to be an associative region that integrates information from multiple sensory modalities and encodes visuomotor transformations—e.g., area V6A is thought to encode object affordances during reaching and grasping tasks ([Fattori et al., 2017](#); [Filippini et al., 2017](#)). Moreover, evidence suggests that the PPC encodes multiple goals in parallel during sequences of actions, even when there is a considerable delay between different goals ([Baldauf et al., 2008](#)).

In short, the agent constantly maintains plausible hypotheses over the causes of its percepts, either bodily states or objects in exteroceptive domains; by manipulating them, the agent dynamically constructs representations of future states, i.e., intentions, which in turn act as priors over the current belief. Thus, if the job of the sensory pathways is to compute sensory-level predictions, we hypothesize that higher levels of the sensorimotor control hierarchy integrate into the PPC previous states of belief with flexible intentions, each predicting the next plausible belief state.

### 3.2. Dynamic goal-directed behavior in Active Inference

For a more formal definition, we assume that the neural system perceives the environment and receives motor feedback through  $J$  noisy sensors  $\mathbf{S}$  comprising multiple domains (most critically, proprioceptive and visual). Under the VB and Gaussian approximations of the recognition density, we also assume that the nervous system operates on beliefs  $\boldsymbol{\mu} \in \mathbb{R}^M$  that define an abstract internal representation of the world. Furthermore, we assume that the agent maintains generalized coordinates up to the 1st order resulting from free energy minimization in the generalized belief space  $\tilde{\boldsymbol{\mu}} = \{\boldsymbol{\mu}, \boldsymbol{\mu}'\}$ .

We then define *intentions*  $\mathbf{h}_k$  as predictions of target goal states over the current belief  $\boldsymbol{\mu}$  computed with the help of  $K$  functions  $\mathbf{i}_k(\boldsymbol{\mu}) \in \mathbb{R}^M$ . Although both belief and intentions could be abstract representations of the world—comprising states in extrinsic and intrinsic coordinates—we assume a simpler scenario in which the intentions operate on beliefs in a common intrinsic motor-related domain, e.g., the joint angles space. As explained before, we assume that there are two conceptually different components in both the belief  $\boldsymbol{\mu}$  and the output of the intention functions  $\mathbf{i}_k$ . The first component could represent the bodily states and serve to drive actions, while the second one could represent the state of other objects—mostly targets to interact with—which can be internally

encoded in the joint angles space as well (the reason for this particular encoding will be clear later). These targets could be observed, but they could also be imagined or set by higher-level cognitive control frontal areas such as the PFC or PMd ([Genovesio et al., 2012](#); [Stoianov et al., 2016](#)).

For the sake of notational simplicity, we group all intentions into a single matrix  $\mathbf{H} \in \mathbb{R}^{M \times K}$ :

$$\mathbf{H} = \mathbf{i}(\boldsymbol{\mu}) = [\mathbf{i}_0(\boldsymbol{\mu}) \dots \mathbf{i}_K(\boldsymbol{\mu})] = [\mathbf{h}_0 \dots \mathbf{h}_K] \quad (20)$$

Intention prediction errors  $\mathbf{e}_{i_k}$  are then defined as the difference between the current belief and every intention:

$$\mathbf{E}_i = \mathbf{H} - \boldsymbol{\mu} = [\mathbf{e}_{i_0} \dots \mathbf{e}_{i_K}] \quad (21)$$

In turn, sensory predictions are produced by a set of generative models  $\mathbf{g}_j$ , one for each sensory modality. We group the predictions into a prediction matrix  $\mathbf{P}$ :

$$\mathbf{P} = \mathbf{g}(\boldsymbol{\mu}) = \begin{bmatrix} \mathbf{g}_0(\boldsymbol{\mu}) \\ \vdots \\ \mathbf{g}_j(\boldsymbol{\mu}) \end{bmatrix} = \begin{bmatrix} \mathbf{p}_0 \\ \vdots \\ \mathbf{p}_j \end{bmatrix} \quad (22)$$

Note that each term  $\mathbf{p}_j$  is a multidimensional sensory-level representation that provides predictions for a particular sensory domain, with its own dimensionality, which we group into a single quantity for notational simplicity. Sensory prediction errors  $\mathbf{e}_{s_j}$  are then computed as the difference between sensations from each domain and the corresponding sensory-level predictions:

$$\mathcal{E}_s = \mathbf{S} - \mathbf{P} = \begin{bmatrix} \mathbf{e}_{s_0} \\ \vdots \\ \mathbf{e}_{s_j} \end{bmatrix} \quad (23)$$

Under the assumption of independence among intentions and sensations, we can factorize the joint probability of the generative model into a product of distributions for each sensory modality and intention, which expands as follows:

$$p(\tilde{\boldsymbol{\mu}}, \mathbf{s}) = p(\boldsymbol{\mu}) \prod_k^K p(\boldsymbol{\mu}'_k | \boldsymbol{\mu}) \cdot \prod_j^J p(\mathbf{s}_j | \boldsymbol{\mu}) \quad (24)$$

In the following, we will not consider the prior probability over the 0th order belief  $p(\boldsymbol{\mu})$ . The other probability distributions are assumed to be Gaussian:

$$\begin{aligned} p(\boldsymbol{\mu}'_k | \boldsymbol{\mu}) &= \mathcal{N}(\boldsymbol{\mu}'_k | \mathbf{f}_k(\boldsymbol{\mu}), \boldsymbol{\gamma}_k^{-1}) \\ p(\mathbf{s}_j | \boldsymbol{\mu}) &= \mathcal{N}(\mathbf{s}_j | \mathbf{g}_j(\boldsymbol{\mu}), \boldsymbol{\pi}_j^{-1}) \end{aligned} \quad (25)$$

Where  $\boldsymbol{\gamma}_k$  and  $\boldsymbol{\pi}_j$  are, respectively, the precisions of intention  $k$  and sensor  $j$ . Here,  $\boldsymbol{\mu}'_k$  and  $\mathbf{f}_k$  correspond to the  $k$ th component of the 1st order dynamics function:

$$\mathbf{f}_k(\boldsymbol{\mu}) = \lambda \mathbf{e}_{i_k} + \mathbf{w}_{\mu_k} \quad (26)$$

Where  $\lambda$  is the gain of intention prediction errors  $\mathbf{E}_i$ . Note that the goal states are embedded into these functions, acting as

belief-level attractors for each intention, so that the agent expects to be pulled toward target states with a velocity proportional to the error. Although the generalized belief allows encoding information about the dynamics of the true generative process, in the simple case delineated the agent does not have any such prior. For example, the agent does not know the trajectory of a moving target in advance (whose prior, in a more realistic scenario, would be present and acquired through learning of past experiences) and will update the belief only relying on the incoming sensory information. Nevertheless, the agent maintains (false) expectations about target dynamics, and it is indeed the discrepancy between the evolution of the (real) generative process and that of the (internal and biased) generative model that makes it able to implement a goal-directed behavior.

The prediction errors of the dynamics functions can be grouped into a single matrix:

$$\mathcal{E}_\mu = \mu' - \lambda E_i \quad (27)$$

From Equation (14), we can now compute the free energy derivative with respect to the belief:

$$\dot{\mu} = \begin{bmatrix} \dot{\mu} \\ \dot{\mu}' \end{bmatrix} = \mathcal{D}\tilde{\mu} - \partial_{\tilde{\mu}}\mathcal{F} = \begin{bmatrix} \mu' + \mathbf{G}^T(\mathbf{\Pi} \odot \mathcal{E}_s) + (\mathbf{F} \odot \mathcal{E}_\mu)\mathbf{\Gamma}^T \\ -\mathcal{E}_\mu \mathbf{\Gamma}^T \end{bmatrix} \quad (28)$$

Here,  $\odot$  is the element-wise product,  $\mathbf{G}$  and  $\mathbf{F}$  enclose the gradients of all sensory generative models and dynamics functions, while  $\mathbf{\Pi}$  and  $\mathbf{\Gamma}$  comprise all sensory and intention precisions:

$$\mathbf{G} = \frac{\partial \mathbf{g}}{\partial \mu} = \begin{bmatrix} \partial g_0 \\ \vdots \\ \partial g_J \end{bmatrix} \quad \mathbf{\Pi} = \begin{bmatrix} \pi_0 \\ \vdots \\ \pi_J \end{bmatrix} \\ \mathbf{F} = \frac{\partial \mathbf{f}}{\partial \mu} = \begin{bmatrix} \partial f_0 & \dots & \partial f_K \end{bmatrix} \quad \mathbf{\Gamma} = \begin{bmatrix} \gamma_0 & \dots & \gamma_K \end{bmatrix} \quad (29)$$

In the following, we will neglect the backward error in the 0th order of Equation (28) since it has a much smaller impact on the overall dynamics, and treat as the actual attractor force the forward error at the 1st order:

$$\dot{\mu} \approx \begin{bmatrix} \mu' + \mathbf{G}^T(\mathbf{\Pi} \odot \mathcal{E}_s) \\ -\mathcal{E}_\mu \mathbf{\Gamma}^T \end{bmatrix} = \begin{bmatrix} \mu' + \epsilon_s \\ \epsilon_i \end{bmatrix} \quad (30)$$

Where  $\epsilon_s$  and  $\epsilon_i$ , respectively, stand for the contributions (in the belief domain) of precision-weighted sensory and intention prediction errors. Considering the 1st order forward error as attractive force instead of the 0th order backward error results in simpler computations since there is no gradient of the dynamics functions to be considered. Further studies are however needed to understand the relationships between these two forces in goal-directed behavior. We can interpret  $\gamma_k$  as a quantity that determines the relative attractor gain of intention  $k$ , so that intentions with greater strength have a more significant impact on the overall update direction; these gains could also be modulated by projections from higher-levels areas applying cognitive control. In turn,  $\pi_j$  corresponds to the confidence about each sensory modality  $j$ , so that the agent relies more on sensors with higher strength.

Similarly, we can compute control signals by minimizing the free energy with respect to the actions, expressing the mapping from sensations to actions by:

$$\frac{\partial s}{\partial \mathbf{a}} = \frac{\partial \mu}{\partial \mathbf{a}} \cdot \mathbf{G} \\ \dot{\mathbf{a}} = -\partial_{\mathbf{a}}\mathcal{F} = -\partial_{\mathbf{a}}\mu^T \epsilon_s \quad (31)$$

Where  $\partial_{\mathbf{a}}\mu$  is an inverse model from belief to actions. If motor signals are defined in terms of joint velocities, we can decompose and approximate the inverse model as follows:

$$\partial_{\mathbf{a}}\mu = \frac{\partial \theta}{\partial \mathbf{a}} \cdot \frac{\partial \mu}{\partial \theta} = \frac{\partial \mathbf{g}_p}{\partial \mathbf{a}} \cdot \frac{\partial \mu}{\partial \mathbf{g}_p} = \Delta_t \mathbf{G}_p^{-1} \quad (32)$$

Where  $\theta$  are the joint angles, the subscript  $p$  indicates the proprioceptive contribution and we approximated  $\partial_{\mathbf{a}}\mathbf{g}_p$  by a time constant  $\Delta_t$  (Oliver et al., 2019). If we assume that the belief over hidden states is encoded in joint angles, the computation of the inverse model may be as simple as finding the pseudoinverse of a matrix. However, if the belief is specified in a more generic reference frame and the proprioceptive generative model is a non-linear function, it could be harder to compute the corresponding gradient, causing additional control problems like temporal delays on sensory signals (Friston, 2011). Alternatively, we can consider a motor control driven only by proprioceptive predictions, so that the control signal is already in the correct domain and may be achieved through simple reflex arc pathways (Adams et al., 2013; Versteeg et al., 2021). In this case, all that is needed is a mapping from proprioceptive predictions to actions:

$$\dot{\mathbf{a}} = -\partial_{\mathbf{a}}\mathcal{F}_p = -\Delta_t \cdot \pi_p \mathbf{e}_p \quad (33)$$

Expressing in Equation (31) the mapping from sensations to actions by the product of the inverse model  $\partial_{\mathbf{a}}\mu$  and the gradient of the generative models allows the control signals to be defined in terms of the weighted sensory contribution  $\epsilon_s$ , already computed during the inference process. Such an approach may have some computational advantages (as will be explained later), but it is unlikely to be implemented in the nervous system as control signals are supposed to convey predictions and not prediction errors (Adams et al., 2013).

Algorithm 1 outlines a schematic description of the flow of dynamic computations. For simplicity, we used the term "intention" also when describing the dynamics functions and their precisions, but one has to keep in mind the difference between the intention prediction errors  $E_i$ , which directly encode the direction toward target states, and the dynamics prediction errors  $\mathcal{E}_\mu$ , which arise from the derivation of the corresponding probability distributions.

### 3.3. Neural implementation

Figure 2 shows a schematic neuronal representation of the proposed agent, which further extends earlier perceptual inference schemes (Bogacz, 2017) to full-blown Active Inference. In this simple model, the intentions consist of a single layer with two neurons, and the goal states are implicitly defined in the dynamics functions; however, in a realistic setting the latter would be



composed of networks of neurons where these states are explicitly encoded, and non-linear functions could also be used to achieve more advanced behaviors. Note also that intentions  $h_k$  and sensory generative models  $g_j$  are all part of the same architecture, the only difference being the location in the cortical hierarchy.

Low-level prediction errors for each sensory modality are represented by neurons whose dynamics depends on both observations and predictions of the sensory generative models:

$$\dot{\epsilon}_{s_j} = s_j - g_j(\mu) - \frac{\epsilon_{s_j}}{\pi_j} \tag{34}$$

Upon convergence of the neural activity, that is,  $\dot{\epsilon}_{s_j} = 0$ , we obtain the prediction error computation derived above. In turn, the internal activity of neurons corresponding to high-level prediction errors is obtained by subtracting the generated dynamics function from the 1st order belief:

$$\dot{\epsilon}_{\mu_k} = \mu' - f_k(\mu) - \frac{\epsilon_{\mu_k}}{\gamma_k} \tag{35}$$

**Input:**  $S, i, g, \partial_a \mu, \Gamma, \Pi, \lambda, \Delta_t$

```

1:  $\mu, \mu', \mu'' \leftarrow \text{InitializeBelief}()$ 
2: while  $t < T$  do
3:    $H \leftarrow i(\mu)$   $\triangleright$  Intentions and sensory predictions
4:    $P \leftarrow g(\mu)$ 
5:    $\mathcal{E}_\mu \leftarrow \mu' - \lambda(H - \mu)$   $\triangleright$  Prediction errors
6:    $\mathcal{E}_s \leftarrow S - P$ 
7:    $\epsilon_i \leftarrow -\mathcal{E}_\mu \Gamma^T$   $\triangleright$  Precision-weighted contributions
8:    $\epsilon_s \leftarrow G^T(\Pi \odot \mathcal{E}_s)$ 
9:    $\dot{\mu} \leftarrow \mu' + \epsilon_s$   $\triangleright$  Belief and action update
10:   $\dot{\mu}' \leftarrow \mu'' + \epsilon_i$ 
11:   $\dot{a} \leftarrow -\partial_a \mu \cdot \epsilon_s$ 
12:   $\ddot{\mu} \leftarrow \ddot{\mu} + \Delta_t \dot{\mu}$   $\triangleright$  Gradient descent
13:   $a \leftarrow a + \Delta_t \dot{a}$ 
14: end while
    
```

Algorithm 1. Active Inference agent with flexible intentions.

