

## Predictive Processing in Cognitive Robotics: A Review

**Alejandra Ciria**

*ciriacontacto@gmail.com*

*Facultad de Psicología, Universidad Nacional Autónoma de México,  
Mexico City, CP 04510, Mexico*

**Guido Schillaci**

*guido.schillaci@santannapisa.it*

*BioRobotics Institute, Scuola Superiore Sant'Anna, 34 56025 Pontedera, Italy*

**Giovanni Pezzulo**

*giovanni.pezzulo@istc.cnr.it*

*Institute of Cognitive Sciences and Technologies, National Research Council,  
44 00185 Rome, Italy*

**Verena V. Hafner**

*hafner@informatik.hu-berlin.de*

*Adaptive Systems Group, Department of Computer Science, Humboldt-Universität  
zu Berlin, D-12489, Germany*

**Bruno Lara**

*bruno.lara@uaem.mx*

*Laboratorio de Robótica Cognitiva, Centro de Investigación en Ciencias, Universidad  
Autónoma del Estado de Morelos, Cuernavaca CP 62209, Mexico*

Predictive processing has become an influential framework in cognitive sciences. This framework turns the traditional view of perception upside down, claiming that the main flow of information processing is realized in a top-down, hierarchical manner. Furthermore, it aims at unifying perception, cognition, and action as a single inferential process. However, in the related literature, the predictive processing framework and its associated schemes, such as predictive coding, active inference, perceptual inference, and free-energy principle, tend to be used interchangeably. In the field of cognitive robotics, there is no clear-cut distinction on which schemes have been implemented and under which assumptions. In this letter, working definitions are set with the main aim of analyzing the state of the art in cognitive robotics research working under the predictive processing framework as well as some related nonrobotic models. The analysis suggests that, first, research in both cognitive robotics implementations and nonrobotic models needs to be extended to the study of how multiple exteroceptive modalities can be integrated into prediction

error minimization schemes. Second, a relevant distinction found here is that cognitive robotics implementations tend to emphasize the learning of a generative model, while in nonrobotics models, it is almost absent. Third, despite the relevance for active inference, few cognitive robotics implementations examine the issues around control and whether it should result from the substitution of inverse models with proprioceptive predictions. Finally, limited attention has been placed on precision weighting and the tracking of prediction error dynamics. These mechanisms should help to explore more complex behaviors and tasks in cognitive robotics research under the predictive processing framework.

## 1 Introduction

---

Predictive processing has become an influential framework in the cognitive sciences. A defining characteristic of predictive processing is that it “depicts perception, cognition, and action as the closely woven product of a single kind of inferential process” (Clark, 2018, 522). This idea has had a profound effect on models and theories in different research communities, from neuroscience to psychology, computational modeling, and cognitive robotics. In the literature, terms such as *predictive processing*, *hierarchical predictive processing*, *active inference*, *predictive coding*, and *free energy principle* are often used interchangeably. Scholars refer to them as either theories or frameworks, occasionally interweaving their core ideas.

In cognitive robotics, a number of architectures and models have claimed to follow the postulates of these frameworks. Research in embodied cognitive robotics focuses on understanding and modeling perception, cognition, and action in artificial agents. It is through bodily interactions with their environment that agents are expected to learn and then be capable of performing cognitive tasks autonomously (Lara et al., 2018; Schillaci, Hafner, & Lara, 2016). The aim of this letter is to set working definitions and delimit the main ideas for each of these frameworks, so as to be able to analyze the literature of cognitive robotics and the different implementations in the literature. This should help to highlight what has been done and what is missing and, above all, what the real impact of these frameworks in the area of robotics and artificial intelligence is. Finally, this letter sets the issues and challenges that these new frameworks bring to the table.

The structure of this letter is as follows. Section 2 sets the relevant working definitions. In section 3, different models and architectures are analyzed in the light of the above mentioned frameworks. Section 4 concludes.

## 2 Working Definitions

---

For the purpose of this article, predictive processing is considered to be the most general set of postulates. It proposes to turn the traditional picture of

perception upside down (Clark, 2015). The standard picture of perceptual processing is dominated by the bottom-up flow of information transduced from sensory receptors. In this picture of perception, as information flows upward, a progressively richer picture of the world is constructed from a low-level feature layer processing perceptual input to a high-level semantics layer interpreting information (Marr, 1982). Altogether, predictive processing claims to unify perception, cognition, and action under the same explanatory scope (Clark, 2013; Hohwy, 2013).

The predictive processing view of perception states that agents are constantly and actively predicting sensory stimulation and that only deviations from the predicted sensory input (prediction errors) are processed bottom-up. Prediction error is newsworthy sensory information that provides corrective feedback on top-down predictions and promotes learning. Therefore, in this view of perception, the core flow of information is top-down, and the bottom-up flow of sensory information is replaced by the upward flow of prediction error. The core function of the brain is minimizing prediction error. This process has become known as prediction error minimization (PEM). In a general sense, PEM has been a scheme used in many machine learning algorithms where the error between the desired output and the output generated by the network is used for learning (see, for instance, backpropagation algorithms for training neural networks). Different strategies of PEM have been used in models for perception and action control in artificial agents (see Schillaci, Hafner et al., 2016, for a review).

Going further, predictive processing suggests that the brain is an active organ that constantly generates explanations about sensory inputs and then tests these hypotheses against incoming sensory information (Feldman & Friston, 2010) in a way that is coherent with Helmholtz's view of perception as an unconscious form of inference.

Recurrent neuronal interactions with descending predictions and ascending prediction errors following the predictive processing postulates are illustrated in a simplified segment of the cortical hierarchy in Figure 1A. Neuronal activity of deep pyramidal cells (represented in black) at higher layers of the cortex encode prior beliefs about the expected states of the superficial pyramidal cells (represented in red) at lower layers. At each cortical level, prior beliefs encode the more likely neuronal activity at lower levels. Superficial pyramidal cells compare descending predictions with the ascending sensory evidence, resulting in what is known as prediction error. The prediction error at superficial pyramidal cells is sent to deep pyramidal cells for belief updating (posterior belief). In Figure 1B, descending modulation determines the relative influence of prediction errors at lower levels of the hierarchy on deep pyramidal cells encoding predictions. Precision beliefs are encoded by a descending neuromodulatory gating or gain control (green) of superficial pyramidal cells. In Bayesian inference, beliefs about precision have a great effect on how posterior beliefs are updated. Precision beliefs are considered an attentional mechanism that weights

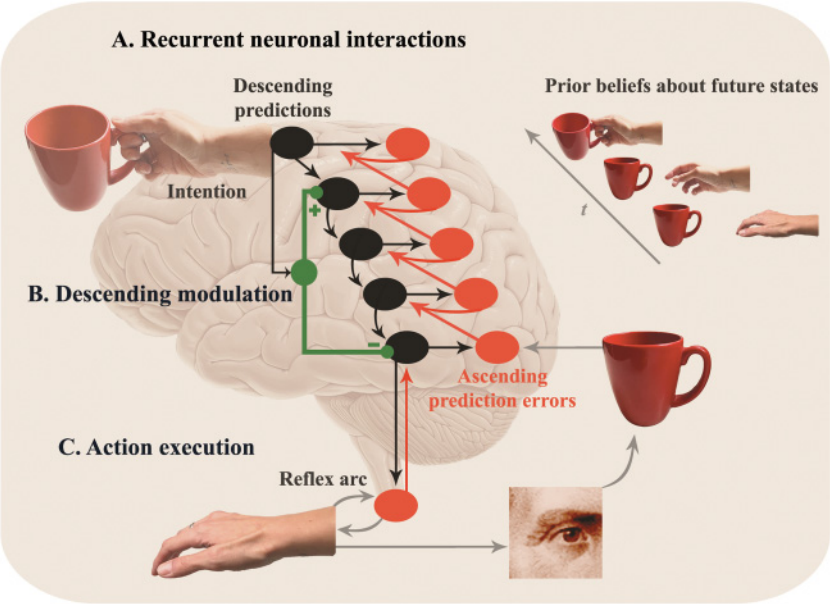


Figure 1: Schematic representation of hierarchical neuronal message under the predictive processing postulates.

predictions and sensory evidence depending on how certain or useful these are for a given task and context. Figure 1C shows a particular example of active inference for prediction error minimization. Perceptual inferences about grasping a cup generate visual, cutaneous, and proprioceptive prediction errors that are then minimized by movement. Descending proprioceptive predictions should be fulfilled by being highly weighted to incite movement. Then, proprioceptive prediction errors are generated at the level of the spinal cord and minimized at the level of peripheral reflexes. At the same time, when the movement trajectory to grasp the cup is performed, visual and cutaneous prediction errors are minimized at all levels of the cortical hierarchy.

Humans and other biological agents deal with a world full of sensory uncertainty. In humans, there is psychophysical evidence that shows how Bayesian models can account for perceptual and motor biases by encoding uncertainty in the internal representations of the brain (Knill & Pouget, 2004).

Several Bayesian approaches center on the idea that perceptual and cognitive processes are supported by internal probabilistic generative models (Clark, 2013, 2015; Friston, 2010a; Hohwy, 2013; Rao & Ballard, 1999). A generative model is a probabilistic model (joint density), mapping hidden

causes in the environment with sensory consequences from which samples are generated (Friston, 2010a). It is usually specified in terms of the likelihood probability distribution of observing some sensory information given its causes and a prior probability distribution of the beliefs about the hidden causes of sensory information (before sampling new observations) (Badcock, Davey, Whittle, Allen, & Friston, 2017). A posterior density is a posterior belief generated by combining the prior and the likelihood weighted according to their precision, defined as the inverse variance (Adams, Stephan, Brown, Frith, & Friston, 2013). A posterior density can be calculated using Bayes' theorem:

$$p(s|O) = \frac{p(O|s)p(s)}{p(O)}, \quad (2.1)$$

where  $p(s|O)$ , also known as the posterior belief, is the probability of hypothesis  $s$  with given evidence or observation  $O$ . Prior beliefs are updated (thus becoming posterior beliefs) when sensory evidence (likelihood) is available.  $p(O|s)$  is the likelihood of relating the sensory observation to the hidden causes, that is, the probability of the specific evidence  $O$ .  $P(s)$  is the prior distribution of any hypothesis  $s$  or prior belief, and it can be seen as the prediction of states.  $P(O)$  is the probability of encountering this evidence or observation.