Having received information coming from the top and bottom of the hierarchy, the belief is updated by integrating every signal:

$$\dot{\mu} = \sum_j^J \partial g_j \epsilon_{s_j} + \sum_k^K \partial f_k \epsilon_{\mu_k} \tag{36}$$

Which parallels the update formula derived above (Equation 28). Correspondingly, the 1st order component of the belief is updated as follows:

$$\dot{\mu}' = - \sum_k^K \epsilon_{\mu_k} \tag{37}$$

The belief is thus constantly pushed toward a direction that matches sensations on one side and intentions on the other. We adopted the idea that the slow-varying precisions are encoded as synaptic strengths (Bogacz, 2017), but alternative views consider them as gains of superficial pyramidal neurons (Bastos et al., 2012). In any case, they could be dynamically optimized during inference in a direction that minimizes free energy—e.g., if a sensory modality does not help predict sensations, its weight will decrease. This is also true for the intention weights: by dynamically changing during the movement, they can act as modulatory signals that select the best intention to realize at every moment, which can be useful for solving simultaneous or sequential tasks. Nonetheless, the distinction is purely conceptual as the agent does not discriminate between modulating a future intention or increasing the confidence of a sensory signal. At the belief level, every element just follows the rules of free energy minimization.

### 4. Method

To demonstrate the feasibility of the approach and its capacity to successfully implement goal-directed behavior in dynamic environments, we simulated an agent consisting of an actuated upper limb with visual and proprioceptive sensors that allow it

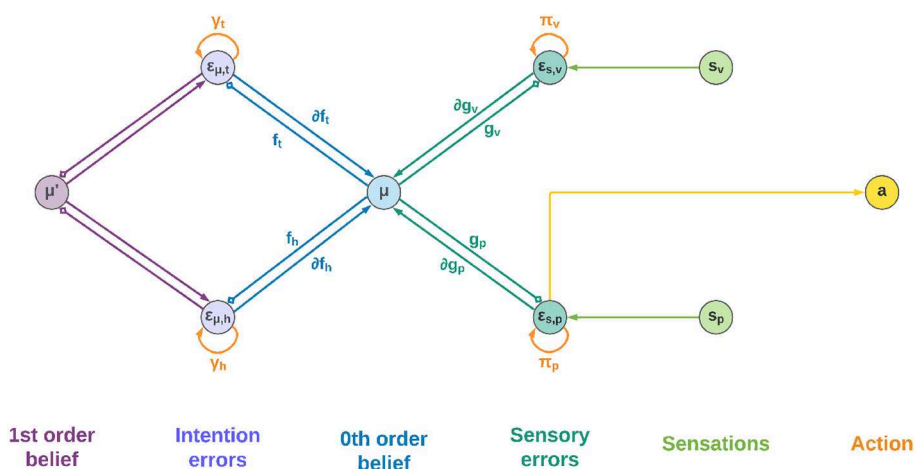
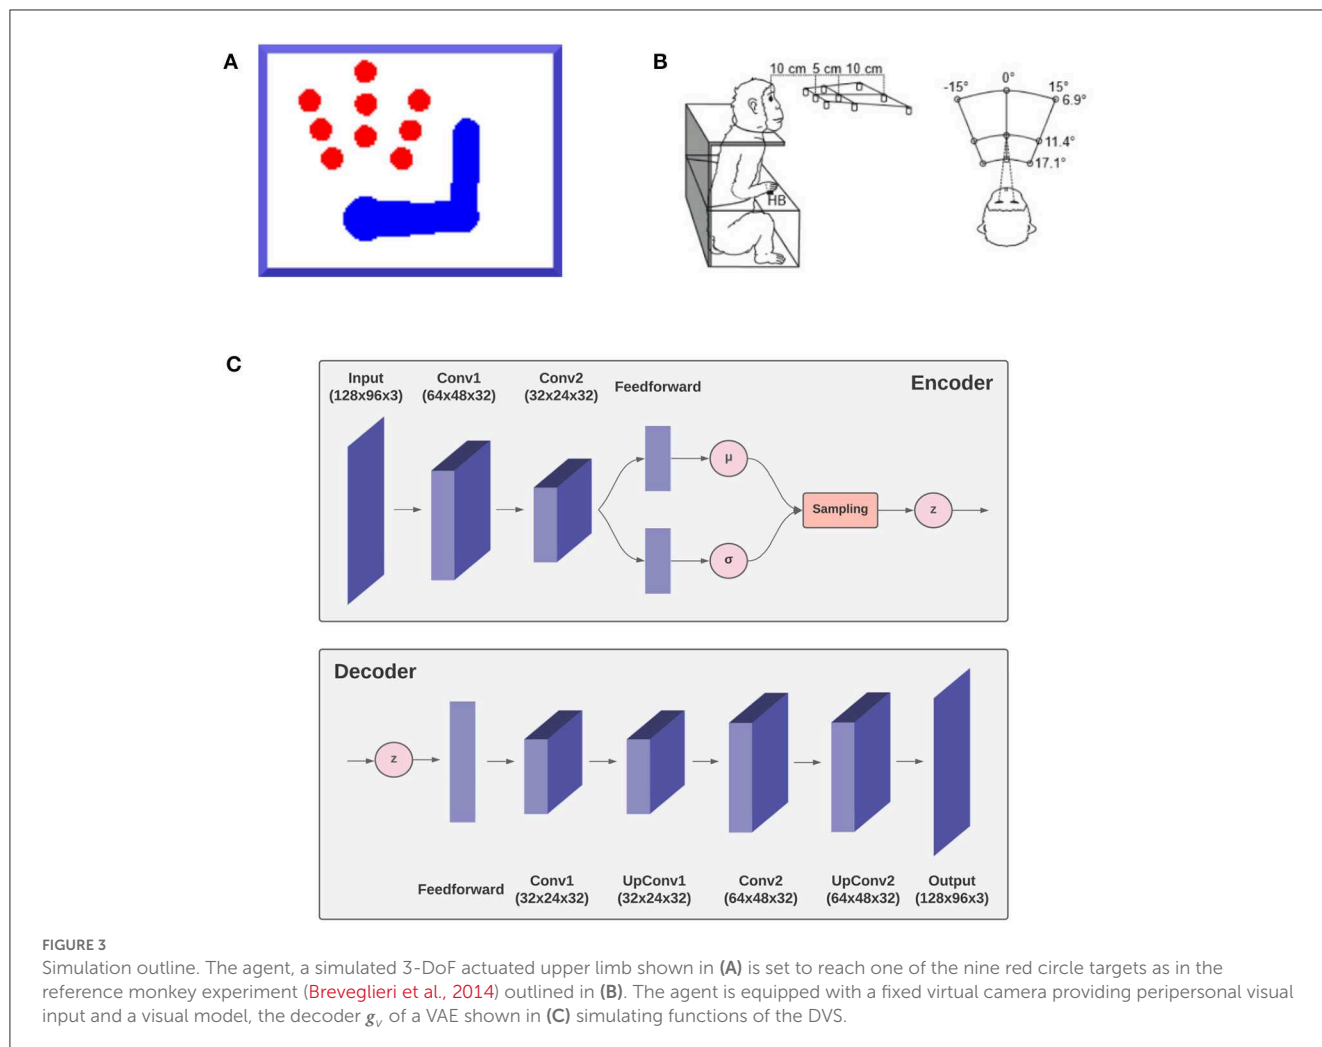


FIGURE 2 Neuronal representation with two intentions. Small squares stand for inhibitory connections.



to perceive and reach static and moving targets within its reach.<sup>1</sup> Figure 3A shows the size and position of the targets, as well as limb size and a sample posture. Since the focus here was on theoretical aspects, we simulated just a coarse 3-DoF limb model moving on a 2D plane. However, the approach easily generalizes to a more elaborated limb model and 3D movements. In the following, we describe the agent, the specific implementation, and the simulated task. Then, in the Results section we assess the agent's perceptual and motor control capabilities in static and dynamic conditions. The static condition simulated a typical monkey reaching task of peripersonal targets as in Figure 3 (Breviglieri et al., 2014). In turn, the dynamic condition involved a moving target that the agent had to track continuously.

#### 4.1. Delayed reaching task

The primary testbed task is a simplified version of a delayed reaching monkey task in which a static target must be reached with a movement that can only start after a delay period (Breviglieri et al., 2014). Delayed actions are used to separately

investigate neural processes related to action preparation (e.g., perception and planning) and execution in goal-directed behavior, and are thus useful to analyze the two main computational components of free energy minimization, namely, perceptual and active inference, which otherwise work in parallel. Delayed reaching could be implemented using various approaches: the update of the posture component of the belief dynamics could be blocked by setting the intention gain  $\lambda$  to zero during inference (implemented here): in this way, there are no active intentions and the belief only follows sensory information. Alternatively, action execution could be temporarily suspended by setting to zero the proprioceptive precision, so that the agent still produces proprioceptive predictions but does not trust their prediction errors: in this scenario, the belief dynamics includes a small component directed toward the intention, but the discrepancy produced is not minimized through movement.

Reach trials start with the hand placed on a home button (HB) located in front of the body center (i.e., the “neck”), and the belief is initialized with this configuration. Then one of the 9 possible targets of the reference experiment (Figure 3) is lit red. Follows a delay period of 100 time steps during which the agent is only allowed to perceive the visible target and the limb, and the inference process can only change the belief. After that, the limb is allowed

<sup>1</sup> Python code provided in <https://github.com/priorelli/PACE>.

to move and the joint angles are updated according to Equation (38). As in the reference task, upon target reaching the agent stops for a sufficiently long period, i.e., a total of 300 time steps per trial. After that, the agent reaches back the HB (not analyzed here). The simulation included 100 repetitions per target, i.e., 900 trials in total.

## 4.2. Body

The body consists of a simulated monkey upper limb composed of a moving torso attached to an anchored neck, an upper arm, and a lower arm, as shown in Figure 3. The three moving segments are schematized as rectangles, each with unit mass, while the joints (shoulder, elbow) and the tips (neck, hand) as circles. The proportions of the limb segment and the operating range of the joint angles were derived from monkey data *Macaca mulatta* (Kikuchi and Hamada, 2009). The state of the limb and its dynamics are described by the joint angles  $\theta$  and their first moment  $\dot{\theta}$ . We assume noisy velocity-level motor control, whereby the motor efferents  $\mathbf{a}$  noisily control the first moment of joint angles with zero-centered Gaussian noise:

$$\dot{\theta} = \mathbf{a} + \mathbf{w}_a \quad (38)$$

## 4.3. Sensors

The agent receives information about its proprioceptive state and visual context. Simplified peripersonal visual input  $\mathbf{s}_v$  was provided by a virtual camera that included three 2D color planes, each of them 128 x 96 pixels in size. The location and orientation of the camera were fixed so that the input provided full vision of peripersonal targets and the entire limb in any possible limb state within its operating range. The limb could occlude the target in some configurations.

As in the simulated limb, the motor control system also receives proprioceptive feedback through sensors  $\mathbf{s}_p$ , providing noisy information on the true state of the limb (Tuthill and Azim, 2018; Versteeg et al., 2021). We further assumed that  $\mathbf{s}_p$  provides a noisy reading of the state of all joints only in terms of joint angles, ignoring other proprioceptive signals such as force and stretch (Srinivasan et al., 2021), which the Active Inference framework can natively incorporate.

## 4.4. Belief

We assume that both the orders of the generalized belief  $\tilde{\boldsymbol{\mu}}$  comprise three components: (i) beliefs  $\tilde{\boldsymbol{\mu}}_a$  over arm joint angles, or posture; (ii) beliefs  $\tilde{\boldsymbol{\mu}}_t$  over the target location represented again in the joint angles space—i.e., the posture corresponding to the arm touching the target; and (iii) beliefs  $\tilde{\boldsymbol{\mu}}_h$  over a memorized HB configuration. Thus,  $\boldsymbol{\mu} = [\boldsymbol{\mu}_a, \boldsymbol{\mu}_t, \boldsymbol{\mu}_h]$ . Note that the last two components can be interpreted as *affordances*, allowing the agent to implement interactions in terms of bodily configurations (Pezzulo and Cisek, 2016).

## 4.5. Sensory model

The sensory generative distribution has two components, one for each sensory modality: a simplified proprioceptive model  $\mathbf{g}_p(\boldsymbol{\mu})$  and a full-blown visual model  $\mathbf{g}_v(\boldsymbol{\mu})$ :

$$\mathbf{g}(\boldsymbol{\mu}) = \begin{bmatrix} \mathbf{g}_p(\boldsymbol{\mu}) \\ \mathbf{g}_v(\boldsymbol{\mu}) \end{bmatrix} \quad (39)$$

Since the belief is already in the joint angles domain, we implemented a simple proprioceptive generative model  $\mathbf{g}_p(\boldsymbol{\mu}) = \mathbf{G}_p \boldsymbol{\mu} = \boldsymbol{\mu}_a$ , where  $\mathbf{G}_p$  is a mapping that only extracts the first component of the belief:

$$\mathbf{G}_p = \begin{bmatrix} \mathbf{I} & \mathbf{0} \end{bmatrix} \quad (40)$$

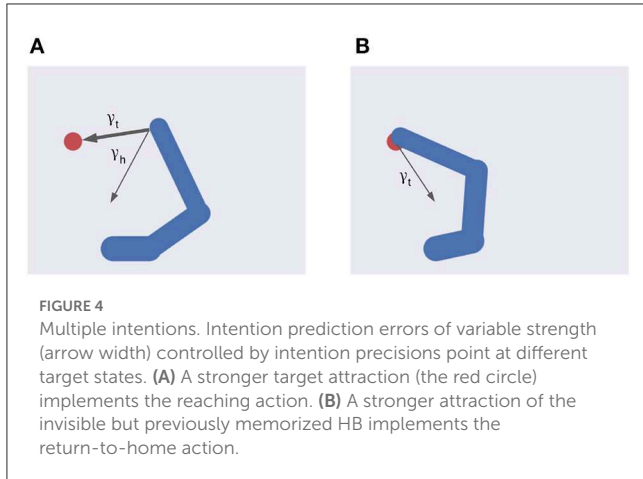
Where  $\mathbf{0}$  and  $\mathbf{I}$  are respectively 3 x 3 zero and identity matrices. Note that  $\mathbf{g}_p(\boldsymbol{\mu})$  could be easily extended to a more complex proprioceptive mapping if the body and/or joint sensors have a more complex structure and the belief has a richer and abstract representation.

In turn, the visual generative model  $\mathbf{g}_v$  is the decoder component of a VAE (see Figure 3C). It consists of one feedforward layer, two transposed convolutional layers, and two standard convolutional layers needed to smooth the output. Its latent space is composed of two elements, representing the joint angles of arm and target (example in Figure 13). The first component is used to generate, in the visual output, an arm with a specific joint configuration, while the second component is used to produce only the image of the target through direct kinematics of every joint angle. The VAE was trained in a supervised manner for 100 epochs on a dataset comprising 20,000 randomly drawn body-target configurations that uniformly spanned the entire operational space, and the corresponding visual images. The target size varied with a radius ranging from 5 to 12 pixels.