This calculation is often practically intractable, and variational Bayes is then used for approximately calculating the posterior. This method introduces an optimization problem that requires an auxiliary probability density termed the recognition density (Buckley, Kim, McGregor, & Seth, 2017).

Prediction error is the difference between the mean of the prior belief and the mean of the likelihood in their respective probability distributions. Information gain is measured as the Kullback-Leibler (KL) divergence between the prior belief and the posterior belief. The prior and likelihood distributions have an expected precision, which is encoded as the inverse of their respective variance. This precision will bias the posterior belief update. In particular, the posterior belief is updated biased toward the prior belief given its higher expected precision as compared to the low expected precision on sensory evidence (see Figure 2A). On the contrary, when the expected precision on prior belief is low and the expected precision on sensory evidence is high, the prediction is more uncertain or unreliable, having less of an impact on how the posterior belief is updated than the sensory evidence (see Figure 2B). In both examples in Figure 2, although the magnitude of the prediction error is equivalent, the information gain is greater in panel B due to the greater divergence between the prior and the posterior beliefs.

In Bayesian inference, there are beliefs about beliefs (empirical priors) in terms of having expectations about the beliefs' precision or uncertainty (Adams, Stephan et al., 2013). Here, attention is seen as a selective sampling




Figure 2: Relevance of the precision of probability distributions in Bayesian inference.

of sensory information in such a way that predictions about the confidence of the signals are made to enhance or attenuate prediction errors from different sensory modalities. In order to attain this sampling, this framework proposes a mechanism known as *precision weighting*. The information coming from different modalities is weighted according to the expected confidence given a certain task in a certain context (Parr & Friston, 2017; Friston, Adams, Perrinet, & Breakspear, 2012; Donnarumma, Costantini, Ambrosini, Friston, & Pezzulo, 2017).

Importantly, precision weights are not only assigned according to their reliability, but also by their context-varying usefulness, and are thus considered to be a mechanism for behavior control (Clark, 2020). In the brain, precision weighting might be mediated by a neuromodulatory gain control that can be conceived as a Bayes-optimal encoding of precision at a synaptic level of neuronal populations encoding prediction errors (Friston, Stephan, Montague, & Dolan, 2014). Prediction errors with high precision have a great impact on belief updating, and priors with high precision are robust in the face of noisy or irrelevant prediction errors.

Bayesian beliefs are treated as inferences about the posterior probability distribution (recognition density) via a process of belief updating (Ramstead, Kirchhoff, & Friston, 2020). The recognition density is an approximate probability distribution of the causes of sensory information, which encodes posterior beliefs as a product of inverting the generative

model (Friston, 2010a). According to the Bayesian brain hypothesis, prior beliefs are encoded as neuronal representations, and in light of the new evidence, beliefs are updated (posterior density) to produce a posterior belief following Bayes' rule (Friston et al., 2014). This means that the brain encodes Bayesian recognition densities within its neural dynamics, which can be conceived as inferences of the hidden causes to find the best guess of the environment (Demekas, Parr, & Friston, 2020).

According to Friston, Kilner, and Harrison (2006), predictive processing must be situated within the context of the free-energy principle (Williams, 2018), given that prediction error minimization, under certain assumptions, corresponds to minimizing free energy (Friston, 2010b). Predictive processing can be seen as a name for a family of related theories, where the free energy principle (FEP) provides a mathematical framework to implement the above ideas. The principle is a biological and a neuroscientific framework in which prediction error minimization is conceived as a fundamental process of self-organizing systems to maintain their sensory states within their physiological bounds in the face of constant environmental changes (Adams, Shipp, & Friston, 2013; Friston, 2009, 2010b).

Essentially, the free-energy principle is a mathematical formulation of how biological agents or systems (like brains) resist a natural tendency to disorder by limiting the repertoire of their physiological and sensory states that define their phenotypes (Friston, 2010b). In other words, to maintain their structural integrity, the sensory states of any biological system must have low entropy. Entropy is the negative log-probability of an outcome or the average surprise of sensory signals under the generative model of the causes of the signals (Friston, Mattout, & Kilner, 2011).

Therefore, biological systems are obliged to minimize their sensory surprise (and implicitly entropy) in order to increase the probability of remaining within their physiological bounds over long timescales (Friston, 2009).

The main aim of minimizing free energy is to guarantee that biological systems spend most of their time in their valuable states, those that they expect to frequent. Prior expectations prescribe a primary repertoire of valuable states with innate value, inherited through genetic and epigenetic mechanisms (Friston, 2010b).

Agents are constantly trying to maximize the evidence for the generative model by minimizing surprise. The FEP claims that because biological systems cannot minimize surprise directly, they need to minimize an upper bound called free energy (Buckley et al., 2017). Free energy can be expressed as the Kullback-Leibler divergence between two probability distributions, subtracted by the natural log of the probability of possible states. As Sajid, Parr, Hope, Price, and Friston (2020) stated, free energy can always be written in terms of complexity and accuracy:

$$\begin{aligned} F &= D_{KL}(Q(s)||P(s|o)) - \ln P(o), \\ &= D_{KL}(Q(s)||P(s)) - E_Q[\ln P(o|s)], \end{aligned} \quad (2.2)$$



where  $Q(s)$  is the recognition density or approximate posterior distribution and encodes the prior beliefs an agent possesses about the unknown variables. The conditional density  $P(s|o)$  is the probability of some (hidden) state ( $s$ ) given a certain observation ( $o$ ), and is referred to as the generative model. The first line of equation 2.2 can be read as evidence bound minus log evidence or divergence minus surprise. Rewritten as in the second line, it is read as complexity, which is the difference between the posterior beliefs and prior beliefs before new evidence is available and accuracy, the expected log likelihood of the sensory outcomes given some posterior about the causes of the data (Sajid et al., 2020).

The recognition density (coded by the internal states) and the generative model are necessary to evaluate free energy (Friston, 2010a). Variational free energy (VFE) provides an upper bound on surprise, and it is formally equivalent to weighted prediction error (Buckley et al., 2017). VFE is a statistical measure of the surprise under a generative model. Negative VFE provides a lower bound on model evidence. Minimizing VFE with respect to the recognition density will also minimize the KL divergence between the recognition density and the true posterior. Therefore, minimizing VFE makes the recognition density, the probabilistic representation of the causes of sensory inputs, an approximate of the true posterior (Friston, 2010a). Optimizing the recognition density makes it a posterior density on the causes of sensory information.

Biological agents can minimize free energy by means of two strategies: changing the recognition density or actively changing their internal states. Changing the recognition density minimizes free energy and thus reduces the perceptual divergence. This is a relevant component of the free energy formulation when expressed as complexity minus accuracy.

Minimizing perceptual divergence increases the complexity of the model, defined as the difference between the prior density and the posterior beliefs encoded by the recognition density (Friston, 2010a). This first strategy is known as perceptual inference, this is, when agents change their predictions to match incoming sensory information. Given that sensory information can be noisy and ambiguous, perceptual inferences are necessary to make the input coherent and meaningful.

The second strategy is the standard approach to action in predictive processing, known as active inference (Adams, Shipp et al., 2013; Brown, Adams, Parees, Edwards, & Friston, 2013), which consists of an agent changing sensory inputs through actions that conform to predictions. This is the same as minimizing the expected free energy (Kruglanski, Jasko, & Friston, 2020). When acting on the world, free energy is minimized by sampling sensory information that is consistent with prior beliefs. An action can be defined as a set of real states that change hidden states in the world, which are closely related to control states inferred by the generative model to explain the consequences of action (Friston, Samothrakis, & Montague, 2012). Therefore, actions directly affect the accuracy of the generative



model, defined as the surprise about sensory information expected under the recognition density (Friston, 2010a). For survival, valuable actions are those that are expected to provide agents with the capability to avoid states of surprise.

Every action serves to maximize the evidence of the generative model in such a way that policies are selected to minimize complexity. The expected action consequences include the expected inaccuracy or ambiguity, and the expected complexity or risk, which are combined into the expected free energy (Kruglanski et al., 2020). Thus, expected free energy is the value of a policy, describing its pragmatic (instrumental) and epistemic value. In other words, actions are valuable if they maximize the utility by exploitation (fulfilling preferences), and if they minimize uncertainty by exploration on model parameters (information gathering, as in intrinsic motivation strategies; Seth & Tsakiris, 2018). Maximizing epistemic value is associated with selecting actions that increase model complexity by changing beliefs, whereas maximizing pragmatic value is associated with actions that change internal states that align with beliefs (Tschantz, Seth, & Buckley, 2020). Consequently, the minimization of expected free energy occurs when pragmatic and epistemic values are maximized.

Priors are constantly optimized because they are linked hierarchically and informed by sensory data in such a way that learning occurs when a system effectively minimizes free energy (Friston, 2010b). Here, motor commands are proprioceptive predictions, as specific muscle movements (internal frame of reference) are mapped onto an external frame of reference (e.g., vision).

Furthermore, it has been suggested that for biological systems, “It becomes important not only to track the constantly fluctuating instantaneous errors, but also to pay attention to the dynamics of error reduction over longer time scales” (Kiverstein, Miller, & Rietveld, 2019, 2856). Rate of change in prediction error is relevant for epistemic value and novelty-seeking situations. In other words, this mechanism permits an agent to monitor how good it is in performing an action, and it has been suggested as the basis for intrinsic motivation and value-related learning (Kiverstein et al., 2019; Kaplan & Friston, 2018). Therefore, prediction error and its reduction rates might signal the expectations on the learnability of particular situations (Van de Cruys, 2017).

Currently, predictive coding is the most accepted candidate to model how predictive processing principles are manifested in the brain, namely, those laid out by the FEP (Friston, 2009; Buckley et al., 2017). It is a framework for understanding redundancy reduction and efficient coding in the brain (Huang & Rao, 2011) by means of neuronal message passing among different levels of cortical hierarchies (Rao & Ballard, 1999). “Hierarchical predictive coding” suggests that the brain predicts its sensory inputs on the basis of how higher levels provide predictions about lower-level activation until eventually making predictions about incoming sensory

information (Friston, 2002, 2005). Active inference enables predictive coding in a prospective way, where actions attempt to fulfill sensory predictions by minimizing prediction error (Friston et al., 2011).

In this framework, the minimization of prediction error occurs through recurrent message passing within the hierarchical inference (Friston, 2010b). Therefore, the changes in higher levels are driven by the forward flow of the resultant prediction errors in the lower level to optimize top-down predictions until the prediction error is minimized (Friston, 2002, 2010b).

Predictive coding is closely related to Bayes' formulations, from the explanation of how "hierarchical probabilistic generative models" are encoded in the brain to the manner in which the whole system deals with uncertainty. Furthermore, the PEM hypothesis suggests that the brain can be conceived as being "literally Bayesian" (Hohwy, 2013, 17).

However, there is an increasing number of predictive coding variants, for example, there are differences in the algorithms and in the type of generative model they use (Spatling, 2017) and in the excitatory or inhibitory properties of the hierarchical connections (e.g., Rao & Ballard, 1999; Spatling, 2008, among others). "These issues matter when it comes to finding definitive empirical evidence for the computational architectures entailed by predictive coding" (Friston, 2019, 3).

All of these frameworks provide new ways to solve the perception-action control problem in cognitive robotics (Schillaci, Hafner et al., 2016). In the previous couple of decades, the standard solution was the use of paired inverse-forward models in what is known as optimal control theory (OCT). In OCT, a copy of a motor command predicted by an inverse model or controller is passed to a forward model that in turn predicts the sensory consequences of the execution of the movement (Wolpert, Ghahramani, & Jordan, 1995; Wolpert & Kawato, 1998; Kawato, 1999). This leads to multiple implementations using artificial agents with different computational approaches (Demiris & Khadhour, 2006; Möller & Schenck, 2008; Escobar-Juárez, Schillaci, Hermosillo-Valadez, & Lara-Guzmán, 2016; Schillaci, Ritter, Hafner, & Lara, 2016). OCT presents a number of difficult issues to solve, such as the ill-posed problem of learning an inverse model.

On the other hand, in predictive processing, optimal movements are understood in terms of inference and beliefs, and not by the optimization of a value function of states as being the causal explanation of movement (Friston, 2011). Therefore, there are no desired consequences, because experience-dependent learning generates prior expectations, which guide perceptual and active inference (Friston et al., 2011). Contrary to OCT, in predictive processing there are no rewards or cost functions to optimize behavior. Optimal behavior minimizes variational free energy, and cost functions are replaced by priors about sensory states and their transitions (Friston, Samothrakis, & Montague, 2012). Understanding movement as a

matter of beliefs for generating inferences removes the problem of learning an inverse model.

Therefore, predictive processing suggests that there is no need for an inverse model and, thus, for any efference copy of the motor command as input to a forward model. The mere existence of the efference copy of the motor command is a controversial issue (Dogge, Custers, & Aarts, 2019; Pickering & Clark, 2014). The core mechanism in predictive processing is an integral forward model (Pickering & Clark, 2014), better known as a generative model, in which motor commands are replaced by proprioceptive top-down predictions, mapping prior beliefs to sensory consequences (Friston, 2011; Clark, 2015; Friston, Samothrakis et al., 2012). Top-down predictions can be seen as control states based on an extrinsic frame of reference (world-centered-limb position) that are translated into intrinsic muscle-based coordinates that are then fulfilled by the classical reflex arcs (Friston, 2011). Minimizing proprioceptive prediction error brings the action about, fulfilling sensory predictions (Friston et al., 2011).

### 3 Implementations

---

In this section, we review implementation studies inspired by the models and frameworks described in the previous section. Review papers can be found in the literature. This work focuses mostly on robotics research, which has been developing quite rapidly in the last couple of years. We also review a number of nonrobotic studies, in particular those having important aspects that have not received enough exploration in robotics. By highlighting them, this work aims at encouraging experimental research in embodied cognitive robotics.

We are certain that there could be work that is not mentioned in this letter. The omission is not intentional. Articles have been selected according to two criteria. First, the authors mention in their work any of the frameworks described in the previous section. Second, although the authors do not explicitly mention these frameworks, it is our understanding that these works could well enter the discussion and bring interesting topics and questions to the table. This includes some nonrobotic works. Deriving from the descriptions in the previous section, the following items have been considered as relevant to analyze the literature in cognitive robotics:

- **(Bay)** Bayesian/probabilistic framework. Does the study adopt a Bayesian or probabilistic formalization?
- **(PW)** Precision weights. Top-down predictions and bottom-up prediction errors are dynamically weighted according to their expected reliability.
- **(FofI)** Flow of information. Predictions flow top-down, while the difference between predictions and real sensory information (i.e., prediction error) flows bottom-up in the model.

- **(HP)** Hierarchical processing. The model presents a hierarchical structure for the processing of information.
- **(IM)** Inverse model. The work discusses the benefits or challenges of using an inverse model, as is the case in OCT.
- **(Mod)** Modalities. Which modalities are tackled in the proposed model.
- **(BC)** Beyond motor control and estimation of body states. Most of the reviewed studies adopt predictive processing frameworks to control robot movements. This attribute is defined to highlight studies that take a step further by addressing aspects of the framework that may help in understanding or implementing higher-level cognitive capabilities.

The selected studies are summarized in Tables 1 and 2. Table 1 classifies each study according to the attributes mentioned above, and Table 2 provides an overview of some implementation details of these works:

- **Training:** The generative model used in the study is either precoded or trained. If applicable, this specifies what type of learning algorithm (i.e., online or offline) has been employed.
- **Data generation:** If applicable, this specifies how the training data have been generated.
- **Agent:** The type of artificial system used in the experiment.
- **Generative model:** The name, or acronym, of the generative model implemented in the study. Some studies may have not implemented any generative model but used instead the forward kinematics provided by the robot manufacturer.
- **Aim:** What cognitive or motor task has been modeled.

**3.1 Robotic Implementations.** The analysis of the literature starts with one of the first robotic implementations of predictive processing. Tani and Nolfi (1999) present a two-layer hierarchical architecture that self-organizes expert modules. Each expert module is a recurrent neural network (RNN). The bottom layer of RNNs is trained and responds to different types of sensory and motor inputs. The upper set of experts serves as a gating mechanism for the lower-level RNNs. The computational model has been deployed onto a simulated mobile robot for a navigation task. The architecture is trained in an online fashion. After a short period of time, the gating experts specialize in navigating through corridors, right and left turns, and T-junctions. The free parameters of the architecture are trained online using the backpropagation-through-time algorithm (Rumelhart, Hinton, & Williams, 1986). However, as the authors point out, a limitation of the architecture is that it uses only the bottom-up flow of information, without integrating top-down predictions to modulate the activation of lower levels. Tani (2019) provides a thorough review of related neurorobotics experiments, many of which were carried out in the author's laboratory. An

Table 1: Summary of the Main Characteristics of the Reviewed Literature.

| Article                        | Bay | PW | Fofl | HP | IM | Mod   | BC | Aim                             |
|--------------------------------|-----|----|------|----|----|-------|----|---------------------------------|
| Robotic studies                |     |    |      |    |    |       |    |                                 |
| Tani and Nolfi (1999)          | -   | -  | -    | ✓  | -  | V     | -  | Safe navigation                 |
| Ahmadi and Tani (2019)         | ✓   | -  | ✓    | ✓  | -  | PV    | -  | Movement imitation              |
| Ahmadi and Tani (2017)         | -   | -  | ✓    | ✓  | -  | PV    | -  | Movement imitation              |
| Baltieri and Buckley (2017)    | ✓   | ✓  | ✓    | -  | ✓  | L     | -  | Gradient following              |
| Hwang et al. (2018)            | -   | -  | ✓    | ✓  | -  | PV    | -  | Gesture imitation               |
| Idei et al. (2018)             | ✓   | ✓  | ✓    | -  | -  | PV    | ✓  | Simulation of autistic behavior |
| Lanillos and Cheng (2018)      | ✓   | -  | ✓    | -  | -  | PV(T) | -  | Body pose estimation            |
| Lanillos et al. (2020)         | ✓   | -  | ✓    | -  | -  | PV    | ✓  | Self-other distinction          |
| Murata et al. (2015)           | ✓   | -  | ✓    | ✓  | -  | PV    | -  | Human-robot interaction         |
| Ohata and Tani (2020)          | ✓   | -  | ✓    | ✓  | ✓  | PV    | ✓  | Multimodal imitation            |
| Oliver et al. (2019)           | ✓   | -  | ✓    | -  | -  | PV    | -  | Visuomotor coordination         |
| Park et al. (2018)             | -   | -  | ✓    | ✓  | -  | PV    | -  | Arm control                     |
| Pezzato et al. (2020)          | ✓   | -  | ✓    | -  | ✓  | P     | -  | Arm control                     |
| Pio-Lopez et al. (2016)        | ✓   | ✓  | ✓    | ✓  | ✓  | PV    | -  | Control and body estimation     |
| Sancaktar and Lanillos (2019)  | ✓   | -  | ✓    | -  | -  | PV    | -  | Control and body estimation     |
| Schillaci, Ciria et al. (2020) | -   | -  | -    | -  | ✓  | PV    | ✓  | Goal regulation, emotion        |
| Annabi et al. (2020)           | ✓   | -  | -    | -  | ✓  | PV    | -  | Simulation arm control          |
| Zhong et al. (2018)            | -   | -  | ✓    | ✓  | -  | PV    | -  | Movement generation             |
| Nonrobotic studies             |     |    |      |    |    |       |    |                                 |
| Allen et al. (2019)            | ✓   | ✓  | ✓    | -  | -  | IV    | ✓  | Emotional inference             |
| Baltieri and Buckley (2019)    | ✓   | -  | ✓    | -  | ✓  | P     | -  | 1 DoF Control                   |
| Friston et al. (2015)          | ✓   | ✓  | ✓    | ✓  | -  | RO    | ✓  | Exploration versus exploitation |
| Huang and Rao (2011)           | ✓   | -  | ✓    | ✓  | -  | V     | -  | Visual perception               |
| Oliva et al. (2019)            | ✓   | ✓  | -    | -  | -  | V     | ✓  | PW development                  |
| Philippssen and Nagai (2019)   | ✓   | ✓  | -    | -  | -  | V     | ✓  | PW & represent. drawing         |
| Tschantz et al. (2020)         | ✓   | -  | ✓    | -  | -  | RO    | ✓  | Epistemic behaviors             |