The proprioceptive gradient  $\partial \mathbf{g}_p$  simplifies to the mapping  $\mathbf{G}_p$  itself, while the visual component  $\partial \mathbf{g}_v$  is the gradient of the decoder computed by backpropagation. Since the Cartesian position of the target is encoded in joint angles, this gradient implicitly performs a kinematic inversion. Therefore, predictions  $\mathbf{P}$  and prediction errors  $\mathcal{E}_s$  take the form:

$$\mathbf{P} = \begin{bmatrix} \boldsymbol{\mu}_a \\ \mathbf{g}_v(\boldsymbol{\mu}) \end{bmatrix} \quad \mathcal{E}_s = \begin{bmatrix} \mathbf{s}_p - \boldsymbol{\mu}_a \\ \mathbf{s}_v - \mathbf{g}_v(\boldsymbol{\mu}) \end{bmatrix} \quad (41)$$

Note that defining sensory predictions on both proprioceptive and visual sensory domains allows the agent to perform efficient goal-directed behavior also in conditions of visual uncertainty, e.g., due to low visibility. Indeed, since the belief is maintained over time, the agent remembers the last known target position and can thus accomplish reaching tasks also in case of temporarily occluded targets.



## 4.6. Intentions

Stepping on the proposed formalization (Equation 20), we define two specific intentions (Figure 4) as follows:

$$H = \begin{bmatrix} i_t(\mu) & i_h(\mu) \end{bmatrix} = \begin{bmatrix} I_t \mu & I_h \mu \end{bmatrix} = \begin{bmatrix} \mu_t & \mu_h \\ \mu_t & \mu_t \\ \mu_h & \mu_h \end{bmatrix} \quad (42)$$

Here  $h_t = i_t(\mu)$  defines the agent's expectation that the arm belief is equal to the joint configuration corresponding to the target to be reached, and it is implemented as a simple mapping  $I_t$  that sets the first belief component equal to the second one. In turn, the intention  $h_h = i_h(\mu)$  encodes the future belief of the agent that the arm will be at the HB position. The two intention mappings are defined by:

$$I_t = \begin{bmatrix} \mathbf{0} & \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{I} \end{bmatrix} \quad I_h = \begin{bmatrix} \mathbf{0} & \mathbf{0} & \mathbf{I} \\ \mathbf{0} & \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{I} \end{bmatrix} \quad (43)$$

The corresponding intention prediction errors are then:

$$E_i = \begin{bmatrix} e_{i_t} & e_{i_h} \end{bmatrix} = \begin{bmatrix} h_t - \mu & h_h - \mu \end{bmatrix} = \begin{bmatrix} \mu_t - \mu_a & \mu_h - \mu_a \\ \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \quad (44)$$

These errors provide an update direction respectively toward the target and HB joint angles. As there is no intention to move the target or the HB, the second and third components of the prediction errors will be zero.

## 4.7. Precisions

Free energy minimization and Predictive Coding in general heavily depend on precisions modulation. To investigate their role, we parameterized the relative precisions of each intention and sensory domain with parameters  $\alpha$  and  $\beta$  as follows:

$$\Pi = \begin{bmatrix} \pi_p \\ \pi_v \end{bmatrix} = \begin{bmatrix} 1 - \alpha \\ \alpha \end{bmatrix} \quad \Gamma = \begin{bmatrix} \gamma_t & \gamma_h \end{bmatrix} = \begin{bmatrix} 1 - \beta & \beta \end{bmatrix} \quad (45)$$

The parameter  $\alpha$  controls the relative strength of the error update due to proprioception and vision, while the parameter  $\beta$  controls the relative attraction by each intention. With these parameters, the sensory and intention weighted contributions are unpacked as follows:

$$\epsilon_s = (1 - \alpha) \cdot \mathbf{G}_p^T \mathbf{e}_{s_p} + \alpha \cdot \partial \mathbf{g}_v^T \mathbf{e}_{s_v} \quad (46)$$

$$\epsilon_i = -\mu' + \lambda[(1 - \beta) \cdot \mathbf{e}_{i_t} + \beta \cdot \mathbf{e}_{i_h}] \quad (47)$$

Equation (46) shows the balance between visual and proprioceptive information. For example, if  $\alpha = 0$  the agent will only use proprioceptive feedback, while for  $\alpha = 1$  the belief will be updated only relying on visual feedback. Note that these are extreme conditions—e.g., the former may correspond to null visibility—and typical sensory systems provide balanced feedback. In turn, Equation (47) spells out the control of belief attraction. The agent will follow the first intention when  $\beta = 0$ , or the second one when  $\beta = 1$  (Figure 4). Note that the introduction of a possible competing reach movement creates a conflict among intentions aiming to fulfill opposing goals (e.g., for intermediate values of  $\beta$ ) while the agent can physically realize only one of them at a time (Figure 4). Thus, we assume that the control of intention selection is realized through mutual inhibition and higher-level bias. Finally, the parameter  $\lambda$  controls the overall attractor magnitude (see also Equation 26).

We can also use the precision parameter  $\alpha$  to manipulate the strength of the free energy derivative with respect to the actions as follows:

$$\dot{a} = -\Delta_t(1 - \alpha) \cdot \mathbf{e}_{s_p} \quad (48)$$

Note that by increasing  $\alpha$ —i.e., more reliability on vision—the magnitude of the belief update remains constant, while action updates decrease because the agent becomes less confident about its proprioceptive information. Also, one could differentially investigate the effect of precision strength on belief and action by directly manipulating the precisions—e.g., visual precision  $\pi_v$  may include different components that follow the belief structure:

$$\pi_v = [\alpha, \pi_{v_t}, \pi_{v_h}] \quad (49)$$

Where we used the parameter  $\alpha$  only for the arm belief. For example, when  $\alpha = 0$  and  $\pi_{v_t} > 0$ , the target belief is updated using visual input while the arm moves only using proprioception, a scenario that emulates movement in darkness with a lit target.

## 5. Results

In the following, we assess the capacities of the intention-driven Active Inference agent to perceive and perform goal-directed actions in reaching tasks with static and dynamic visual targets. The main testbed task was delayed reaching, but we simulated several other conditions.

Sensorimotor control that implements goal-directed behavior was investigated in various sensory feedback conditions, including pure proprioceptive or mixed visual and proprioceptive, in which the VAE decoder provided support for dynamic estimation of visual targets and bodily states. The latter is the typical condition of

performing reaching actions and allows greater accuracy (Keele and Posner, 1968). In an additional *baseline* (BL) condition, the target was estimated by the decoder, but the movement was performed without visual feedback or proprioceptive noise, to allow comparisons with the typical approach in previous continuous Active Inference studies, e.g., Pio-Lopez et al. (2016). We also investigated the effects of sensory and intention precisions, motor control type, and movement onset policy. Finally, we analyzed the visual model and the nature of its gradients to provide critical information about the causes of the observed motor behavior.

Action performance was assessed with the help of several measures: (i) *reach accuracy*: success in approaching the target within 10 pixels of its center, i.e., the hand touching the target; (ii) *reach error*:  $L^2$  hand-target distance at the end of the trial; (iii) *reach stability*: standard deviation of  $L^2$  hand-target distance during the period from target reach to the end of the trial, in successful trials; (iv) *reach time*: number of time steps needed to reach the target in successful trials. We also assessed target perception through analog measures based on the  $L^2$  distance between the target location and its estimate transformed from joint angles into visual position by applying the geometric (forward) model. Specifically, we defined the following measures: (v) *perception accuracy*: success in estimating the target location within 10 pixels; (vi) *perception error*:  $L^2$  distance between the true and estimated target position at the end of the trial; (vii) *perception stability*: standard deviation of the  $L^2$  distance between the target position and its estimation during the period starting from successful estimation until end of the trial; (viii) *perception time*: number of time steps needed to successfully estimate the target position.

Figure 5 illustrates key points of the delayed reaching task. During the delay period (Figures 5A–C), the posture does not change since the joint angles only follow the arm belief, which is kept fixed, while the target belief is attracted by the sensory evidence and gradually shifts toward it. When movement is allowed (Figures 5D–F) by setting  $\lambda > 0$  and  $\beta = 0.1$ , the combined attractor produces a force that moves the arm belief toward the target, generating proprioceptive predictions—therefore motor commands—that let the real arm follow this trajectory. Reaching performance is summarized in Figure 6. Panels A–D show spatial statistics of the final hand location (with the corresponding belief) for each target, separately for reaching with proprioception only or proprioceptive and visual sensory feedback. Descriptive statistics revealed an important benefit of visual feedback (Figures 6E–H), in parallel with classical behavioral observations (Keele and Posner, 1968): reach accuracy was higher (with: 88.28%; without: 83.72%) and both reach stability and arm belief error were considerably better with visual feedback as well (stability: 1.35; error: 1.98px) compared to the condition with only proprioception (stability: 1.78; error: 2.87px).

## 5.1. Precision balance

The effects of sensory feedback led to a further systematic assessment of the effects of sensory and intention precisions  $\alpha$ ,  $\pi_v$ , and  $\lambda$  (see Equations 45, 49). The assessment was carried out following the structure of the delayed reaching task. We varied the

above precisions one at a time, using levels shown on the abscissas in Figure 7, while keeping the non-varied precisions at their default values. Note that  $\alpha = 0$  corresponds to reaching without visual feedback, while the conditions  $\alpha > 0$  may be interpreted as reaching with different levels of arm visibility. We recall that the baseline condition (BL) performs reaching movements without visual feedback and proprioceptive noise, i.e.,  $\alpha = 0$  and  $w_p = 0$ . To obtain a systematic evaluation, each condition was run on a rich set of 1,000 randomly selected targets that covered the entire operational space. Finally, we only considered the target-reaching intention, i.e.,  $\beta = 0$ ; everything else was the same as in the main task.

The results are shown in Figure 7. The panels in the left column show the effect of  $\alpha$  compared to the BL agent with noiseless proprioception. Active Inference with only proprioception (i.e.,  $\alpha = 0$ ) has a lower reach accuracy and higher error, while the best performance is obtained with balanced proprioceptive and visual input, in corroboration with the observations of the basic delayed reaching task. In the latter case, the motor control circuitry continuously integrates all available sensory sources to implement visually-guided behavior (Saunders and Knill, 2003). However, accuracy and stability rapidly decrease for excessively high values of  $\alpha$ , due to the discrepancy in update directions between the belief—which makes use of all available sensory information, including the more precise visual feedback—and action—which in this case relies on excessively noisy proprioception. Furthermore, as in the main experiment, the effects of visual precision are evident in the stability of the arm belief, which gradually improves with increasing values (Figure 8): In addition to the reliability of the visual input, this effect is also a consequence of the smaller action updates due to the reduced proprioceptive precision.

In turn, the panels in the middle column of Figure 7 reveal the effects of the attractor gain  $\lambda$ ; to remind the reader, the greater the gain, the greater the contribution of intention prediction errors in the belief updates. The results show that as the intention gain  $\lambda$  increases reach accuracy generally improves, and the number of time steps needed to reach the target decreases. However, beyond a certain level, the accuracy tends to decrease since the trajectory dynamics becomes unstable; thus, excessively strong action drag is counterproductive to the implementation of smooth movements. Finally, the panels in the right column of Figure 7 show the effects of the target precision  $\pi_v$ , which directly affects the quality of target perception. Note that better performances are generally obtained in terms of accuracy, error, and perception time for values of  $\pi_v$  higher than the arm visual precision, which corresponds to a classical effect of contrast on perception, but also means here that the arm and target beliefs follow different dynamics.

## 5.2. Motor control

We described earlier two different ways of implementing motor control in Active Inference: making use of all sensory information, or proprioception only. The first method requires significantly more computations since the agent needs to know the inverse mapping from every sensory domain to compute the control signals. However, given the assumptions we made,

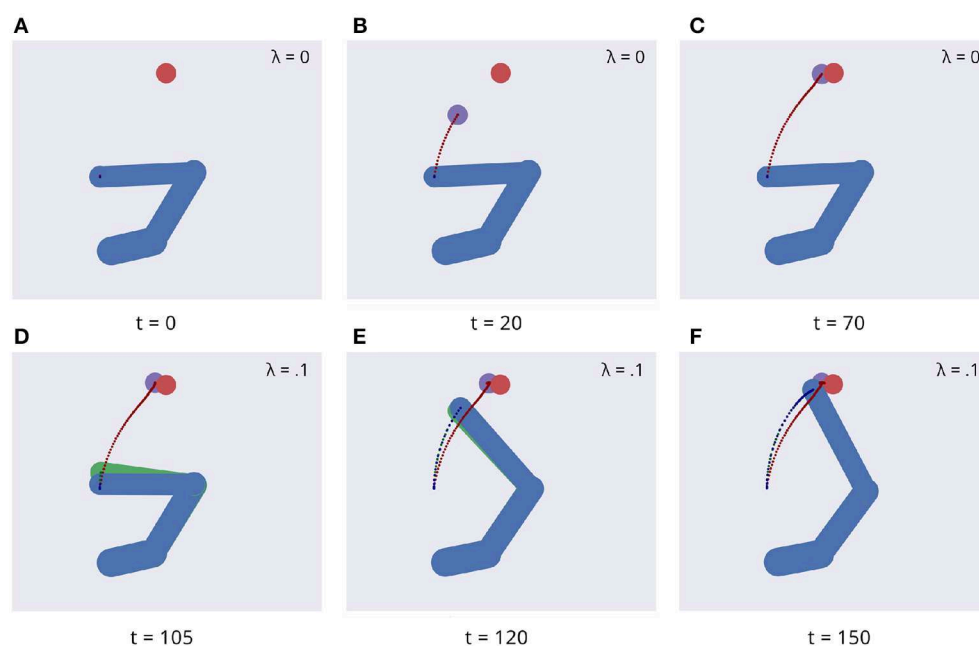


FIGURE 5

Dynamics of the delayed reaching task. At trial onset (A) a visual target (red circle) appears, the arm (in blue) is located on the HB position, and the arm belief (in green) is set at the true arm state. During the delay period, the perceptual inference process gradually drives the target belief (purple circle) toward the real position (B, C). During this phase, the intention gain  $\lambda$  is set to 0, so that movement is inhibited and the arm belief does not change given the unchanged proprioceptive evidence. After movement onset, the arm freely follows its belief (D, E) until they both arrive at the goal state (F).

this approach is potentially more stable because it updates both belief and action with the same information. On the other hand, a pure proprioception-based control mechanism could produce potentially incorrect movements because the motor control commands result from comparing proprioceptive predictions with noisy observations. Greater cost-effectiveness of the second method thus might come at the cost of worsened performances, which we investigate here.

Figure 9 shows a comparison of the two control methods and the BL agent, evaluated under the same conditions we used to investigate precision balance, including 1,000 random targets. Performance was measured in terms of belief and reach stability and reach accuracy. The results reveal, first, that the expected decreased belief stability of the full model with respect to the BL agent (Figure 9C) does not affect hand stability (Figure 9B), although the proprioceptive noise apparently contributed to decreased reach accuracy (Figure 9A). More importantly, the results confirm our expectations that pure proprioception control has considerably lower reach stability caused by incorrect update directions of the motor control signals, resulting in a greater decreased reach accuracy relative to the full model.