Note: Bay: Bayesian / probabilistic framework; PW: implements precision-weighting; Fofl: tackles bottom-up/top-down flows of information; HP: implements hierarchical processing; IM: discusses the need of inverse models; Mod: modalities addressed in the experiment (P: proprioception, V: visual; T: tactile; I: interoceptive; L: luminance as chemo-trail; RO: simulated rewards and observation; BC: the study goes beyond motor control and estimation of body states.

Table 2: Summary of the Implementations in the Reviewed Literature.

| Article                        | Training   | Data generation                    | Agent              | Generative model                |
|--------------------------------|------------|------------------------------------|--------------------|---------------------------------|
| Robotic studies                |            |                                    |                    |                                 |
| Tani and Nolfi (1999)          | Online     | Direct learning                    | Mobile agent       | RNN                             |
| Ahmadi and Tani (2019)         | Offline    | Direct teaching                    | Humanoid           | PV-RNN                          |
| Ahmadi and Tani (2017)         | Offline    | Direct teaching                    | Humanoid           | MTRNN                           |
| Baltieri and Buckley (2017)    | Online     | Exploration                        | Mobile agent       | Agent dynamics                  |
| Hwang et al. (2018)            | Offline    | Direct teaching                    | Simulated humanoid | VMDNN                           |
| Idei et al. (2018)             | Offline    | Recorded sequences                 | Humanoid           | S-CTRNN with PB                 |
| Lanillos and Cheng (2018)      | Offline    | Random movements                   | Humanoid           | Gaussian Process Regression     |
| Lanillos et al. (2020)         | Retrain    | Left-right arm movement            | Humanoid           | Mixture density network         |
| Murata et al. (2015)           | Offline    | Motionese                          | Humanoid           | S-CTRNN                         |
| Ohta and Tani (2020)           | Offline    | Human demonstrations               | Humanoid           | Multiple PV-RNN                 |
| Oliver et al. (2019)           | None       | N.A.                               | Humanoid           | Forward kinematics              |
| Park et al. (2018)             | Dev.learn. | Sets of actions                    | Humanoid           | RNN/PB                          |
| Pezzato et al. (2020)          | None       | N.A.                               | Industr.rob.       | Set-points                      |
| Pio-Lopez et al. (2016)        | None       | N.A.                               | Humanoid           | Forward kinematics              |
| Sancaktar and Lanillos (2019)  | Offline    | andom explanation, direct teaching | Humanoid           | Convolutional decoder           |
| Schillaci, Ciria et al. (2020) | Online     | Goal-directed exploration          | Simulated robot    | Convolutional AE, SOM, DeepNN   |
| Annabi et al. (2020)           | Offline    | Exploration                        | Simulated arm      | SOM, RNN                        |
| Zhong et al. (2018)            | Offline    | Recorded sequences                 | Simulated robot    | Convolutional LSTM              |
| Nonrobotic studies             |            |                                    |                    |                                 |
| Allen et al. (2019)            | None       | N.A.                               | Minim.agent        | Markov Decision Process         |
| Baltieri and Buckley (2019)    | Online     | N.A.                               | 1 DoF agent        | System dynamics                 |
| Friston et al. (2015)          | None       | N.A.                               | Simulated rat      | POMDP                           |
| Huang and Rao (2011)           | Offline    | Image data set                     | –                  | Hierarchical neural model       |
| Oliva et al. (2019)            | Offline    | Precoded trajectories              | Simulated drawing  | S-CTRNN                         |
| Philippssen and Nagai (2019)   | Offline    | Human demonstrations               | Simulated drawing  | S-CTRNN                         |
| Tschantz et al. (2020)         | Online     | RL exploration                     | OpenAlsim          | Gaussian, Laplace approximation |

Note: Training: Which type of training, if applicable, has been performed on the generative model; Data generation: How the training data were generated; Agent: Which type of artificial system has been used; Generative model: The name of the machine learning tool, if applicable, that was adopted for training the generative model; Aim: Which cognitive or motor task has been modeled. N.A.: Not applicable.

interesting implementation is described in Hwang, Kim, Ahmadi, Choi, and Tani (2018), which the authors refer to as a predictive coding model. The adopted network is a multilayer hierarchical architecture encoding visual and proprioceptive information. Although the work is far from the formulations laid in the free-energy principle (Friston, 2009), the VMDNN (predictive visuo-motor deep dynamic neural network) performs very similar operations. These include the generation of actions following a prediction error minimization scheme and the usage of the same model structure for action generation and recognition. Hwang et al. (2018) claim that “the proposed model provides an online prediction error minimization mechanism by which the intention behind the observed visuo-proprioceptive patterns can be inferred by updating the neurons’ internal states in the direction of minimizing the prediction error” (3). It is worth noting that such an update does not refer to model weights, only to the state of the neurons. The training of the model is performed in a supervised fashion. The error being minimized is the difference between a signal generated through kinesthetic teaching (i.e., where a human experimenter manually directs the movements of the robot limb) and the model predictions. The lateral connections between modalities at each layer of the hierarchy are an interesting aspect of the network.

Another relevant work from the same group (Ahmadi & Tani, 2019) stands out for its formulation of active inference and a training strategy based on variational Bayes recurrent neural networks.

Finally, Ahmadi and Tani (2017) propose a multiple timescale recurrent neural network (MTRNN), which consists of multiple levels of subnetworks with specific temporal constraints on each layer. The model processes data from three modalities and is capable of generating long-term predictions in both open-loop and closed-loop fashions. During closed-loop output generation, internal states of the network can be inferred through error regression. The network is trained in an open-loop manner, modifying free parameters using the error between desired states and real activation values.

A common characteristic of the implementations reviewed so far is that learning and testing are decoupled. During the testing phase, prediction errors flow bottom-up, and the network’s “internal state is modified in the direction of minimizing prediction error via error regression” (Ahmadi & Tani, 2017, 4). This implies that the network’s weights are not modified after training. In most of their work, Tani and colleagues use mathematical formulations based on connectionist networks, which are different from those proposed by Friston (2009); nonetheless, the work is conceptually related to predictive coding and active inference. More recently, authors have used explicitly variational inference (e.g., Matsumoto & Tani, 2020; Jung, Matsumoto, & Tani, 2019). An illustrative architecture, which comprises most of the characteristics of the networks used by these authors, can be seen in Figure 1 in Hwang et al. (2018).



A similar approach has been presented by Murata et al. (2015), who propose an RNN-based model, stochastic continuous-time RNN (S-CTRNN). The framework integrates probabilistic Bayesian schemes in a recurrent neural network. Network training is performed offline using temporal sequences under two learning conditions: with and without presenting actions that reveal distinctive characteristics amplifying or exaggerating meaning and structure within bodily motions (also named *motionese*; Brand, Baldwin, & Ashburn, 2002). Training data are obtained through kinesthetic teaching on the robot directed by an experimenter. The loss function of the optimization process considers the sum of log uncertainty and precision-weighted prediction error. This is formally equivalent to free energy as proposed in active inference.

In trying to explain the underlying mechanisms causing different types of behavioral rigidity of the autism spectrum, Idei et al. (2018) adopt an S-CTRNN with parametric bias (PB) as the computational model for simulating aberrant sensory precision in a humanoid robot. In this study, S-CTRNNs learn to estimate sensory variance (precision) and adapt to different environments using prediction error minimization schemes. Learning is performed in an offline fashion using prerecorded perceptual sequences. "The objective of the learning is to find the optimal values of the parameters (synaptic weights, biases, and internal states of PB units) minimizing negative log-likelihood, or precision weighted prediction error" (Idei et al., 2018). Once trained, the network is capable of reproducing target visuo-proprioceptive sequences. In the test phase following the learning one, only the internal states of the PB units are updated in an online fashion while keeping the other parameters fixed. The study simulates increased and decreased sensory precision by altering estimated sensory variance (inverse of their precision). This is performed by modulating a constant in the activation function of the variance units of the trained model. Interestingly, the authors report abnormal behaviors in the robot, such as freezing and inappropriate repetitive behaviors, correlated to specific modulation of the sensory variance. In particular, increased sensory variance reduces the precision of prediction error, thus freezing the PB states of the network and, consequently, the robot behavior. Decreasing sensory variance instead leads to unlearned repetitive behavior, likely due to the fixation of the PB states on suboptimal local solution during prediction error minimization.

Ohata and Tani (2020) extend the predictive coding-inspired variational recurrent neural network (PV-RNN) presented by Ahmadi and Tani (2019) in a multimodal imitative interaction experiment with a humanoid robot. Modalities (proprioception and vision), each encoded with a multilayered PV-RNN, are connected through an associative PV-RNN module. The associative module generates the top-down prior, which is then fed to both the proprioception and vision modules. Each sensory module also generates top-down priors conditioned by the other flows. The authors show how metapriors assigned to the proprioception and vision modules have an impact on the learning process and the performance of the error

regression. Modulating the Kullback-Leibler divergence (KLD) term in the error minimization scheme leads to better regulation of multimodal perception, which would be otherwise biased toward a single modality. Stronger regulation of the KLD term also leads to higher adaptivity in a human-robot imitation experiment.

Park, Lim, Choi, and Kim (2012) propose an architecture based on self-organizing maps and transition matrices for studying three different capabilities and phenomena: performing trajectories, object permanence, and imitation. Interestingly, the architecture features a hierarchical self-organized representation of state spaces. However, no bidirectional (top-down/bottom-up) flow of information as in the previous studies is implemented. Moreover, the models are in part precoded. In a more recent study, Park, Kim, and Nagai (2018) adopt a recurrent neural network with parametric bias (RNNPB) with recurrent feedback from the output layer to the input layer. As in Tani (2019), training and testing are decoupled and the optimization is based on the backpropagation-through-time algorithm. The optimization of the network parameters uses the prediction error between a generated motor action and a reference action. Remarkably, this work analyzes the developmental dynamics of the parameter space in terms of prediction error. Experiments are carried out on a simulated two degrees-of-freedom robot arm and a Nao humanoid robot, where goal-directed actions are generated using the RNNPB.

An interesting series of studies has been produced by Lanillos and colleagues. Lanillos and Cheng (2018) present an architecture that combines generative models and a probabilistic framework inspired by some of the principles of predictive processing. The architecture is employed to estimate body configurations of a humanoid robot, using three modalities (proprioceptive, vision, and touch). In the literature, the way the brain integrates multimodal streams in similar error minimization schemes is still under debate. Some authors suggest that the integration of different streams of unimodal sensory surprise occurs in hierarchically higher multimodal areas (Limanowski & Blankenburg, 2013; Apps & Tsakiris, 2014; Clark, 2013; Pezulo, Rigoli, & Friston, 2015), and therefore multimodal predictions and prediction errors would be generated (Friston, 2012). Lanillos and Cheng (2018) apply an additive formulation of unimodal prediction errors: (1) prior error, that is, the “error between the most plausible value of the body configuration and its prior belief”; (2) proprioceptive error, that is, the distance between joint angle readings and joint angle samples generated by a Normal distribution; and (3) visual error, that is, the distance between observed end-effector image coordinates and those predicted by a visual generative model.

The proposed minimization scheme adjusts the prior on body configuration by summing up the additive multimodal error, while the system is exposed to multimodal observations. As in Tani’s work, training and testing are decoupled. The generative models are pretrained using gaussian process regression. In particular, a visual forward model maps

proprioceptive data (position of three joints) to visual data (image coordinates of the end effector), whereas a proprioceptive model generates joint angles from a Normal distribution representing the joint states. Training data are recorded offline from a humanoid robot executing random trajectories. Another generative model is created for the tactile modality as a function of the visual generative model. This model is used in a second experiment to translate the end-effector positions to the spatial locations on the robot arm touched by an experimenter, in order to correct visual estimations.

A follow-up work (Oliver, Lanillos, & Cheng, 2019) applies an active inference model for visuomotor coordination in the humanoid robot iCub. The framework controls two subsystems of the robot body, the head and one arm. An attractor model drives actions toward goals. Goals are specified in a visual domain—encoded as linear velocity vectors toward a goal, whose 3D position is estimated using stereo vision and a color marker—and transformed using a Moore-Penrose pseudoinverse Jacobian matrix into linear velocities in the 4D joint space of the robot. Similarly, visual goals are transformed into joint velocity goals for the head subsystem. The authors assume normally distributed noise in the sensory inputs. Sensor variances and action gains are pretuned and fixed during the experiments. Although no generative models are trained in this experiment (iCub's forward kinematics functions are used), the authors show that minimizing Laplace-encoded free energy through gradient descent leads to reaching behaviors and visuomotor coordination. Similarly, Pezzato, Ferrari, and Corbato (2020) present an active inference framework using a precoded controller and a generative function. The study aims at controlling the movements of an industrial robotic platform using active inference and comparing its adaptivity and robustness to another state-of-the-art controller for robotic manipulators, namely, the model reference adaptive controller (MRAC).

Lanillos, Cheng, and Pages (2020) extend the active inference implementation presented in Oliver et al. (2019). In this study, the visual generative model is pretrained using a probabilistic neural network (mixture density network, MDN). Inverse mapping is performed through the backward pass of the MDN of the most plausible gaussian kernel. The system retrains the network from scratch whenever the sensory inputs are too far from its predictions. Differently from Oliver et al. (2019), visual inputs consist of movements estimated through an optical flow algorithm. The generative model thus maps joint angles to the 2D centroid of a moving blob detected from the camera. A deep learning classifier is then trained to label joint velocities and optical flow inputs as self-generated or not.

Sancaktar and Lanillos (2019) apply a similar approach on the humanoid robot Nao. The minimization scheme uses a pretrained generative model for the visual input, that is, a convolutional decoder-like neural network. Training data are collected through a combination of random babbling and kinesthetic teaching. The generative model maps joint angles to visual inputs, as in Lang, Schillaci, and Hafner (2018). When computing the

likelihood for the gradient descent, the density defining the visual input is created as a collection of independent gaussian distributions centered at each pixel. In the minimization scheme, the visual prediction error multiplied by the inverse of the variance is calculated by applying a forward pass and a backward pass to the convolutional decoder. The approach is interesting, but studies have pointed at questionable aspects of the biological plausibility of backpropagating errors. This refers, in particular, to the lack of local error representations in ANNs and at the symmetry between forward and backward weights, which is not always present in cortical networks (Whittington & Bogacz, 2019). As in the previous series of experiments, active inference is used to control the robot arm movement in a reaching experiment.

Pio-Lopez, Nizard, Friston, and Pezzulo (2016) present a proof-of-concept implementation of a control scheme based on active inference using the 7 degrees-of-freedom arm of a simulated PR2 humanoid robot. The control scheme is adopted to perform trajectories towards predefined goals. Authors highlight that such a scheme eliminates the need of an inverse model for motor control as “action realizes the (sensory) consequences of (prior) causes” (Pio-Lopez et al., 2016, 9). A generative model maps causes to actions, where causes are seen as “forces that have some desired fixed point or orbit” (Pio-Lopez et al., 2016, 9), as sensed by proprioception. Proprioceptive predictions are thus realized in an open-loop fashion, by means of reflex arcs.

This framework, which employs a hierarchical generative model minimizes the KL-divergence between the distribution of the agent’s priors and that of the true posterior distribution, which represents the updated belief given the evidence. Pio-Lopez et al. (2016) point out that more complex behaviors require the design of equations of motion. The question on the scalability of such an approach for cognitive robotics remains open.

Although not adopting an active inference approach, Schillaci, Ciria, & Lara (2020) present a study where intrinsically motivated behaviors are driven by error-minimization schemes in a simulated robot. The proposed architecture generates exploratory behaviors toward self-generated goals, leverages computational resources, and regulates goal selection and the balance between exploitation and exploration through a multilevel monitoring of prediction error dynamics. The work is framed within the study of the underlying mechanisms of motivation and the emergence of emotions that drive behaviors and goal selection to promote learning. Scholars such as Van de Cruys (2017), Kiverstein et al. (2019), and Hsee and Abelson (1991) argue that what motivates engagement in a behavior is not just the final outcome but the satisfaction that emerges from the pattern and the velocity of an outcome over time.<sup>1</sup> “If one . . . assumes that people not only

<sup>1</sup> Here we intend the desired outcome of an event or of an activity. As for the velocity of an outcome, we intend the velocity, or the rate, at which such desired goal is achieved.

passively experience satisfaction, but actively seek satisfaction, then one can infer an interesting corollary from the velocity relation: People engage in a behavior not just to seek its actual outcome, but to seek a positive velocity of outcomes that the behavior creates over time" (Hsee & Abelson, 1991, 346).

The system proposed by Schillaci, Ciria et al. (2020) monitors prediction error dynamics over time and at different levels, driving behaviors toward those goals that are associated with specific patterns of prediction error dynamics. The system also modulates exploration noise and leverages computational resources according to the dynamics of the overall learning performance. Learning is performed in an online fashion, where image features, compressed using a pretrained convolutional autoencoder, are fed into a self-organizing neural network for unsupervised goal generation and into an inverse-forward models pair for movement generation and prediction error monitoring. The models are updated in an online fashion, and an episodic memory system is adopted to reduce catastrophic forgetting issues. Actions are generated toward goals associated with the steepest descent in low-level prediction error dynamics.

A similar approach for the self-generation of goals has been employed by Annabi, Pitti, and Quoy (2020) in a simulated experiment where a two-degrees-of-freedom robotic arm has to learn how to write digits. The proposed architecture learns sequences of motor primitives based on a free energy minimization approach. The system combines recurrent neural networks for trajectories encoding with a self-organizing system for goal estimation, which is trained on data generated through random behaviors. In the experiments, the system incrementally learns motor primitives and policies, using a predefined generative forward model. Free energy minimization is used for action selection.

Zhong, Cangelosi, Zhang, and Ogata (2018) present a hierarchical model consisting of a series of repeated stacked modules to implement active inference in simulated agents. Each layer of the network contains different modules, including generative units implemented as convolutional recurrent networks (long short-term memory networks, LSTM). In the hierarchical architecture, predictions and prediction errors flow in top-down and bottom-up directions, respectively. Generative units are trained in an offline learning session during two simulated experiments.

It is worth noting that all the work reviewed in this section makes use of different forms of prediction error minimization schemes to obtain working models and controllers.

---

In the context of learning, a goal could be merely the reduction of prediction error. The velocity of the outcome here would correspond to the rate of reduction of the prediction error, that is, how fast or slow is prediction error minimized.