### 5.3. Movement onset policy

We also investigated the effects of movement onset using several policies, which differ by the duration of the period of pure perception preceding full Active Inference. One such policy we investigate here is characteristic of actions performed under time pressure, in which movement starts along with perception, i.e.,

action is *immediate*. Another policy that could be considered typical for acting under normal conditions has movement beginning with the satisfaction of a certain perception criterion. This policy *dynamically* deliberates the onset of movement. Various perception criteria could be used: here, the action starts when the norm of the target belief  $\mu_t$  remained below a given threshold (i.e., 0.01) for a certain period (i.e., 5 time steps). These parameters were arbitrarily chosen in consideration of exploratory delayed reaching simulations. Finally, we include the previously used delayed reaching policy in which movement onset is delayed by a *fixed* period (here, 100 time steps, sufficient to obtain a precise target estimation). To obtain systematic observations, each policy was again run on 1,000 randomly selected targets. Measurements included reach and perception accuracy, motor control stability after reach, target perception stability, as well as reach time since the beginning of the trial or after movement onset.

Figure 10 shows the results with the three different policies. Although the reach error is approximately the same in all tested conditions, the agent controlled by the immediate policy reached the target within the lowest total number of time steps: target perception and intention setting were dynamically computed along with movement onset. However, if we consider the total task time, the number of time steps is the highest in this condition, since the arm belief and the arm itself move along with the slow visual target estimation. In turn, if the agent starts the movement when the uncertainty about the target position is already minimized (either in the *dynamic* or *fixed* condition), the movement time decreases, although if added on top of the perception time results in slower actions relative to the immediate movement condition. Finally, we note that target perceptual stability somewhat decreases for dynamic and fixed policies; this somewhat unexpected result is

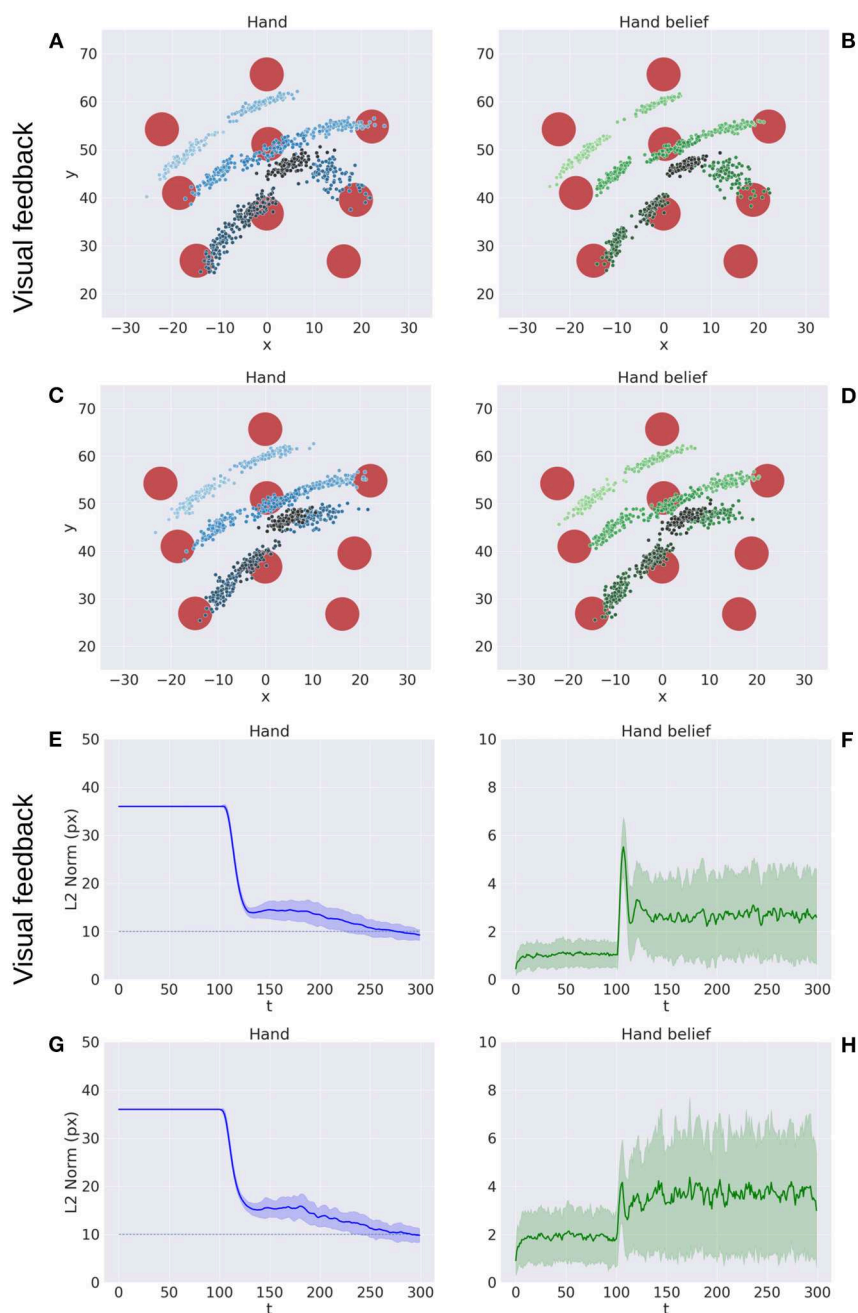


FIGURE 6

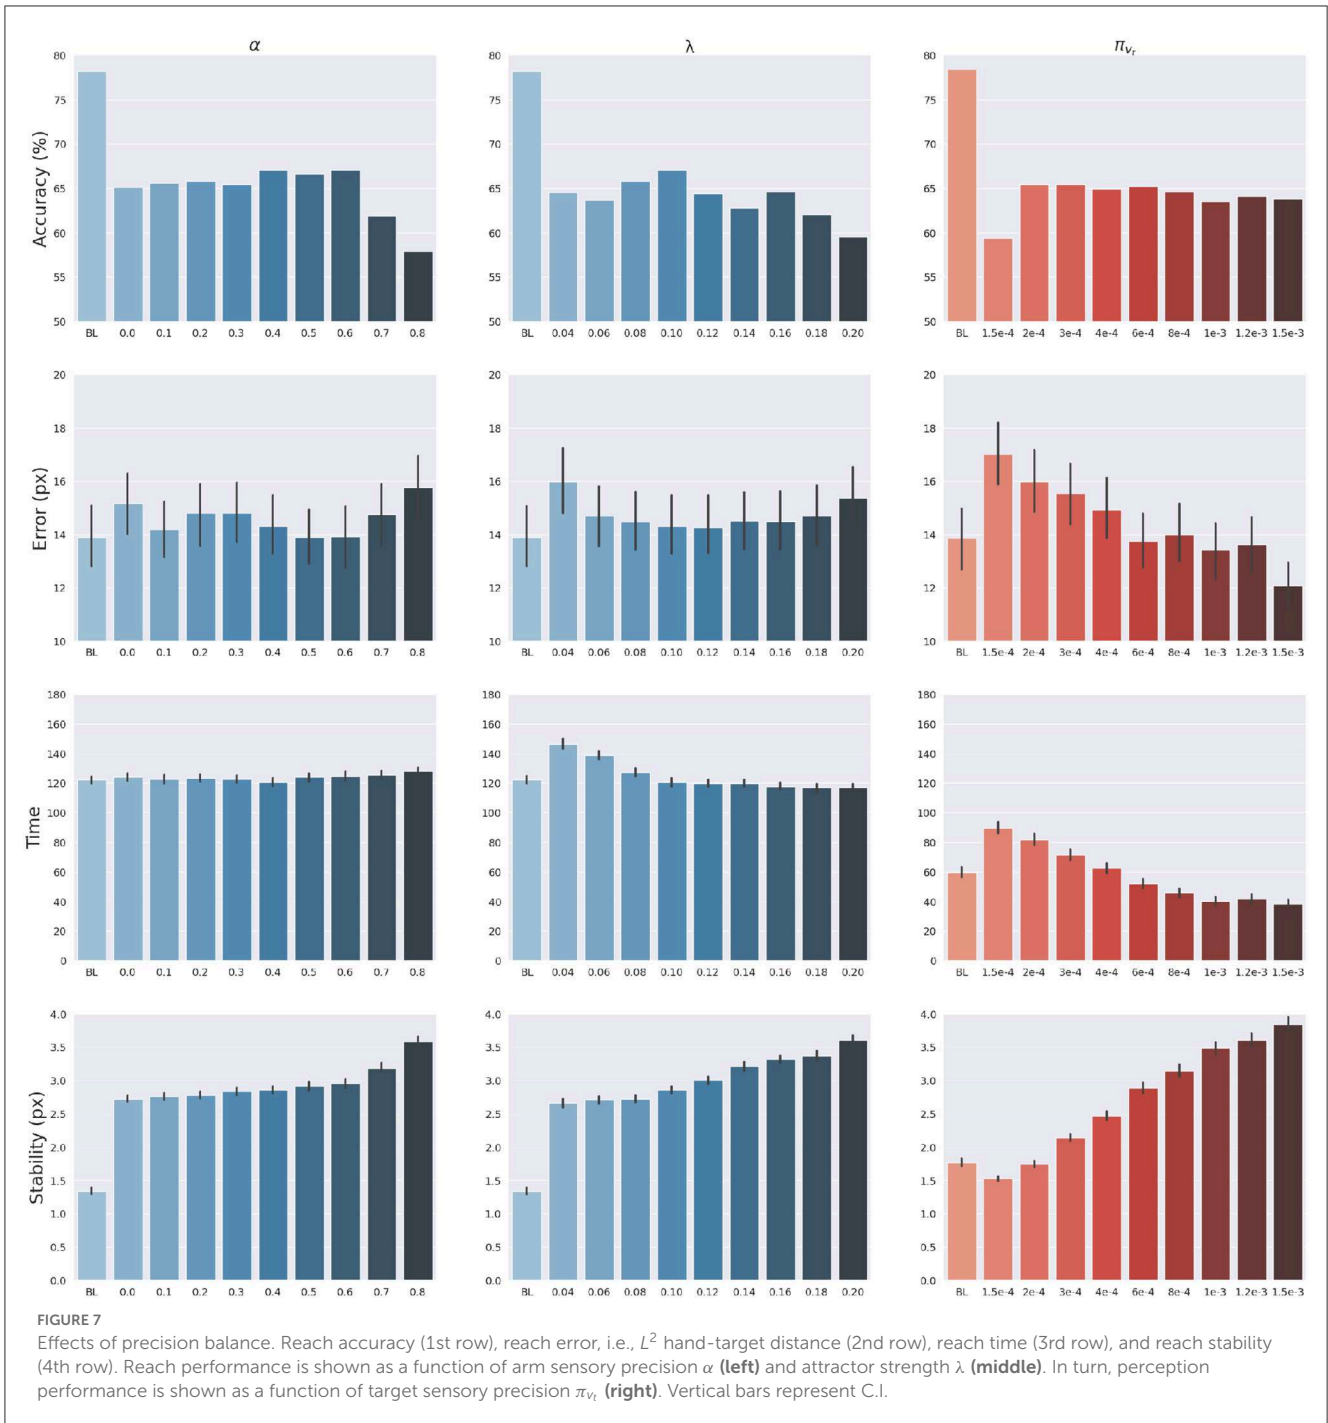
Performance of the delayed reaching task. (A–D) Spatial distribution of hand positions (A, C) and corresponding beliefs (B, D) per target at the end of the reach movements, with (A, B) and without (C, D) visual feedback. Each point represents a trial (100 trials per target). Reach error (E, G) and belief error (F, H) over time, with (E, F) and without (G, H) visual feedback (bands represent C.I.). The reach criterion of the hand-target distance is visualized as a dotted line.  $L^2$  norm for the hand belief is computed by the difference between real and estimated hand positions. Reaching with visual feedback resulted in a more stable hand belief.

encouraging for dynamic target tracking tasks in which immediate movement onset is mandatory.

#### 5.4. Tracking dynamic targets

In a second testbed task, the agent was required to track a smooth-moving target whose initial location was randomly

chosen from the entire operational space. In each trial, the targets received an initial velocity of 0.1px per step in a direction uniformly spanning the 0–360° range. When the target reached a border, its movement was reflected. As in the previous simulations, the belief was initialized at the HB configuration and the trial time limit was 300 time steps. However, for the agent to correctly follow the targets, both the belief and action were dynamically



and continuously inferred in parallel, i.e., without a pure perceptual period.

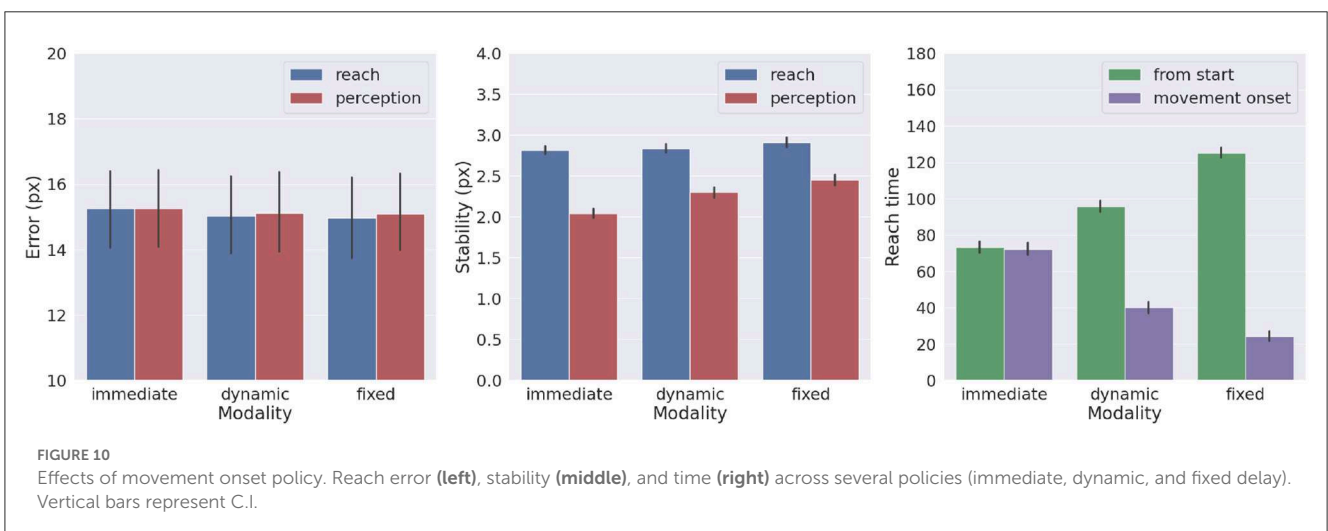
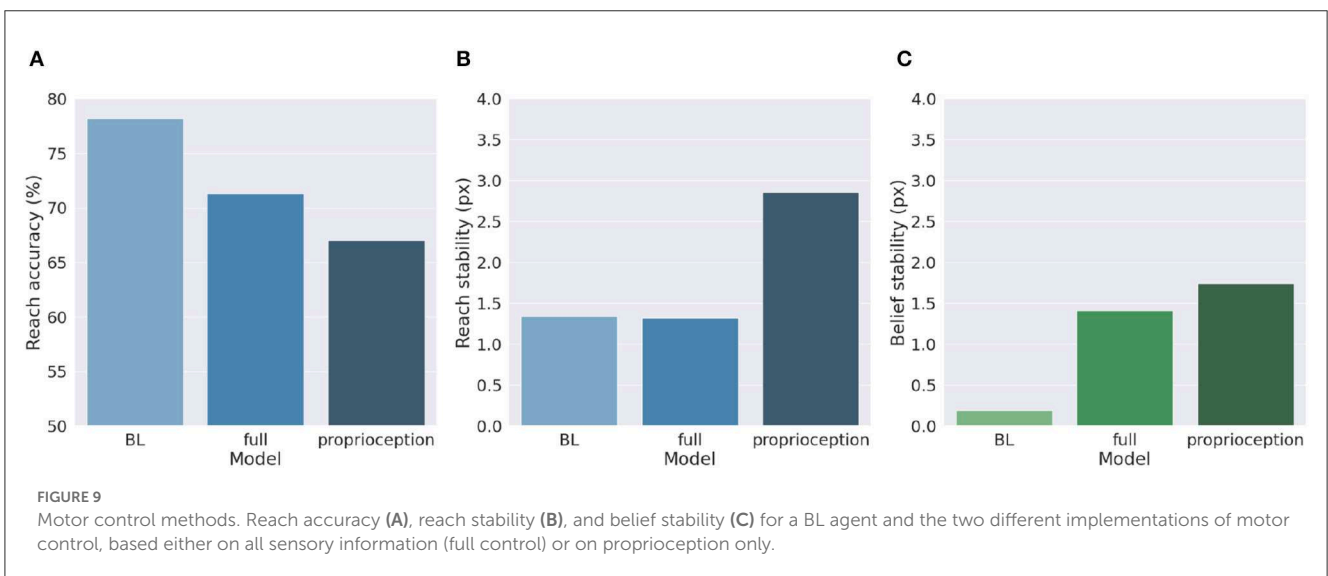
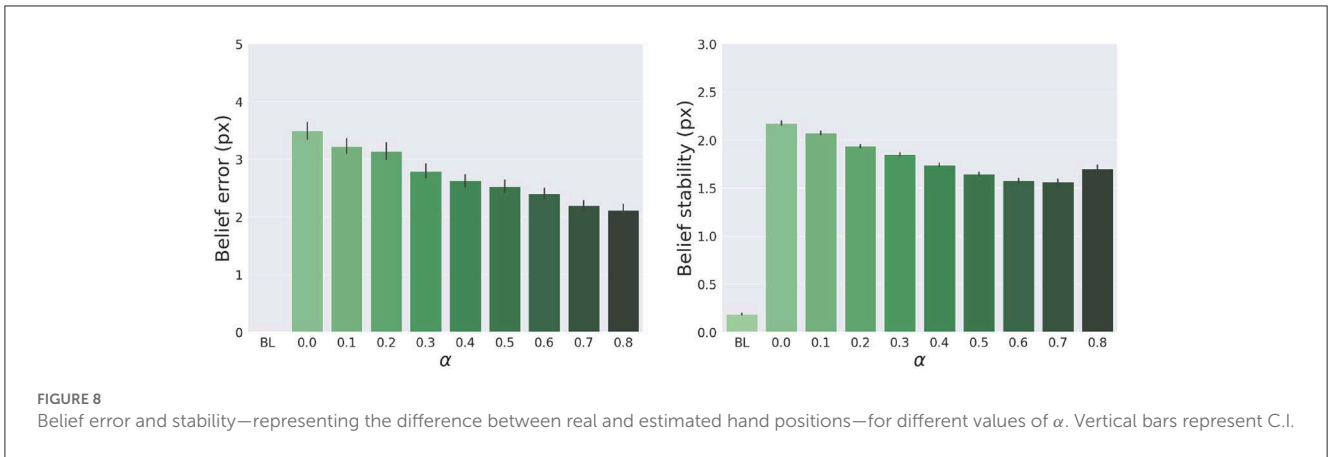
Figure 11 shows the reach trajectory in dynamic target tracking for 10 random trials. The left panel shows the evolution over time of  $L^2$  hand-target distance, while the right panel represents the error between estimated and true target positions. The results suggest that the agent is generally able to correctly and dynamically estimate the beliefs over both target and arm for almost every trial, also in the case of moving targets. In some cases however, mainly when the target is out of reach, it is temporarily or permanently "lost" in terms of its belief, which has also the consequences of losing the

target in terms of reach. Further analysis with a more realistic bodily configuration and visual sensory system—as well as comparisons with actual kinematic data—should provide further insights into the capabilities of Active Inference to perform dynamic reaching.

### 5.5. Free energy minimization

Here we illustrate the dynamics of free energy minimization in delayed reaching, which is at the heart of continuous Active Inference. To that aim, we run 10 new reaching trials with static





and dynamic targets and recorded the free energy derivatives with respect to generalized belief and action.

Figures 12A–F shows the trajectory of the free energy derivatives with respect to the arm and target components during

delayed reaching of a static target; the two columns show the trends for the last two joints, i.e., the arm and forearm segments, that most strongly articulate the reaching action. Note that the gradients of the free energy with respect to the target belief are

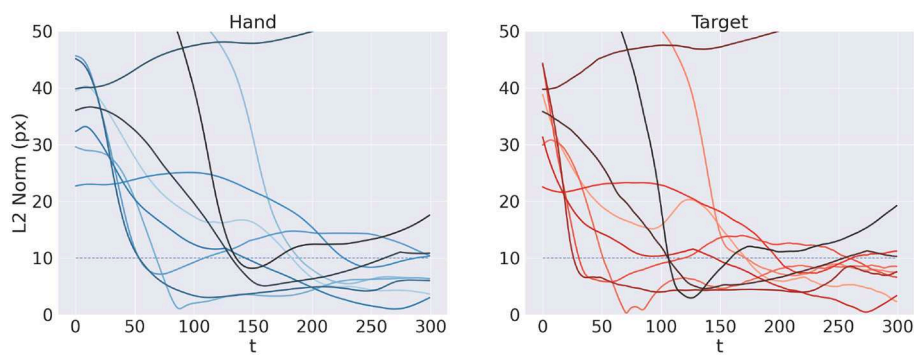


FIGURE 11  
Tracking dynamic targets. Reach (left) and perception (right) error over time for 10 random trials.

rapidly minimized during the initial perceptual phase while the arm gradients remain still. Upon action execution (indicated by a vertical line), the arm gradients rapidly change as well, resulting in updated proprioceptive predictions that drive arm movements. However, arm movements cause changes in the visual scene, resulting in a secondary effect over the just minimized free energy on target belief. Figures 12G–J goes even deeper, showing a direct comparison between  $\hat{\mu}$ ,  $\mu'$ , and the difference  $\hat{\mu} - \mu'$ , on sample static (G–H) and dynamic (I–J) targets. We recall that free energy minimization implies that the two reference frames (the path of the mode  $\hat{\mu}$  and the mode of the path  $\mu'$ ) should overlap at some point in time, when the agent has inferred the correct trajectory of the generalized hidden states. This is crucial especially in dynamic reaching, in which the aim is to capture the instantaneous trajectory of every object in the scene. The decreasing free energy gradients (blue lines) show that this aim is indeed successfully achieved in both static and dynamic tasks.

## 5.6. Visual model analysis

Here, we provide an assessment of the visual model whose performance is critical for accurate visually-guided motor control. To recall, the visual model is implemented with a VAE trained offline to reconstruct images of arm-target configurations such as the one in Figure 13A. A critical VAE parameter is the variance of the recognition (encoder) density  $\Sigma_\phi$  (see Equation 17). We therefore evaluated its effect on perception and action by training several VAEs with different variance levels. VAE performance was assessed on other 10,000 randomly selected configurations that uniformly sampled the space, with a target size of 5 pixels (the default condition for the Active Inference tests).

Most critical was the VAE capacity to generate adequate images of joint arm-target configurations, which we measured with the help of the  $L^2$  norm between visual observations, and VAE-generated images. To provide more insights on the two VAE processes, decoding and encoding, we proceeded as follows: first, *decoding* was assessed by generating images for given body-target states such as that in Figure 13B. The decoded images were compared with the ground truth images produced by applying the geometric model for the same state of the body target (Figure 13A).

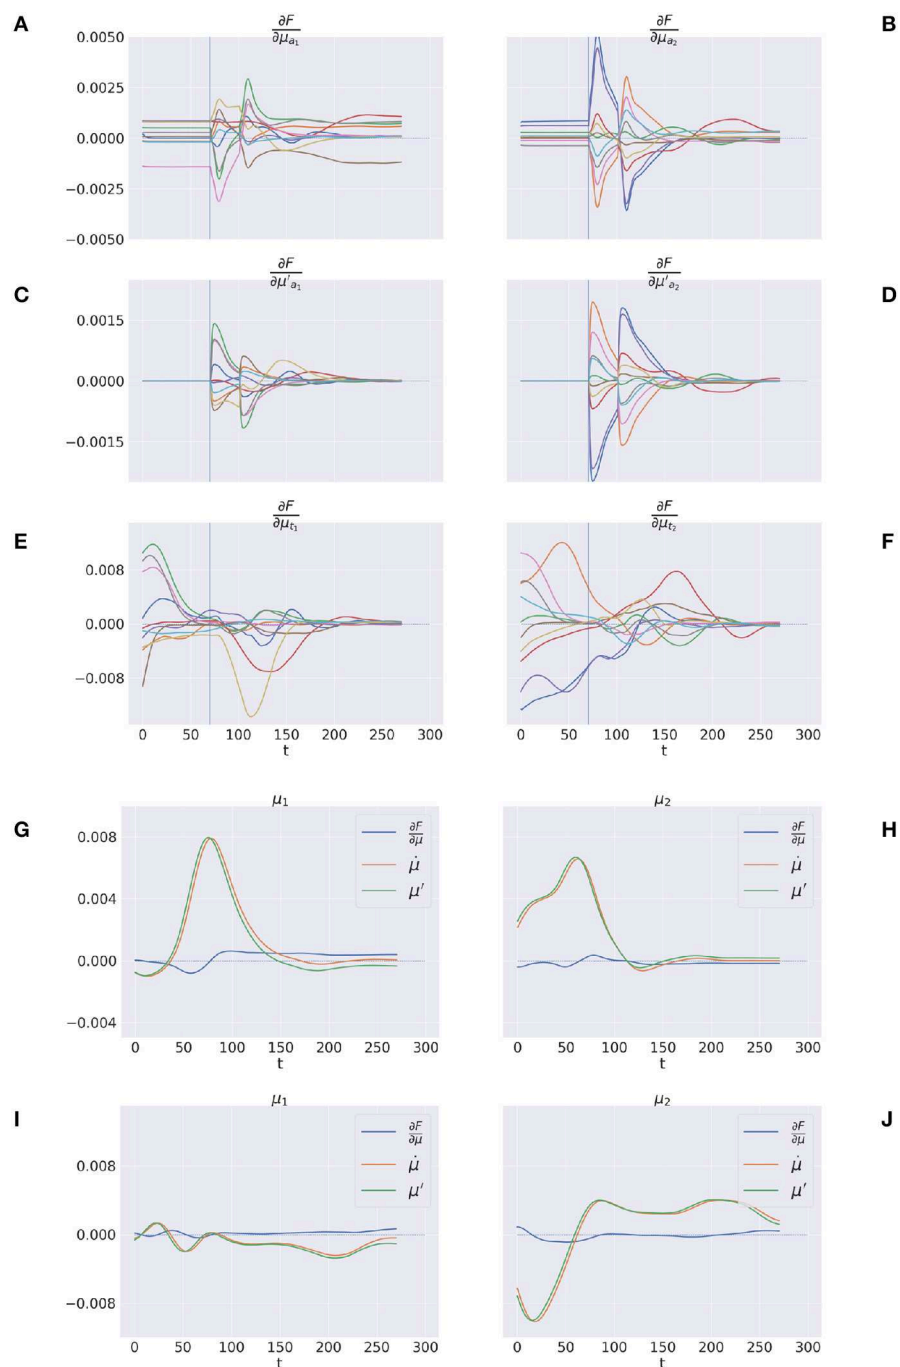
Second, full VAE performance was assessed by computing the average  $L^2$  norm between observed images and their full VAE *reconstruction*, i.e., first encoding and then decoding them (as in Figure 13C). Third, we directly assessed the specific effects of the recognition density variance on Active Inference using the BL condition of the delayed reaching task as a measure.

Figure 13D represents the results of the perceptual assessment tests, showing the  $L^2$  norm between the original and generated images as a function of recognition density variance. As expected, lower variances generally resulted in lower errors with respect to both pure decoding and full encoding-decoding. Surprisingly, however, the accuracy of Active Inference in the reaching task behaved somewhat differently: the best accuracy was obtained not for predictions with low variance, but for intermediate variance levels (Figure 13E). This could be explained by the fact that low-variance images imply highly non-linear gradients that prevent correct gradient descent on free energy. On the other hand, as the variance increases the reconstructed image becomes somewhat blurred, which helps obtain a smoother gradient that correctly drives free energy minimization and therefore improves movement accuracy (more on this in the next section). However, as the variance continues to increase, the reconstructed images become too blurry, degrading both belief inference and motor control.

## 5.7. Visual gradient analysis

To further investigate the cause of the unexpected low variance issue, we analyzed the consistency of the visual gradient  $\partial g_v$  of the decoder for several encoder variance values. To this aim, we computed the gradients for different reference states over the entire operational space.

Figures 14A–C reveals that a decoder with intermediate variance values (green line) causes smaller but smoother gradients, while a too-low variance (orange line) causes sharp peaks near the reference point and even incorrect gradient directions in some cases. Therefore, too low encoder variances seem to make the decoder prone to overfitting, while higher variance values help extract a smoother relationship between irregular multidimensional sensory domains and regular low-dimension causes. Figures 14D–G further illustrates the arm and target



**FIGURE 12** Free energy minimization. (A–F) Free energy derivative with respect to the 0th and 1st order belief for arm (A–D) and target (E, F). (G–J) Comparison between the reference frames of the belief—the path of the mode  $\dot{\mu}$  and the mode of the path  $\mu'$ —for sample static (G, H) and dynamic (I, J) trials. The left/right columns refer to the arm/forearm segments. Trials data are smoothed with a 30 time-step moving average.

gradients relative to a sample reference posture and target location (the result is similar for other configurations) in both Cartesian and polar coordinates; the polar plot shows the two joints most relevant to the reaching action.