**3.2 Nonrobotic Implementations.** A wide number of nonrobotic studies on predictive processing have been produced. This section opens only a small window on this literature. Nevertheless, promising directions for cognitive robotics research on predictive processing can be characterized from the few samples reported here.

The issue of scalability highlighted on the active inference study of Pio-Lopez et al. (2016) is also apparent in the work of Baltieri and Buckley (2019), where the authors design an active inference-based linear quadratic gaussian controller to manipulate a one-degree-of-freedom system. The study aims at showing that such a controller can achieve goal positions without the need of an efference copy, as in optimal control theory (OCT).

Similar basic proofs-of-concept are presented by Tschantz et al. (2020) and Baltieri and Buckley (2017), where active inference is used to model bacterial chemo-taxis in a minimal simulated agent. Tschantz et al. (2020) focus on an action-oriented model that employs goal-directed (instrumental) and information-seeking (epistemic) behaviors when learning a generative model. Different error-minimization strategies are tested, generating epistemic, instrumental, random behaviors or expected free energy-driven ones. Tschantz et al. (2020) show that active inference balances exploration and exploitation and suggest that “[they] are both complementary perspectives of the same objective function—the minimization of expected free energy” (19). The model is not hierarchical, but it fully exploits the proposals of active inference. In the other interesting proof-of-concept, Baltieri and Buckley (2017) present a Braitenberg-like vehicle where behaviors are modulated according to predefined precision weights.

Friston et al. (2015) also address the exploration-exploitation dilemma. They argue that when adopting Bayes’ optimal behavior under the free energy principle, epistemic, intrinsic value is maximized until there is no further information gain, after which exploitation is assured through maximization of extrinsic value (i.e., the utility of the result of an action). In fact, epistemic actions can bring the agent far from a goal. Nonetheless, they can be used to plan a path to a goal with greater confidence. Adopting the formalism of partially observed Markov decision processes, Friston et al. (2015) present a simulated experiment where an agent (i.e., a rat) navigates through a T-shaped maze to show the role of epistemic value in resolving uncertainty about goal-directed behavior. Moreover, the authors discuss an aspect of the Bayesian framework, that is, the role of the precision (i.e., the inverse of the variance) of the posterior belief, which is estimated from the prior belief and the likelihood of the evidence, about control states<sup>2</sup> as a message-passing channel. According to this view, precision is associated with dopaminergic responses, which has been interpreted in terms of

<sup>2</sup>In the generative model, a control state corresponds to the hidden cause of an action. “This means the agent has to infer its behavior by forming beliefs about control states, based upon the observed consequences of its action” (Friston et al., 2015, 190).



changes in expected value (e.g., reward prediction errors). In brief, changes in precision would correlate with changes in exploratory or exploitative behaviors.

In a follow-up study, Schwartenbeck et al. (2019) present an architecture that has an implicit weighting of the exploitation and (goal-directed) exploration tendencies, determined by the precision of prior beliefs and the degree of uncertainty about the world. Two mechanisms for goal-directed exploration are implemented in the rat-within-a-maze simulated setup: model parameter exploration and hidden state exploration. In the former active learning strategy, the agents forage for information about the correct parameterization of the observation model, represented as a Markovian model in the study. Here, parameters are the set of arrays encoding the Markovian transition probabilities, that is, the mapping between hidden states and observations and the transition between hidden states. In the latter active inference strategy, agents aim at gathering information about the current (hidden) state of the world—for example, the current context. In particular, they sample the outcomes associated with high uncertainty, only when these are informative for the representation of the task structure. Much like a standard intrinsic motivation approach, Schwartenbeck et al. (2019) appeal to the need for random sampling when the uncertainty about model parameters and hidden states (goal-exploration strategies) fails to inform behavior. The aim of this work is to understand “the generative mechanisms that underlie information gain and its trade-off with reward maximization” (Schwartenbeck et al., 2019, 45), but as they note, how to scale up these mechanisms to more complicated tasks is an open challenge.

Precision weighting is also one of the main focuses of the predictive coding study carried out by Oliva, Philippsen, and Nagai (2019). Interestingly, they analyze the variations of the precision of prior prediction of a recurrent (S-CTRNN) generative model over a developmental process. The model learns to estimate stochastic time series (two-dimensional trajectory drawings), thus providing an estimate of the variance of the input data. The framework “shares crucial properties with the developmental process of humans in that it naturally switches from a strong reliance on sensory input at an early learning stage to a proper integration of sensory input and own predictions at later learning stages” (Oliva et al., 2019, 254). This is correlated to a reduction of the prediction error and the estimated (prior) variance over time during learning. Some formulations of the problem in this work are, however, problematic. In particular, in Oliva et al. (2019) the posterior is computed naively by multiplying the likelihood and the prior using the basic Bayesian formula, and learning is performed only for maximizing the likelihood. In a follow-up work (Philippsen & Nagai, 2019), the framework is applied to simulate the generation of representational drawings—i.e., drawings that represent objects—in infants and chimpanzees. The authors observe that stronger reliance on the prior (hyperprior) enables the network to perform representational drawings like those produced by



children, whereas a weak reliance on the prior produces highly accurate lines but fails to produce missing parts of the representational drawings, as observed in chimpanzees. Results suggest that chimpanzees' and humans' "differences in representational drawing behavior might be explainable by the degree to which they take prior information into account" (Philippesen & Nagai, 2019, 176).

Allen, Levy, Parr, and Friston (2019) study active inference in a multimodal domain, simulating interactions between interoceptive cardiac cycle and exteroceptive (visual) perception. The work hypothesizes that effects of cardiac timing on perception could arise as a function of periodic sensory attenuation. This study does not involve any robotic implementation or any learning or control task. However, related implementations are mostly missing in the literature; therefore, we believe it is worth mentioning in this letter.

#### 4 Discussion

---

This work has reviewed a series of studies on robotics and nonrobotics that have adopted the paradigm of predictive processing under different forms. Tables 1 and 2 provided a general overview of the main aspects as well as the differences of these studies.

It is certainly standing out to what length the robotics research and the nonrobotics models have addressed tasks that go beyond perception and motor control, traditionally the focus of predictive processing studies. Limited cognitive robotics research has addressed the scaling up of the predictive processing paradigm toward higher cognitive capabilities. Computational studies on minimal simulated systems have suggested that specific aspects, such as precision weighting, may bridge this gap.

Embodied robotic systems seem to be the most appropriate experimental platforms not only for studying cognitive development within the predictive processing framework but also for extending this framework to a broader range of modalities and behavioral possibilities. In fact, another aspect of the robotics research reviewed in this letter and worth highlighting is that almost the majority of them<sup>3</sup> address only proprioception and a single exteroceptive modality, vision. Little attention in the robotics community has been posed on how multiple exteroceptive modalities (e.g., vision, haptic, and auditory), as well as interoceptive ones (Seth & Tsakiris, 2018), can be integrated in prediction error-minimization schemes. Studies such as those from Tschantz et al. (2020), Friston et al. (2015), Schwartenbeck et al. (2019), and Schillaci, Ciria et al. (2020) have discussed epistemic and emotional value, homeostatic drives, and intrinsic motivation that

---

<sup>3</sup>Lanillos and Cheng (2018) address also the tactile modality in their study, but do not fully integrate it in the error minimization scheme.

regulate behaviors. Interesting research directions for robotics should include extending this to multimodal self-generated goals and to combinations of fixed homeostatic goals and dynamic ones.

Another important point concerns precision weighting. As in predictive processing, this is assigned a prominent role in behavior and goal regulation, as well as in perceptual optimization processes. Further cognitive robotics study should explore this path. Most of the nonrobotic implementations adopt a Bayesian or probabilistic formalization of error-minimization schemes. This allows an elegant formulation of the precision in weighting schemes, which consists of the inverse of the variance of the prior and posterior distributions. However, alternative strategies are available for implementing precision weighting-like processes in nonprobabilistic models, including the modulation of neuronal activation or of synaptic weights in artificial neural networks, modulation of firing rates in spiking neural networks, dopaminergic modulation, and the like. There is a wide literature on sensor fusion techniques in the machine learning community that focuses on very related challenges, such as the learning and modulation of the relevance of single sensors in multimodal and predictive settings (Fayyad, Jaradat, Gruyer, & Najjaran, 2020).

A common denominator in all the reviewed implementations is the use of predictions for guiding behavior. However, the implementations adopt different machine learning tools. Works that strictly follow the active inference principles make use of Bayes as their main tool. It is still an open question how all other approaches should be considered in the wider predictive processing framework. So far, most robotics implementations make use of nonvariational deep networks as their main tool. However, the bias of using the Bayesian framework in nonrobotics implementations might hinder the search for other approaches that could have advantages, importantly, in terms of computational cost and the complexity of designing generative models to produce coherent and scaled-up behaviors.

Predictive processing emphasizes the prediction-based learning of a generative model, which predicts incoming sensory signals (Clark, 2015). In optimal control theory, a high computational complexity is required for learning to predict sensory consequences by means of the efference copy and the inverse model. In predictive processing accounts, this complexity is mapped to the learning of a generative model during hierarchical perceptual and active inference (Friston, 2011; Friston, Samothrakis et al., 2012). In this regard, it is still unclear how generative models should be learned due to the complexity that implies modeling the richness of the entire environment (Tschantz et al., 2020). Action-oriented models are a common approach for solving this issue by learning and generating inferences that allow adaptive behavior, even when the world is not modeled in a precise manner (Tschantz et al., 2020; Baltieri & Buckley, 2017; Pezzulo, Donnarumma, Iodice, Maisto, & Stoianov, 2017). It is worth highlighting that despite the relevance of learning for belief updating, most nonrobotic

computational work focuses on inference, not on learning. Actually, learning is almost absent here.