The plots reveal greater arm gradients (upper panels) in the vicinity of the target location; in that subspace, the decoder has less uncertainty about which direction to choose to minimize the error. Notably, the gradients tend to compose curved directions,

a characteristic of biological motions. The polar plot provides critical insights into the causes of the circular pattern: the gradients are mostly parallel to the horizontal axis, which corresponds to a movement consisting essentially of pure shoulder rotation. Thus, they provide a strong driving force on the shoulder almost throughout the operational space, while the area in which the elbow is controlled is limited to the vicinity of the target location. These gradients result in a two-phase reaching of static targets in

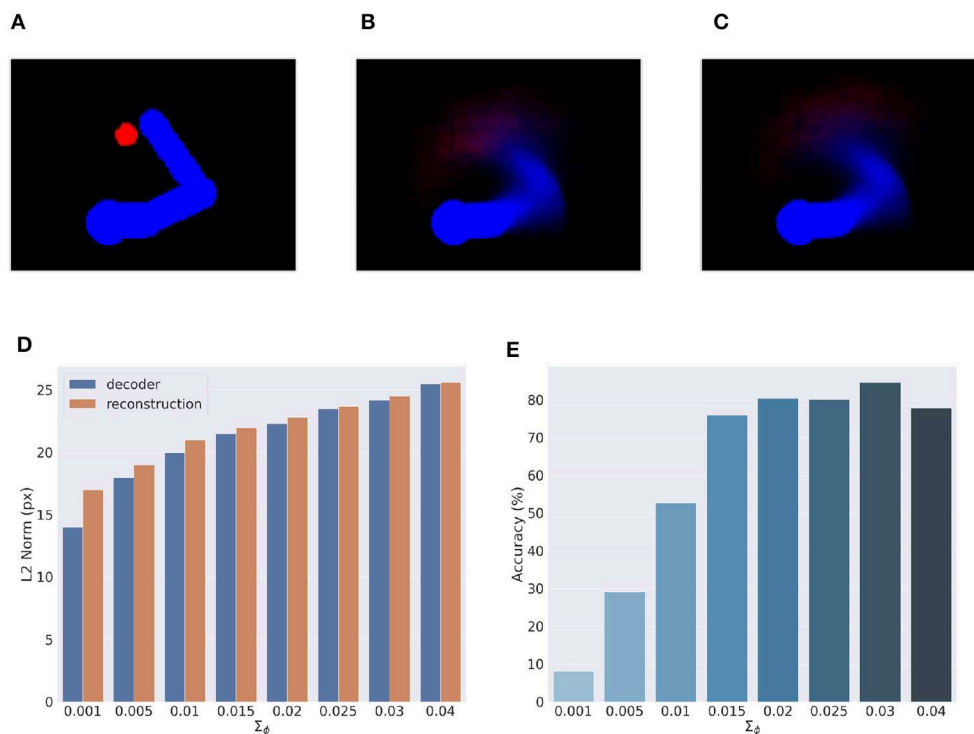


FIGURE 13

Visual model analysis. (A–C) Sample visual observation (A) and its decoding from joint angles (B) and through a complete encoding-decoding process (C). (D, E) Visual model performance. Quality of perception is measured as the  $L^2$  norm between observed and reconstructed images (D) and accuracy of Active Inference (E), as a function of the recognition density variance  $\Sigma_\phi$ .

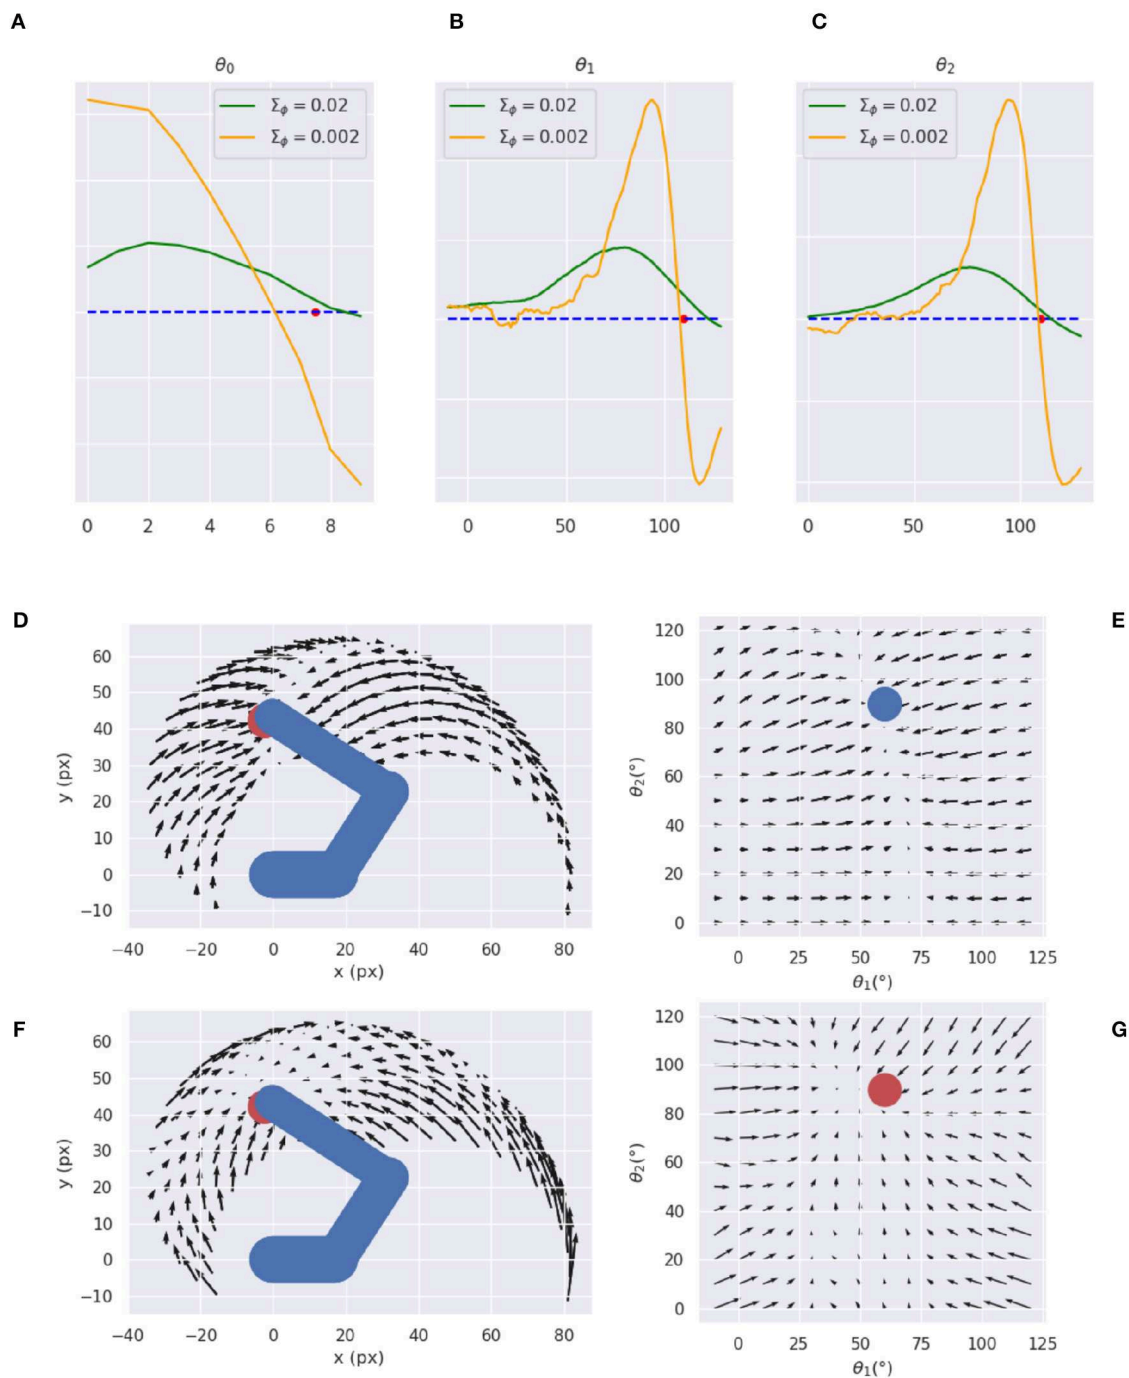
which the agent first rotates the shoulder—resulting in a horizontal positioning—and then starts to rotate the elbow as soon as the latter enters its attraction area. The same gradients can explain the linear motion pattern of an arm tracking dynamically moving targets when the arm is close to the target: in that case, all gradients provide motion force as explained above. On the other hand, the gradients of the target belief (bottom panels) behave somewhat differently: since this belief is unconstrained and can freely move in the environment, update directions more directly approach the target in all angular coordinates (see the polar plot to the right). Yet, linear belief updates in the polar space still translate to curve directions in the Cartesian space.

## 6. Discussion

We presented a normative computational theory based on Active Inference of how the neural circuitry in the PPC and DVS may support visually-guided actions in a dynamically changing environment. Our focus was on the computational basis of encoding dynamic action goals in the PPC through flexible motor intentions and its putative neural basis in the PPC. The theory is based on Predictive Coding (Doya, 2007; Hohwy, 2013), Active Inference (Friston, 2010), and evidence suggesting that the PPC performs visuomotor transformations (Cisek and Kalaska, 2010; Fattori et al., 2017; Galletti and Fattori, 2018) and encodes motor plans (Andersen, 1995; Snyder et al., 1997). Accordingly, the PPC

is proposed to maintain dynamic expectations of both current and desired latent states over the environment and use them to generate proprioceptive predictions that ultimately generate movements through reflex arcs (Adams et al., 2013; Versteeg et al., 2021). In turn, the DVS encodes a generative model that translates latent state expectations into visual predictions. Discrepancies between sensory-level predictions and actual sensations produce prediction errors sent back through the cortical hierarchy to improve the internal representation. The theory unifies research on intention coding (Snyder et al., 1997) and current views that the PPC estimates the body and environmental states (Medendorp and Heed, 2019), providing specific computational hypotheses regarding the involvement in goal-directed behavior. It also extends some perception-bound Predictive Coding interpretations of the PPC dynamics (FitzGerald et al., 2015) and provides a more comprehensive account of movement planning (Erlhagen and Schöner, 2002), tightly integrated into the overall sensorimotor control process.

The core novelty with respect to state-of-the-art implementations in continuous time Active Inference is that we first considered an internal belief over not only bodily states but also every object in the scene, where the latter are encoded in the joint angles space as well, simulating a visuomotor reference frame that the PPC is supposed to encode. Then, we decomposed the belief dynamics into a set of independent intentions each depending on the current belief and predicting the next plausible state. Such formalization has several advantages. First, since



**FIGURE 14** Visual gradient analysis. Marginal gradients for each joint, i.e., neck (A), shoulder (B), and elbow (C), and for two values of the recognition density variance  $\Sigma_\phi$  (green/orange line) computed by backpropagating the error between images with different arm configurations (abscissa: joint angle) and a reference image (whose angle is represented by the red dot on the abscissa). (D–G) Gradients for arm (D, E) and target (F, G) in both Cartesian (D, F) and joint (E, G) space.

attractors are dynamically generated at each time step, the agent can also follow moving targets and interact with a constantly changing environment, in contrast to static reaching tasks where a desired fixed state is specified in the belief dynamics (Baïoumy et al., 2020). Second, expressing the target position in terms of a possible joint configuration—either imposed by higher

levels for realizing specific affordances or freely inferred by the exteroceptive models—results in simple intentions, without the need to directly use sensory information or duplicating lower-level generative models, which leads to implausible scenarios (Lanillos et al., 2020; Sancaktar et al., 2020). It should be however noted that, although an intrinsic-only attractor is faster and more

parsimonious, continuous activation of visual low-level attractors may provide more precise motor control. Indeed, it seems that motor areas are able, at a certain "neural energy" cost, to interact with and generate predictions in multiple sensory domains. The key difference is that in the former a single prediction—which is already biased toward a future state—is compared with the sensory input; in the latter, a prediction of the current state is compared with the desired exteroceptive goal, biasing the belief through the backpropagated gradient. Further studies are thus needed to implement low-level attractors in a biologically plausible way—e.g., intentions could generate through parallel pathways their own future sensory predictions that are compared with the observations in the usual way, with a particular intention that can be viewed as trying to continuously predict the current sensory input—and analyze the differences between the two modalities. Last, maintaining different belief components also allows easy encoding of previously memorized states which can be especially useful when implementing a sequence of actions, since only the intention precisions have to be adjusted. Indeed, it seems that the PPC explicitly encodes and maintains such goals during the whole unfolding of sequential actions (Baldauf et al., 2008). A specific goal is selected among other competitive intentions possibly under the control of the PFC and PMd (Stoianov et al., 2016) and fulfilled by setting it as a predominant belief trajectory as an attractor with a strong gain (see Equations 44, 26 and Figure 4). For example, in a typical reaching task, the goal of reaching a specific visual target corresponds to the future expectation that the agent's arm will be over that target; thus, if the agent maintains a belief over the latter, the corresponding intention links the expected belief over the future body posture with the inferred target, expressed in joint angles, encoding a specific interaction to realize.

We tested the computational feasibility of the theory on a delayed reaching task—a classical experiment in electrophysiology—in which a monkey is required to reach with its hand a visual-spatial target, starting the movement from an HB (Breviglieri et al., 2014). To do this, we simulated an agent consisting of a coarse 3-DoF limb model and noisy visual and proprioceptive sensors (Figure 3A). Simplified proprioceptive sensors provided a noisy reading of the state of the limb in joint angles, while visual input was provided by a fixed camera and consisted of an image of the target and limb. Predictive visual sensory processing simulating the DVS was implemented with a VAE trained to infer body state and target location, both in the joint angles domain (Figure 3C). The limbs were animated at the velocity level with motor control signals computed by the visually-guided Active Inference controller. The computational analysis showed, first and most importantly, that the controller could correctly infer the position of the visual targets (Figure 5,  $t = 70$ ), use it to compute and set motor goals in terms of prior beliefs on the future body state through intention functions (Figure 5,  $t = 105$ ), and perform adequate and smooth reach movements (Figure 5  $t = 105$ –150), with and without visual feedback (Figure 6). The greater accuracy obtained with visual feedback parallels classical results in a similar classical behavioral comparison of reaching (Keele and Posner, 1968).

We then systematically investigated the effects of noise on various functional components (Figure 7), starting with the balance

of the precision between proprioceptive and visual sensory models: a noiseless Active Inference agent (BL condition) resulted in the best performance, with a stable final approach and accuracy only limited by the quality of the visual target estimation. Among the noisy conditions, pure proprioceptive control resulted in the lowest performance, as expected. Motor control driven by both proprioceptive and visual feedback with balanced precision between the two domains resulted in improved reach accuracy and greatly improved arm belief stability (Figure 8). The effect on accuracy was mainly due to the inclusion of visual information in the inference process, but also to slower updates of the motor control signals due to decreased confidence about proprioceptive input. The increased stability of the arm belief did not improve movement stability as increasing confidence about visual input also increased the discrepancy between belief and action updates, the latter only relying on noisy proprioceptive observations. In fact, we showed that if we remove the plausibility constraint that motor control is driven only by proprioceptive predictions and thus let actions minimize prediction errors from all sensory domains, the reach performance greatly increases (Figure 9). Nonetheless, any combination of visual and proprioceptive feedback improved performance relative to a control driven by feedback from a single sensory domain. The instability due to the difference in update directions between belief and action could be balanced by other mechanisms that we have not considered here. For example, we assumed that the same pathway is used for both control and belief inference, but it seems that the motor cortex generates different predictions depending on the brain areas which it interacts with: purely proprioceptive predictions for motor control, whose prediction error is suppressed at the lowest level of the hierarchy, and rich somatosensory predictions for latent state inference, which integrates somatic sensations at different hierarchical levels (Adams et al., 2013). Intention precisions or attractor gains affected performance as well. First, they affected reach time: as expected, the greater the gain, the faster the movement. However, fast movements come at a cost: increased gains generally resulted in less precise movements and decreased stability during the final reach period. Finally, higher visual target precisions decreased perception time and improved perception accuracy but decreased perception stability.