The few nonrobotic models that focus on learning generative models are based on the expected free energy formulations and use simplified agents and behaviors (Tschantz et al., 2020; Baltieri & Buckley, 2017; Ueltzhöffer, 2018; Millidge, 2020). On the contrary, some cognitive robotics implementations do have the emphasis slightly shifted toward the learning of generative models (Lanillos et al., 2020; Ahmadi & Tani, 2017; Idei et al., 2018; Schillaci, Ciria et al., 2020; Schillaci, Pico Villalpando et al., 2020). Yet learning and testing are decoupled in many of these studies, in particular, in those adopting probabilistic methods. This is likely due to the challenges of implementing online learning of probabilistic models, especially in the context of high-dimensional sensory and motor spaces.

It is worth pointing out that in cognitive robotics, a variety of learning methods are used, and just a few of these are equivalent to the free energy principle formulations. Nonetheless, the agents and behaviors that are used are much more complex. For cognitive robotics, it is relevant to explore the reach and possibilities of using generative models for perception, action, and planning. More important, there is a special interest in the tools and methods that can be used for the learning of these models, an area that has been unattended in nonrobotic models using predictive processing principles.

Finally, attention to the temporal aspect of prediction error dynamics has been limited (Kiverstein et al., 2019; Tschantz et al., 2020). Prediction error patterns may be associated with emotional experience (Joffily & Coricelli, 2013). In artificial systems, they are essential components for implementing intrinsically motivated exploration behaviors and artificial curiosity (Oudeyer, Kaplan, & Hafner, 2007; Schillaci, Pico Villalpando et al., 2020; Baldassarre & Mirolli, 2013; Graziano et al., 2011). Recent studies suggest that error dynamics may influence the regulation of computational resources (Schillaci, Ciria et al., 2020) and the emotional valence of actions (Joffily & Coricelli, 2013). We believe that prediction error dynamics represent a promising tool in the exploration of more complex behaviors and tasks in cognitive robotics under the predictive processing paradigm.

## Acknowledgments

G.S. has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement 838861 (Predictive Robots). Predictive Robots is an associated project of the Deutsche Forschungsgemeinschaft (German Research Foundation) Priority Programme, The Active Self. V.H. has received funding from the Deutsche Forschungsgemeinschaft Priority Programme, The Active Self (402790442: Prerequisites for the Development of an Artificial Self). B.L. and A.C. have received funding from the Alexander von

Humboldt Foundation from the Predictive Autonomous Behaviour Internal Models and Predictive Self-Regulation project. We thank the anonymous reviewer for his or her thorough reading of our manuscript and comments, which helped greatly to improve the first version submitted.

## References

- Adams, R. A., Shipp, S., & Friston, K. J. (2013). Predictions not commands: Active inference in the motor system. *Brain Structure and Function*, 218(3), 611–643.
- Adams, R. A., Stephan, K. E., Brown, H. R., Frith, C. D., & Friston, K. J. (2013). The computational anatomy of psychosis. *Frontiers in Psychiatry*, 4, 47.
- Ahmadi, A., & Tani, J. (2017). How can a recurrent neurodynamic predictive coding model cope with fluctuation in temporal patterns? Robotic experiments on imitative interaction. *Neural Networks*, 92, 3–16.
- Ahmadi, A., & Tani, J. (2019). A novel predictive-coding-inspired variational RNN model for online prediction and recognition. *Neural Computation*, 31(11), 2025–2074.
- Allen, M., Levy, A., Parr, T., & Friston, K. J. (2019). *In the body's eye: The computational anatomy of interoceptive inference*. bioRxiv:603928.
- Annabi, L., Pitti, A., & Quoy, M. (2020). *Autonomous learning and chaining of motor primitives using the free energy principle*. arXiv:2005.05151.
- Apps, M. A., & Tsakiris, M. (2014). The free-energy self: A predictive coding account of self-recognition. *Neuroscience and Biobehavioral Reviews*, 41, 85–97.
- Badcock, P. B., Davey, C. G., Whittle, S., Allen, N. B., & Friston, K. J. (2017). The depressed brain: An evolutionary systems theory. *Trends in Cognitive Sciences*, 21(3), 182–194.
- Baldassarre, G., & Mirolli, M. (2013). *Intrinsically motivated learning in natural and artificial systems*. Berlin: Springer.
- Baltieri, M., & Buckley, C. L. (2017). An active inference implementation of phototaxis. In *Artificial Life Conference Proceedings*, 14 (pp. 36–43). Cambridge, MA: MIT Press.
- Baltieri, M., & Buckley, C. L. (2019). *Active inference: Computational models of motor control without efference copy*. researchgate.
- Brand, R. J., Baldwin, D. A., & Ashburn, L. A. (2002). Evidence for “motionese”: Modifications in mothers’ infant-directed action. *Developmental Science*, 5(1), 72–83.
- Brown, H., Adams, R. A., Parees, I., Edwards, M., & Friston, K. (2013). Active inference, sensory attenuation and illusions. *Cognitive Processing*, 14(4), 411–427.
- Buckley, C. L., Kim, C. S., McGregor, S., & Seth, A. K. (2017). The free energy principle for action and perception: A mathematical review. *Journal of Mathematical Psychology*, 81, 55–79.
- Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36(3), 181–204.
- Clark, A. (2015). *Embodied prediction*. Frankfurt am Main: MIND Group.
- Clark, A. (2018). A nice surprise? Predictive processing and the active pursuit of novelty. *Phenomenology and the Cognitive Sciences*, 17(3), 521–534.

- Clark, A. (2020). Beyond desire? Agency, choice, and the predictive mind. *Australasian Journal of Philosophy*, 98(1), 1–15.
- Demekas, D., Parr, T., & Friston, K. J. (2020). An investigation of the free energy principle for emotion recognition. *Frontiers in Computational Neuroscience*, 14.
- Demiris, Y., & Khadhour, B. (2006). Hierarchical attentive multiple models for execution and recognition of actions. *Robotics and Autonomous Systems*, 54(5), 361–369.
- Dogge, M., Custers, R., & Aarts, H. (2019). Moving forward: On the limits of motor-based forward models. *Trends in Cognitive Sciences*, 23(9), 743–753.
- Donnarumma, F., Costantini, M., Ambrosini, E., Friston, K., & Pezzulo, G. (2017). Action perception as hypothesis testing. *Cortex*, 89, 45–60.
- Escobar-Juárez, E., Schillaci, G., Hermosillo-Valadez, J., & Lara-Guzmán, B. (2016). A self-organized internal models architecture for coding sensory-motor schemes. *Frontiers in Robotics and AI*, 3, 22.
- Fayyad, J., Jaradat, M. A., Gruyer, D., & Najjaran, H. (2020). Deep learning sensor fusion for autonomous vehicle perception and localization: A review. *Sensors*, 20(15), 4220.
- Feldman, H., & Friston, K. (2010). Attention, uncertainty, and free-energy. *Frontiers in Human Neuroscience*, 4, 215.
- Friston, K. (2002). Functional integration and inference in the brain. *Progress in Neurobiology*, 68(2), 113–143.
- Friston, K. (2005). A theory of cortical responses. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 360(1456), 815–836.
- Friston, K. (2009). The free-energy principle: A rough guide to the brain? *Trends in Cognitive Sciences*, 13(7), 293–301.
- Friston, K. (2010a). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, 11(2), 127–138.
- Friston, K. (2010b). Is the free-energy principle neurocentric? *Nature Reviews Neuroscience*, 11(8), 605.
- Friston, K. (2011). What is optimal about motor control? *Neuron*, 72(3), 488–498.
- Friston, K. (2012). Prediction, perception and agency. *International Journal of Psychophysiology*, 83(2), 248–252.
- Friston, K. J. (2019). Waves of prediction. *PLOS Biology*, 17(10).
- Friston, K., Adams, R., Perrinet, L., & Breakspear, M. (2012). Perceptions as hypotheses: Saccades as experiments. *Frontiers in Psychology*, 3, 151.
- Friston, K., Kilner, J., & Harrison, L. (2006). A free energy principle for the brain. *Journal of Physiology–Paris*, 100(1–3), 70–87.
- Friston, K., Mattout, J., & Kilner, J. (2011). Action understanding and active inference. *Biological Cybernetics*, 104(1–2), 137–160.
- Friston, K., Rigoli, F., Ognibene, D., Mathys, C., Fitzgerald, T., & Pezzulo, G. (2015). Active inference and epistemic value. *Cognitive Neuroscience*, 6(4), 187–214.
- Friston, K., Samothrakis, S., & Montague, R. (2012). Active inference and agency: Optimal control without cost functions. *Biological Cybernetics*, 106(8–9), 523–541.
- Friston, K. J., Stephan, K. E., Montague, R., & Dolan, R. J. (2014). Computational psychiatry: The brain as a phantastic organ. *Lancet Psychiatry*, 1(2), 148–158.
- Graziano, V., Glasmachers, T., Schaal, T., Pape, L., Cuccu, G., Leitner, J., & Schmidhuber, J. (2011). Artificial curiosity for autonomous space exploration. *Acta Futura*, 4, 41–51.