We also investigated the effects of movement onset policies: response delay allows investigating perceptual and motor preparatory processes separately from the motor control and action execution. We found that delayed response decreased movement time with respect to a policy that requires an immediate response (Figure 10), which fits the behavioral pattern (Shenoy et al., 2013). Apparently, this is due to the need to estimate, in the latter condition, the target position "on the fly," and constantly adapt the intention according to the updated target estimate. The advantage of allowing some preparatory time becomes clear in an anecdotal fly-catching task, which results in faster movement and increased chances of success. This comes with the critical contribution of PPC neurons that systematically modulate their activity during the preparatory period (Shenoy et al., 2013), which here provided specific predictions for the computations performed in the PPC. Notably, the immediate-response policy allowed the Active Inference controller to perform actions under dynamic

environmental conditions, such as tracking moving objects. Free energy minimization resulted in rapid target detection also in this case, and maintained subthreshold perception error on moving targets (Figure 11) which allowed precise tracking after an initial reaching period.

Intention-driven Active Inference in continuous time largely compares to classical neural-level hypotheses of motor planning such as the Dynamic Field Theory (Erlhagen and Schöner, 2002), with the advantage of stepping on an established Predictive Coding framework, dynamic approximate probabilistic inference, and end-to-end sensorimotor control. The Dynamic Field Theory estimates the parameters of the desired movement—such as movement direction and target velocity—from sensory and task features encoding environmental descriptors, which closely compares to motor goal coding through flexible intentions in our model. The two theories have in common a dynamic activation of the internal representations in continuous time, governed by a dynamic system, but they differ in the nature of the signals and their coding. Movement descriptors in Dynamic Field Theory are represented by a dynamically activated multidimensional space, each encoded on a population of competing neurons, while Active Inference approximates movement properties with their central moments (belief) and dispersion (precisions). While population coding allows a complete description of a probabilistic distribution, it could be overshoot when used to code single magnitudes (although it is essential when encoding discrete categories). Yet, such representation allows coding multiple competing targets on the same population of neurons, while in our scheme each target should be encoded by a dedicated unit. Notably, the brain encodes scalar variables using a variety of number coding schemes, including monotonic and distributed (Stoianov and Zorzi, 2012). The latter, known as a “mental number line” (Stoianov et al., 2008), could be an interesting hypothesis to explore also in the context of feature coding in continuous Active Inference. Currently, distributed coding is used only in discrete Active Inference and other probabilistic models to investigate computationally high-level cognitive functions such as planning, navigation, and control (Stoianov et al., 2016, 2022; Pezzulo et al., 2018). The two theories also differ in the nature of the input to their dynamic systems. In Active Inference, system input encodes generalized prediction errors, which are integrated into higher-level moments. Instead, input in Dynamic Field Theory directly encodes state values. Coding based on prediction errors has the advantage of minimizing the quantity of transmitted information—hence, energy. Finally, the theories also differ in scope: Active Inference provides a full account of the entire sensorimotor control process, while Dynamic Field Theory describes only movement planning.

## 6.1. Precision balance and conditions for disorders


Based on our computational analysis, it becomes clear that some motor and behavioral disorders could be due to the lack of proper sensory and intention precisions (Adams et al., 2021). Here, we illustrate the normal condition and two types of potentially improper precision balance that could become a causal condition

for neurological disorders. Figure 15A illustrates the condition for normal functioning, which is such that the contribution of a single intention to the belief update (which, as a reminder to the reader, is proportional to the gain of intentions  $\lambda$ ) is sufficiently small with respect to the sensory contribution. In this case, during free energy minimization, the system dynamics smoothly moves the belief toward the strongest goal, along with precise tracking of the true latent state and sensory signal of the limbs, allowing thus to compute correct motor control errors and perform smooth action execution. A critical abnormal condition arises when the intention gain  $\lambda$  is too strong, as illustrated in Figure 15B. In this case, the belief moves too rapidly toward the goal without being able to match the proprioceptive observations, which results in computing incorrect motor control signals. Another abnormal condition is caused by too close precisions  $\gamma_k$  of competitive intentions, which is likely to result in opposing belief updates and thus prevent the fulfillment of any of the competing goals (as in Figure 15C). This situation might manifest in terms of motor onset failure or oscillatory behavior.

## 6.2. Neural-level predictions

One peculiarity of Active Inference based theories of motor control is that proprioceptive predictions are sent through efferents down to the spinal cord and that specific muscle control signals are computed at that level by reflex arcs, so that action attempts to suppress proprioceptive prediction errors (Adams et al., 2013). This prediction critically differs from competing modern theories such as the Optimal Control (Todorov and Jordan, 2002), according to which the efferents convey muscle control signals computed at the cortical level. A general aspect of Active Inference regards the dynamic inferential process, which predicts with increasing precision the internal representation of the sensorium—including estimation of targets and body posture—starting from noisy priors that gradually converge to ideal states. This kind of precision trend should be observed in an experiment with multiple repetitions of the same action and target, with variability of cell activity encoding the target and body that gradually decreases in time within trials. While this prediction is generally shared with Predictive Coding based theories, classical stimulus-response theories would predict invariant variability of cell activity across time. Another general aspect regards coding of prediction errors. In fact, body-environment transitions involving a change of states and tasks result in transient bursts of activity in error-conveying cells until the error is minimized. Prediction errors conveying upstream information are supposed to be encoded by pyramidal cortical cells in superficial layers while downstream predictions are encoded by deep pyramidal neurons (Parr et al., 2022).

In light of the considerations so far, we predict several different types of correlates that should be found in the PPC related to coding environment, task, and bodily states. The former two include correlates of potential spatial targets and selected motor goals, which indeed have been consistently found in the PPC (Andersen, 1995; Snyder et al., 1997; Filippini et al., 2018). The latter includes correlates of intention-biased bodily state estimates, which thus are not precise representations of the true states. To this concern,



a key expected neural correlate of the proposed mechanism in the PPC includes signals encoding intention prediction errors between the current belief and future states corresponding to targets to interact with, both encoded in a visuomotor reference frame. To investigate this, one can manipulate high-level priors, e.g., by inducing an abrupt change of the intention, which should then be observed as a fast decaying change of the corresponding prediction error. A related hypothesis is that in tasks comprising several targets—like the classical monkey experiment analyzed here—each goal generates its own intention and prediction error. In normal conditions only one intention is selected at a time, and this behavior should be observed in the relative dynamics between all the intention prediction errors encoded simultaneously. Finally, the use of generalized beliefs in Active Inference predicts that the PPC encodes not only static states but also a detailed estimate of body dynamics, up to a few temporal orders. Indeed, a body of literature report motion-sensitive, or Vision-for-Action activity in the DVS and PPC (Galletti and Fattori, 2018). The validation of all these correlates will be the subject of further studies with real monkey experiments similar to the one described in Figure 3B.

### 6.3. Limitations and future directions

Our focus here was on intention coding in the PPC, which directly deals with motor plans and motor control. Further elaborations will extend the theory with higher-level aspects of cognitive control, including intention structuring (dealt by PM), phasing (SMA) (Gallego et al., 2022), planning, and goal selection (HC, PFC) (Stoianov et al., 2016, 2018; Pezzulo et al., 2019). Motor control operated here in an inner belief space belonging to the joint angles domain, which is generally suboptimal in the external Cartesian space. Although a kinematic transformation was implicitly performed by the VAE, we assumed that neural activity in the PPC encodes generalized beliefs over targets and body only in a motor-related domain; however, neural data suggest that neurons in the motor cortex encode motor trajectories also in extrinsic

coordinates (Cohen and Andersen, 2002; Adams et al., 2013), and a more realistic model should include representations encoding states in both intrinsic and extrinsic reference frames. A functional correlate of the motor cortex should represent future states—which were defined here implicitly in the intention prediction errors and dynamics functions—and transform desired trajectories from Cartesian coordinates to proprioceptive predictions in the intrinsic state-space. This transformation is different from Optimal Control planning since the optimization of a classical inverse model reduces to a more manageable inference problem.

Since our focus was on the theoretical introduction of intentionality in Active Inference, every analysis was only partially characterized by a simple reaching task. However, fundamental properties of the physical model, including geometry, mass, and friction, strongly influence the resulting motion dynamics—hence the entire inferential process. This implementation does not adopt other important neural and biomechanical specificities such as signal delay and joint friction (Wolpert and Flanagan, 2016), and just partially covers the three main domains of sensorimotor learning through a predictive forward control; for example, it does not fully include reactive, stimuli-driven control such as obstacle avoidance, although we showed that it can successfully perform static and dynamic tasks. However, it could be easily extended to accommodate additional sensory modalities—e.g., tactile sensations—with rich generative models such as the VAE implemented here. Further planned computational analyses will use a richer belief space, a more realistic physical arm model, and additional actuators, and expand the complexity of the intention functions to investigate the capacity of the theory to explain in-depth neural levels, cognitive, and kinematic phenomena related to motor learning, motion perception, motor planning, and so on. Planned future studies with a more articulated agent will also challenge the theory at the behavioral and neural level against other empirical findings regarding movement preparation and motor control, in either delayed or direct response settings. For example, we will test the model for stimulus-stimulus congruency and stimulus-response compatibility effects (Kornblum et al., 1990). As for the former, it is intuitive that a greater sensory



dimension overlap predicts faster target-belief convergence—thus faster intention setting. Less intuitive is that stimulus-response compatibility effects should emerge due to differences in the dynamic transition from one belief state to another in the proprioceptive domain. For example, a belief over the effector state should change more, requiring more time to converge when reaching a contralateral position than an ipsilateral one.

Although we considered an Active Inference model with just a single layer of intentions, the structure represented in Figure 2 could be scaled hierarchically and intermediate goals could be considered between high-level intentions and low-level sensory generative models, e.g., by combining discrete and continuous Active Inference for planning and movement execution (Friston et al., 2017a,b; Parr et al., 2020; Sajid et al., 2021). According to the free energy principle, the agent will then choose goals and subgoals and rely on specific sensory modalities such that free energy is minimized at every hierarchical level based on prediction errors coming from the level below. This formalization will provide an explicit basis for motor planning, including tasks like object manipulation. Indeed, although the current implementation performs well on spatial tasks like reaching in a dynamically changing environment, it cannot implement composite goals which the brain needs to handle. On the other hand, an agent that can encode higher-level goals in a discrete domain and infer policies based on the *expected free energy* will be able to dynamically modify its behavior and react to environmental changes. An extended implementation of this kind—showing the interplay between discrete goals and continuous intentions—will be the subject of future work.

## Data availability statement

Publicly available datasets were analyzed in this study. This data can be found at: <https://github.com/priorelli/PACE>.

## References

- Adams, R. A., Aponte, E., Marshall, L., and Friston, K. J. (2015). Active inference and oculomotor pursuit: the dynamic causal modelling of eye movements. *J. Neurosci. Methods* 242, 1–14. doi: 10.1016/j.jneumeth.2015.01.003
- Adams, R. A., Shipp, S., and Friston, K. J. (2013). Predictions not commands: active inference in the motor system. *Brain Struct. Funct.* 218, 611–643. doi: 10.1007/s00429-012-0475-5
- Adams, R. A., Vincent, P., Benrimoh, D., Friston, K. J., and Parr, T. (2021). Everything is connected: Inference and attractors in delusions. *Schizophrenia Res.* 245, 5–22. doi: 10.1016/j.schres.2021.07.032
- Andersen, R. A. (1995). Encoding of intention and spatial location in the posterior parietal cortex. *Cereb. Cortex* 5, 457–469. doi: 10.1093/cercor/5.5.457
- Baioumy, M., Duckworth, P., Lacerda, B., and Hawes, N. (2020). Active inference for integrated state-estimation, control, and learning. *arXiv*. doi: 10.1109/ICRA48506.2021.9562009
- Baldauf, D., Cui, H., and Andersen, R. A. (2008). The posterior parietal cortex encodes in parallel both goals for double-reach sequences. *J. Neurosci.* 28, 10081–10089. doi: 10.1523/JNEUROSCI.3423-08.2008
- Baltieri, M., and Buckley, C. L. (2019). PID control as a process of active inference with linear generative models. *Entropy* 21, 257. doi: 10.3390/e21030257
- Bastos, A. M., Usrey, W. M., Adams, R. A., Mangun, G. R., Fries, P., and Friston, K. J. (2012). Canonical microcircuits for predictive coding. *Neuron* 76, 695–711. doi: 10.1016/j.neuron.2012.10.038
- Bishop, C. M. (2006). *Pattern Recognition and Machine Learning*. New York, NY: Springer.
- Bogacz, R. (2017). A tutorial on the free-energy framework for modelling perception and learning. *J. Math. Psychol.* 76, 198–211. doi: 10.1016/j.jmp.2015.11.003
- Breveglieri, R., Galletti, C., Dal Bò, G., Hadjidimitrakis, K., and Fattori, P. (2014). Multiple aspects of neural activity during reaching preparation in the medial posterior parietal area V6A. *J. Cogn. Neurosci.* 26, 879–895. doi: 10.1162/jocn\_a\_00510
- Buckley, C. L., Kim, C. S., McGregor, S., and Seth, A. K. (2017). The free energy principle for action and perception: a mathematical review. *J. Math. Psychol.* 81, 55–79. doi: 10.1016/j.jmp.2017.09.004
- Cisek, P., and Kalaska, J. F. (2010). Neural mechanisms for interacting with a world full of action choices. *Annu. Rev. Neurosci.* 33, 269–298. doi: 10.1146/annurev.neuro.051508.135409
- Cohen, Y. E., and Andersen, R. A. (2002). A common reference frame for movement plans in the posterior parietal cortex. *Nat. Rev. Neurosci.* 3, 553–562. doi: 10.1038/nrn873
- Corbetta, M., and Shulman, G. L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nat. Rev. Neurosci.* 3, 201–215. doi: 10.1038/nrn755

## Author contributions

MP developed the computational method, wrote the code, run the simulations, analyzed the results, and wrote the draft. IS developed theoretical and methodological ideas and wrote the draft. All authors contributed to the article and approved the submitted version.

## Funding

This research was received funding from the European Union's Horizon 2020 Framework Programme for Research and Innovation under H2020-EIC-FETPROACT-2019 Grant Agreement 951910 (MAIA) to IS, Grant Agreement No 945539 (Human Brain Project SGA3), the European Research Council under Grant Agreement No. 820213 (ThinkAhead), and from the Italian Ministry for Research MIUR under Grant Agreement PRIN 2017KZNZLN (PACE) to IS.

## Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

## Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

- Desmurget, M., Epstein, C. M., Turner, R. S., Prablanc, C., Alexander, G. E., and Grafton, S. T. (1999). PPC and visually directing reaching to targets. *Nature* 2, 563–567. doi: 10.1038/9219
- Doya, K. (2007). *Bayesian Brain: Probabilistic Approaches to Neural Coding*. Cambridge, MA: The MIT Press.
- Erlhagen, W., and Schöner, G. (2002). Dynamic field theory of movement preparation. *Psychol. Rev.* 109, 545–572. doi: 10.1037/0033-295X.109.3.545
- Fattori, P., Breveglieri, R., Bosco, A., Gamberini, M., and Galletti, C. (2017). Vision for prehension in the medial parietal cortex. *Cereb. Cortex* 27, 1149–1163. doi: 10.1093/cercor/bhv302
- Filippini, M., Breveglieri, R., Ali Akhras, M., Bosco, A., Chinellato, E., and Fattori, P. (2017). Decoding information for grasping from the macaque dorsomedial visual stream. *J. Neurosci.* 37, 4311–4322. doi: 10.1523/JNEUROSCI.3077-16.2017
- Filippini, M., Breveglieri, R., Hadjidimitrakis, K., Bosco, A., and Fattori, P. (2018). Prediction of reach goals in depth and direction from the parietal cortex. *Cell Rep.* 23, 725–732. doi: 10.1016/j.celrep.2018.03.090
- FitzGerald, T. H., Moran, R. J., Friston, K. J., and Dolan, R. J. (2015). Precision and neuronal dynamics in the human posterior parietal cortex during evidence accumulation. *Neuroimage* 107, 219–228. doi: 10.1016/j.neuroimage.2014.12.015
- Fogassi, L., Ferrari, P. F., Gesierich, B., Rozzi, S., Chersi, F., and Rizzolatti, G. (2005). Parietal lobe: from action organization to intention understanding. *Science* 308, 662–667. doi: 10.1126/science.1106138
- Franklin, D. W., and Wolpert, D. M. (2011). Computational mechanisms of sensorimotor control. *Neuron* 72, 425–442. doi: 10.1016/j.neuron.2011.10.006
- Friston, K. (2008). Hierarchical models in the brain. *PLoS Comput. Biol.* 4, e1000211. doi: 10.1371/journal.pcbi.1000211
- Friston, K. (2010). The free-energy principle: a unified brain theory? *Nat. Rev. Neurosci.* 11, 127–138. doi: 10.1038/nrn2787
- Friston, K. (2011). What is optimal about motor control? *Neuron* 72, 488–498. doi: 10.1016/j.neuron.2011.10.018
- Friston, K. (2012). The history of the future of the Bayesian brain. *Neuroimage* 62, 1230–1233. doi: 10.1016/j.neuroimage.2011.10.004
- Friston, K., and Kiebel, S. (2009). Predictive coding under the free-energy principle. *Philos. Trans. R. Soc. B Biol. Sci.* 364, 1211–1221. doi: 10.1098/rstb.2008.0300
- Friston, K., Mattout, J., Trujillo-Barreto, N., Ashburner, J., and Penny, W. (2007). Variational free energy and the Laplace approximation. *Neuroimage* 34, 220–234. doi: 10.1016/j.neuroimage.2006.08.035
- Friston, K. J. (2002). Functional integration and inference in the brain. *Progr. Neurobiol.* 68, 113–143. doi: 10.1016/S0301-0082(02)00076-X
- Friston, K. J. (2005). A theory of cortical responses. *Philos. Trans. R. Soc. Lond B Biol. Sci.* 360, 815–836. doi: 10.1098/rstb.2005.1622
- Friston, K. J., Daunizeau, J., and Kiebel, S. J. (2009). Reinforcement learning or active inference? *PLoS ONE* 4, e6421. doi: 10.1371/journal.pone.0006421
- Friston, K. J., Daunizeau, J., Kilner, J., and Kiebel, S. J. (2010). Action and behavior: a free-energy formulation. *Biol. Cybern.* 102, 227–260. doi: 10.1007/s00422-010-0364-z
- Friston, K. J., Mattout, J., and Kilner, J. (2011). Action understanding and active inference. *Biol. Cybern.* 104, 137–160. doi: 10.1007/s00422-011-0424-z
- Friston, K. J., Parr, T., and de Vries, B. (2017a). The graphical brain: belief propagation and active inference. *Netw. Neurosci.* 1, 381–414. doi: 10.1162/NETN\_a\_00018
- Friston, K. J., Rosch, R., Parr, T., Price, C., and Bowman, H. (2017b). Deep temporal models and active inference. *Neurosci. Biobehav. Rev.* 77, 388–402. doi: 10.1016/j.neubiorev.2017.04.009
- Friston, K. J., Samothrakis, S., and Montague, R. (2012). Active inference and agency: optimal control without cost functions. *Biol. Cybern.* 106, 523–541. doi: 10.1007/s00422-012-0512-8
- Friston, K. J., Trujillo-Barreto, N., and Daunizeau, J. (2008). DEM: A variational treatment of dynamic systems. *Neuroimage* 41, 849–885. doi: 10.1016/j.neuroimage.2008.02.054
- Gallejo, J. A., Makin, T. R., and McDougle, S. D. (2022). Going beyond primary motor cortex to improve brain-computer interfaces. *Trends Neurosci.* 45, 176–183. doi: 10.1016/j.tins.2021.12.006
- Galletti, C., and Fattori, P. (2018). The dorsal visual stream revisited: Stable circuits or dynamic pathways? *Cortex* 98, 203–217. doi: 10.1016/j.cortex.2017.01.009
- Galletti, C., Gamberini, M., and Fattori, P. (2022). The posterior parietal area V6A: an attentionally-modulated visuomotor region involved in the control of reach-to-grasp action. *Neurosci. Biobehav. Rev.* 141, 104823. doi: 10.1016/j.neubiorev.2022.104823
- Gamberini, M., Passarelli, L., Filippini, M., Fattori, P., and Galletti, C. (2021). Vision for action: thalamic and cortical inputs to the macaque superior parietal lobule. *Brain Struct. Funct.* 226, 2951–2966. doi: 10.1007/s00429-021-02377-7
- Genovesio, A., Tsujimoto, S., and Wise, S. P. (2012). Encoding goals but not abstract magnitude in the primate prefrontal cortex. *Neuron* 74, 656–662. doi: 10.1016/j.neuron.2012.02.023
- Goodfellow, I. J., Bengio, Y., and Courville, A. (2016). *Deep Learning*. Cambridge, MA: MIT Press.
- Haar, S., and Donchin, O. (2020). A revised computational neuroanatomy for motor control. *J. Cogn. Neurosci.* 32, 1823–1836. doi: 10.1162/jocn\_a\_01602
- Hohwy, J. (2013). *The Predictive Mind*. Oxford: Oxford University Press UK. doi: 10.1093/acprof:oso/9780199682737.001.0001
- Kaplan, R., and Friston, K. J. (2018). Planning and navigation as active inference. *Biol. Cybern.* 112, 323–343. doi: 10.1007/s00422-018-0753-2
- Keele, S. W., and Posner, M. I. (1968). Processing of visual feedback in rapid movements. *J. Exp. Psychol.* 77, 155–158. doi: 10.1037/h0025754
- Kikuchi, Y., and Hamada, Y. (2009). Geometric characters of the radius and tibia in Macaca mulatta and Macaca fascicularis. *Primates* 50, 169–183. doi: 10.1007/s10329-008-0120-3
- Kingma, D. P., and Welling, M. (2014). “Auto-encoding variational bayes,” in *2nd International Conference on Learning Representations, ICLR 2014-Conference Track Proceedings* (Banff), 1–14. doi: 10.48550/arXiv.1312.6114
- Kornblum, S., Hasbroucq, T., and Osman, A. (1990). Dimensional overlap: cognitive basis for stimulus-response compatibility—a model and taxonomy. *Psychol. Rev.* 97, 253–270. doi: 10.1037/0033-295X.97.2.253
- Lanillos, P., and Cheng, G. (2018). “Adaptive robot body learning and estimation through predictive coding,” in *IEEE International Conference on Intelligent Robots and Systems* (Madrid: IEEE), 4083–4090.
- Lanillos, P., Pages, J., and Cheng, G. (2020). “Robot self/other distinction: active inference meets neural networks learning in a mirror,” in *ECAI 2020* (Santiago de Compostela). doi: 10.48550/arXiv.2004.05473
- Lau, H. C., Rogers, R. D., Haggard, P., and Passingham, R. E. (2004). Attention to Intention. *Science* 303, 1208–1210. doi: 10.1126/science.1090973
- Levine, S. (2018). Reinforcement learning and control as probabilistic inference: tutorial and review. *ArXiv [Preprint]*. doi: 10.48550/arXiv.1805.00909
- Limanowski, J., and Friston, K. (2020). Active inference under visuo-proprioceptive conflict: simulation and empirical results. *Sci. Rep.* 10, 1–14. doi: 10.1038/s41598-020-61097-w
- Ma, W. J., Beck, J. M., Latham, P. E., and Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nat. Neurosci.* 9, 1432–1438. doi: 10.1038/nn1790
- Medendorp, W. P., and Heed, T. (2019). State estimation in posterior parietal cortex: distinct poles of environmental and bodily states. *Progr. Neurobiol.* 183, 101691. doi: 10.1016/j.pneurobio.2019.101691
- Millidge, B., Tschantz, A., Seth, A. K., and Buckley, C. L. (2020). On the relationship between active inference and control as inference. *Commun. Comput. Inf. Sci.* 1326, 3–11. doi: 10.1007/978-3-030-64919-7\_1
- Oliver, G., Lanillos, P., and Cheng, G. (2019). Active inference body perception and action for humanoid robots. *ArXiv [Preprint]*. doi: 10.48550/arXiv.1906.03022
- Parr, T., and Friston, K. J. (2018). The anatomy of inference: Generative models and brain structure. *Front. Comput. Neurosci.* 12, 90. doi: 10.3389/fncom.2018.00090
- Parr, T., Pezzulo, G., and Friston, K. J. (2022). *Active Inference: The Free Energy Principle in Mind, Brain, and Behavior*. Cambridge, MA: The MIT Press. doi: 10.7551/mitpress/12441.001.0001
- Parr, T., Rikhye, R. V., Halassa, M. M., and Friston, K. J. (2020). Prefrontal computation as active inference. *Cereb. Cortex* 30, 682–695. doi: 10.1093/cercor/bhz118
- Pezzulo, G., and Cisek, P. (2016). Navigating the affordance landscape: feedback control as a process model of behavior and cognition. *Trends Cogn. Sci.* 20, 414–424. doi: 10.1016/j.tics.2016.03.013
- Pezzulo, G., Donnarumma, F., Dindo, H., D’Ausilio, A., Konvalinka, I., and Castelfranchi, C. (2019). The body talks: sensorimotor communication and its brain and kinematic signatures. *Phys. Life Rev.* 28, 1–21. doi: 10.1016/j.plrev.2018.06.014
- Pezzulo, G., Donnarumma, F., Iodice, P., Maisto, D., and Stoianov, I. (2017). Model-based approaches to active perception and control. *Entropy* 19, 266. doi: 10.3390/e19060266
- Pezzulo, G., Rigoli, F., and Friston, K. J. (2018). Hierarchical active inference: a theory of motivated control. *Trends Cogn. Sci.* 22, 294–306. doi: 10.1016/j.tics.2018.01.009
- Pio-Lopez, L., Nizard, A., Friston, K., and Pezzulo, G. (2016). Active inference and robot control: a case study. *J. R. Soc. Interface* 13, 122. doi: 10.1098/rsif.2016.0616
- Rao, R. P., and Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nat. Neurosci.* 2, 79–87. doi: 10.1038/4580
- Rood, T., van Gerven, M., and Lanillos, P. (2020). “A deep active inference model of the rubber-hand illusion,” in *Active Inference. IWAI 2020. Communications in Computer and Information Science*, Vol. 1326, eds T. Verbelen, P. Lanillos, C. L. Buckley and C. De Boom (Cham: Springer).

- Sajid, N., Ball, P. J., Parr, T., and Friston, K. J. (2021). Active inference: demystified and compared. *Neural Comput.* 33, 674–712. doi: 10.1162/neco\_a\_01357
- Sancaktar, C., van Gerven, M. A. J., and Lanillos, P. (2020). "End-to-end pixel-based deep active inference for body perception and action," in *2020 Joint IEEE 10th International Conference on Development and Learning and Epigenetic Robotics (ICDL-EpiRob)* (Valparaiso: IEEE), 1–8.
- Saunders, J. A., and Knill, D. C. (2003). Humans use continuous visual feedback from the hand to control fast reaching movements. *Exp. Brain Res.* 152, 341–352. doi: 10.1007/s00221-003-1525-2
- Shadmehr, R., and Krakauer, J. W. (2008). A computational neuroanatomy for motor control. *Exp. Brain Res.* 185, 359–381. doi: 10.1007/s00221-008-1280-5
- Shenoy, K. V., Sahani, M., and Churchland, M. M. (2013). Cortical control of arm movements: a dynamical systems perspective. *Annu. Rev. Neurosci.* 36, 337–359. doi: 10.1146/annurev-neuro-062111-150509
- Snyder, L. H., Batista, A. P., and Andersen, R. A. (1997). Coding of intention in the posterior parietal cortex. *Nature* 386, 167–170. doi: 10.1038/386167a0
- Snyder, L. H., Batista, A. P., and Andersen, R. A. (2000). Intention-related activity in the posterior parietal cortex: a review. *Vision Res.* 40, 1433–1441. doi: 10.1016/S0042-6989(00)00052-3
- Srinivasan, S. S., Gutierrez-Arango, S., Teng, A. C. E., Israel, E., Song, H., Bailey, Z. K., et al. (2021). Neural interfacing architecture enables enhanced motor control and residual limb functionality postamputation. *Proc. Natl. Acad. Sci. U.S.A.* 118, e2019555118. doi: 10.1073/pnas.2019555118
- Stoianov, I., Genovesio, A., and Pezzulo, G. (2016). Prefrontal goal codes emerge as latent states in probabilistic value learning. *J. Cogn. Neurosci.* 28, 140–157. doi: 10.1162/jocn\_a\_00886
- Stoianov, I., Kramer, P., Umiltà, C., and Zorzi, M. (2008). Visuospatial priming of the mental number line. *Cognition.* 106, 770–779. doi: 10.1016/j.cognition.2007.04.013
- Stoianov, I., Maisto, D., and Pezzulo, G. (2022). The hippocampal formation as a hierarchical generative model supporting generative replay and continual learning. *Progr. Neurobiol.* 217, 1–20. doi: 10.1016/j.pneurobio.2022.102329
- Stoianov, I., Pennartz, C., Lansink, C., and Pezzulo, G. (2018). Model-based spatial navigation in the hippocampus-ventral striatum circuit: a computational analysis. *PLoS Comput. Biol.* 14, 1–28. doi: 10.1371/journal.pcbi.1006316
- Stoianov, I., and Zorzi, M. (2012). Emergence of a 'visual number sense' in hierarchical generative models. *Nat. Neurosci.* 15, 194–196. doi: 10.1038/nn.2996
- Todorov, E. (2004). Optimality principles in sensorimotor control. *Nat. Neurosci.* 7, 907–915. doi: 10.1038/nn1309
- Todorov, E., and Jordan, M. I. (2002). Optimal feedback control as a theory of motor coordination. *Nat. Neurosci.* 5, 1226–1235. doi: 10.1038/nn963
- Toussaint, M., and Storkey, A. (2006). Probabilistic inference for solving discrete and continuous state Markov Decision Processes. *ACM Int. Conf. Proceed. Ser.* 148, 945–952. doi: 10.1145/1143844.1143963
- Tuthill, J. C., and Azim, E. (2018). Proprioception. *Curr. Biol.* 28, R194–R203. doi: 10.1016/j.cub.2018.01.064
- Velliste, M., Perel, S., Spalding, M. C., Whitford, A. S., and Schwartz, A. B. (2008). Cortical control of a prosthetic arm for self-feeding. *Nature* 453, 1098–1101. doi: 10.1038/nature06996
- Versteeg, C., Rosenow, J. M., Bensmaia, S. J., and Miller, L. E. (2021). Encoding of limb state by single neurons in the cuneate nucleus of awake monkeys. *J. Neurophysiol.* 126, 693–706. doi: 10.1152/jn.00568.2020
- Wolpert, D. M., and Flanagan, J. R. (2016). Computations underlying sensorimotor learning. *Curr. Opin. Neurobiol.* 37, 7–11. doi: 10.1016/j.conb.2015.12.003