- Hohwy, J. (2013). *The predictive mind*. New York: Oxford University Press.
- Hsee, C. K., & Abelson, R. P. (1991). Velocity relation: Satisfaction as a function of the first derivative of outcome over time. *Journal of Personality and Social Psychology*, 60(3), 341.
- Huang, Y., & Rao, R. P. (2011). Predictive coding. *Wiley Interdisciplinary Reviews: Cognitive Science*, 2(5), 580–593.
- Hwang, J., Kim, J., Ahmadi, A., Choi, M., & Tani, J. (2018). Dealing with large-scale spatio-temporal patterns in imitative interaction between a robot and a human by using the predictive coding framework. *IEEE Transactions on Systems, Man, and Cybernetics: System*, 50, 1918–1931.
- Idei, H., Murata, S., Chen, Y., Yamashita, Y., Tani, J., & Ogata, T. (2018). A neuro-robotics simulation of autistic behavior induced by unusual sensory precision. *Computational Psychiatry*, 2, 164–182.
- Joffily, M., & Coricelli, G. (2013). Emotional valence and the free-energy principle. *PLOS Comput. Biol.*, 9(6), e1003094.
- Jung, M., Matsumoto, T., & Tani, J. (2019). *Goal-directed behavior under variational predictive coding: Dynamic organization of visual attention and working memory*. arXiv:1903.04932.
- Kaplan, R., & Friston, K. J. (2018). Planning and navigation as active inference. *Biological Cybernetics*, 112(4), 323–343.
- Kawato, M. (1999). Internal models for motor control and trajectory planning. *Current Opinion in Neurobiology*, 9(6), 718–727.
- Kiverstein, J., Miller, M., & Rietveld, E. (2019). The feeling of grip: Novelty, error dynamics, and the predictive brain. *Synthese*, 196(7), 2847–2869.
- Knill, D. C., & Pouget, A. (2004). The Bayesian brain: The role of uncertainty in neural coding and computation. *Trends in Neurosciences*, 27(12), 712–719.
- Kruglanski, A. W., Jasko, K., & Friston, K. (2020). All thinking is “wishful” thinking. *Trends in Cognitive Sciences*, 23, 413–424.
- Lang, C., Schillaci, G., & Hafner, V. V. (2018). A deep convolutional neural network model for sense of agency and object permanence in robots. In *Proceedings of the 2018 Joint IEEE 8th International Conference on Development and Learning and Epigenetic Robotics* (pp. 257–262). Piscataway, NJ: IEEE.
- Lanillos, P., & Cheng, G. (2018). Adaptive robot body learning and estimation through predictive coding. In *Proceedings of the 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 4083–4090). Piscataway, NJ: IEEE.
- Lanillos, P., Cheng, G., & Pages, J. (2020). *Robot self/other distinction: Active inference meets neural networks learning in a mirror*. arXiv:2004.05473.
- Lara, B., Astorga, D., Mendoza-Bock, E., Pardo, M., Escobar, E., & Ciria, A. (2018). Embodied cognitive robotics and the learning of sensorimotor schemes. *Adaptive Behavior*, 26(5), 225–238.
- Limanowski, J., & Blankenburg, F. (2013). Minimal self-models and the free energy principle. *Frontiers in Human Neuroscience*, 7, 547.
- Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. Cambridge, MA: MIT Press.
- Matsumoto, T., & Tani, J. (2020). Goal-directed planning for habituated agents by active inference using a variational recurrent neural network. *Entropy*, 22(5), 564.



- Millidge, B. (2020). Deep active inference as variational policy gradients. *Journal of Mathematical Psychology*, 96, 102348.
- Möller, R., & Schenck, W. (2008). Bootstrapping cognition from behavior: A computerized thought experiment. *Cognitive Science*, 32(3), 504–542.
- Murata, S., Tomioka, S., Nakajo, R., Yamada, T., Arie, H., Ogata, T., & Sugano, S. (2015). Predictive learning with uncertainty estimation for modeling infants' cognitive development with caregivers: A neurorobotics experiment. In *Proceedings of the 2015 Joint IEEE International Conference on Development and Learning and Epigenetic Robotics* (pp. 302–307). Piscataway, NJ: IEEE.
- Ohata, W., & Tani, J. (2020). *Investigation of multimodal and agential interactions in human-robot imitation, based on frameworks of predictive coding and active inference*. arXiv:2002.01632.
- Oliva, D., Philippsen, A., & Nagai, Y. (2019). How development in the Bayesian brain facilitates learning. In *Proceedings of the 2019 Joint IEEE 9th International Conference on Development and Learning and Epigenetic Robotics* (pp. 1–7). Piscataway, NJ: IEEE.
- Oliver, G., Lanillos, P., & Cheng, G. (2019). *Active inference body perception and action for humanoid robots*. arXiv:1906.03022.
- Oudeyer, P.-Y., Kaplan, F., & Hafner, V. V. (2007). Intrinsic motivation systems for autonomous mental development. *IEEE Transactions on Evolutionary Computation*, 11(2), 265–286.
- Park, J.-C., Kim, D.-S., & Nagai, Y. (2018). Learning for goal-directed actions using RNNPB: Developmental change of “what to imitate.” *IEEE Transactions on Cognitive and Developmental Systems*, 10(3), 545–556.
- Park, J.-C., Lim, J. H., Choi, H., & Kim, D.-S. (2012). Predictive coding strategies for developmental neurorobotics. *Frontiers in Psychology*, 3, 134.
- Parr, T., & Friston, K. J. (2017). Working memory, attention, and salience in active inference. *Scientific Reports*, 7(1), 1–21.
- Pezzato, C., Ferrari, R., & Corbato, C. H. (2020). A novel adaptive controller for robot manipulators based on active inference. *IEEE Robotics and Automation Letters*, 5(2), 2973–2980.
- Pezzulo, G., Donnarumma, F., Iodice, P., Maisto, D., & Stoianov, I. (2017). Model-based approaches to active perception and control. *Entropy*, 19(6), 266.
- Pezzulo, G., Rigoli, F., & Friston, K. (2015). Active inference, homeostatic regulation and adaptive behavioural control. *Progress in Neurobiology*, 134, 17–35.
- Philippsen, A., & Nagai, Y. (2019). A predictive coding model of representational drawing in human children and chimpanzees. In *Proceedings of the 2019 Joint IEEE 9th International Conference on Development and Learning and Epigenetic Robotics* (pp. 171–176). Piscataway, NJ: IEEE.
- Pickering, M. J., & Clark, A. (2014). Getting ahead: Forward models and their place in cognitive architecture. *Trends in Cognitive Sciences*, 18(9), 451–456.
- Pio-Lopez, L., Nizard, A., Friston, K., & Pezzulo, G. (2016). Active inference and robot control: A case study. *Journal of the Royal Society Interface*, 13(122), 20160616.
- Ramstead, M. J., Kirchhoff, M. D., & Friston, K. J. (2020). A tale of two densities: Active inference is enactive inference. *Adaptive Behavior*, 28(4), 225–239.



- Rao, R. P., & Ballard, D. H. (1999). Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2(1), 79.
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning representations by back-propagating errors. *Nature*, 323(6088), 533–536.
- Sajid, N., Parr, T., Hope, T. M., Price, C. J., & Friston, K. J. (2020). Degeneracy and redundancy in active inference. *Cerebral Cortex*, 30, 5750–5766.
- Sancaktar, C., & Lanillos, P. (2019). *End-to-end pixel-based deep active inference for body perception and action*. arXiv:2001.05847.
- Schillaci, G., Ciria, A., & Lara, B. (2020). Tracking emotions: Intrinsic motivation grounded on multi-level prediction error dynamics. In *Proceedings of the 10th Joint International Conference on Development and Learning and Epigenetic Robotics*. arXiv:2007.14632.
- Schillaci, G., Hafner, V. V., & Lara, B. (2016). Exploration behaviors, body representations, and simulation processes for the development of cognition in artificial agents. *Frontiers in Robotics and AI*, 3, 39.
- Schillaci, G., Pico Villalpando, A., Hafner, V. V., Hanappe, P., Colliaux, D., & Wintz, T. (2020). Intrinsic motivation and episodic memories for robot exploration of high-dimensional sensory spaces. *Adaptive Behavior*. <https://doi.org/10.1177/1059712320922916>.
- Schillaci, G., Ritter, C.-N., Hafner, V. V., & Lara, B. (2016). Body representations for robot ego-noise modelling and prediction: Towards the development of a sense of agency in artificial agents. *Artificial Life Conference Proceedings*, 28, 390–397.
- Schwartenbeck, P., Passecker, J., Hauser, T. U., FitzGerald, T. H., Kronbichler, M., & Friston, K. J. (2019). Computational mechanisms of curiosity and goal-directed exploration. *eLife*, 8, e41703.
- Seth, A. K., & Tsakiris, M. (2018). Being a beast machine: The somatic basis of selfhood. *Trends in Cognitive Sciences*, 22(11), 969–981.
- Spratling, M. W. (2008). Predictive coding as a model of biased competition in visual attention. *Vision Research*, 48(12), 1391–1408.
- Spratling, M. W. (2017). A review of predictive coding algorithms. *Brain and Cognition*, 112, 92–97.
- Tani, J. (2019). Accounting for the minimal self and the narrative self: Robotics experiments using predictive coding. In *AAAI Spring Symposium: Towards Conscious AI Systems*. Stanford, CA: AAAI.
- Tani, J., & Nolfi, S. (1999). Learning to perceive the world as articulated: An approach for hierarchical learning in sensory-motor systems. *Neural Networks*, 12(7–8), 1131–1141.
- Tschantz, A., Seth, A. K., & Buckley, C. L. (2020). Learning action-oriented models through active inference. *PLOS Computational Biology*, 16(4), e1007805.
- Ueltzhöffer, K. (2018). Deep active inference. *Biological Cybernetics*, 112(6), 547–573.
- Van de Cruys, S. (2017). *Affective value in the predictive mind*. Johannes Gutenberg-Universität Mainz.
- Whittington, J. C., & Bogacz, R. (2019). Theories of error back-propagation in the brain. *Trends in Cognitive Sciences*, 23, 235–250.
- Williams, D. (2018). Predictive processing and the representation wars. *Minds and Machines*, 28(1), 141–172.

- Wolpert, D. M., Ghahramani, Z., & Jordan, M. (1995). An internal model for sensorimotor integration. *Science*, 269(5232), 188–1882.
- Wolpert, D. M., & Kawato, M. (1998). Multiple paired forward and inverse models for motor control. *Neural Netw.*, 11(7–8), 1317–1329.
- Zhong, J., Cangelosi, A., Zhang, X., & Ogata, T. (2018). Afa-prednet: The action modulation within predictive coding. In *Proceedings of the 2018 International Joint Conference on Neural Networks* (pp. 1–8). Piscataway, NJ: IEEE.

---

Received August 12, 2020; accepted December 31, 2020.