

MIT Open Access Articles

World model learning and inference

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Friston, Karl, Moran, Rosalyn J, Nagai, Yukie, Taniguchi, Tadahiho, Gomi, Hiroaki et al. 2021. "World model learning and inference." *Neural Networks*, 144.

As Published: 10.1016/J.NEUNET.2021.09.011


Publisher: Elsevier BV

Persistent URL: <https://hdl.handle.net/1721.1/150396>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of use: Creative Commons Attribution-NonCommercial-NoDerivs License





2021 Special Issue on AI and Brain Science: Perspective

World model learning and inference

Karl Friston^a, Rosalyn J. Moran^b, Yukie Nagai^c, Tadahiro Taniguchi^d, Hiroaki Gomi^{e,*},
Josh Tenenbaum^{f,g}

^a Wellcome Centre for Human Neuroimaging, Institute of Neurology, University College London (UCL), WC1N 3BG, UK

^b Department of Neuroimaging, Institute of Psychiatry, Psychology & Neuroscience, King's College London, SE5 8AF, UK

^c International Research Center for Neurointelligence (IRCN), The University of Tokyo, Tokyo, Japan

^d College of Information Science and Engineering, Ritsumeikan University, Shiga, Japan

^e NTT Communication Science Labs., Nippon Telegraph and Telephone, Kanawaga, Japan

^f Department of Brain and Cognitive Sciences, MIT, Cambridge, MA, USA

^g The Center for Brains, Minds and Machines, MIT, Cambridge, MA, USA

ARTICLE INFO

Article history:

Available online 21 September 2021

Keywords:

Generative model
Probabilistic inference
Predictive coding
Bayesian inference
Free energy principle
Cognitive development

ABSTRACT

Understanding information processing in the brain—and creating general-purpose artificial intelligence—are long-standing aspirations of scientists and engineers worldwide. The distinctive features of human intelligence are high-level cognition and control in various interactions with the world including the self, which are not defined in advance and are vary over time. The challenge of building human-like intelligent machines, as well as progress in brain science and behavioural analyses, robotics, and their associated theoretical formalisations, speaks to the importance of the world-model learning and inference. In this article, after briefly surveying the history and challenges of internal model learning and probabilistic learning, we introduce the free energy principle, which provides a useful framework within which to consider neuronal computation and probabilistic world models. Next, we showcase examples of human behaviour and cognition explained under that principle. We then describe symbol emergence in the context of probabilistic modelling, as a topic at the frontiers of cognitive robotics. Lastly, we review recent progress in creating human-like intelligence by using novel probabilistic programming languages. The striking consensus that emerges from these studies is that probabilistic descriptions of learning and inference are powerful and effective ways to create human-like artificial intelligent machines and to understand intelligence in the context of how humans interact with their world.

© 2021 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

1. Introduction

During their lives humans constantly interact with the physical environment, as well as with themselves and others. How do we generate our actions to interact with various environments? Important brain functions were modelled by computational learning schemes in the 80–90s (Jordan & Rumelhart, 1992; Kawato, 1999; Kawato, Furukawa, & Suzuki, 1987; Wolpert, Ghahramani, & Jordan, 1995); these studies demonstrated that internal models of the body and environment dynamics are effective for achieving movement goals. In addition, stochastic generative modelling (Dayan, Hinton, Neal, & Zemel, 1995) and reinforcement learning (Barto, Sutton, & Anderson, 1983; Barto,

Sutton, & Brouwer, 1981; Sutton & Barto, 1987) were proposed as general models of biological learning that implicitly acquire structures and statistical regularities inherent in sets of data generated by the world. These successful theoretical frameworks support the idea that acquiring internal models is a natural way to realise optimal exchange with the environment, and therefore constitutes a reasonable explanation of how the brain functions. Furthermore, physiological representations of internal models were widely examined (Doya, 1999; Imamizu et al., 2000; Kawato & Gomi, 1992; Miall & Wolpert, 1996; Shidara, Kawano, Gomi, & Kawato, 1993) as ways to understand brain computations for acting on the world.

In the past decades, studies of neural networks and statistical learning have progressed by developing new model learning algorithms (Hinton, Osindero, & Teh, 2006; Kingma & Welling, 2013) and fast and parallel computational engines. These trends have stimulated not only machine learning researchers to realise human-like recognition and action control by creating world models (Eslami et al., 2018; Ha & Schmidhuber, 2018), but have

* Corresponding author.

E-mail addresses: k.friston@ucl.ac.uk (K. Friston), rosalyn.moran@kcl.ac.uk (R.J. Moran), nagai.yukie@mail.u-tokyo.ac.jp (Y. Nagai), taniguchi@em.ci.ritsumei.ac.jp (T. Taniguchi), hiroaki.gomi.ga@hco.ntt.co.jp (H. Gomi), jbt@mit.edu (J. Tenenbaum).

also inspired neuroscience researchers to use these techniques for decoding and creating brain-like processing using natural data from the real world (Yamins et al., 2014).

In this article, Karl Friston (Section 2) introduces the free energy principle as a framework to integrate related theories, with a special emphasis on Bayesian inference and learning. Next, Rosalyn Moran (Section 3) and Yukie Nagai (Section 4) explain how human actions and cognitive behaviour can be explained under that principle. Tadahiro Taniguchi (Section 5) discusses current and novel approaches to create intelligence and world models for real robots. Lastly, Josh Tenenbaum (Section 6) compares human and artificial intelligence in light of the previous sections. He considers the outstanding challenges of building a general-purpose AI by stressing the importance of world modelling and probabilistic inference. We conclude with a brief discussion of what is needed to realise a brain-like artificial intelligence that can interact naturally with the real world and our society.

2. Theoretical examination of the world modelling

This section considers the fundamentals of learning and inference under world models; under the banner of the free energy principle (Friston, 2010). The free energy principle (FEP) and its corollary—active inference (depicted in Fig. 1)—is a normative framework for inference and learning that takes the Bayesian brain into the realm of actions and decisions (Doya, Ishii, Pouget, & Rao, 2007; Knill & Pouget, 2004). The free energy principle inherits from variational Bayes (Beal, 2003; Winn & Bishop, 2005), where the world model is known as a generative model (Dayan et al., 1995; Le Roux, Heess, Shotton, & Winn, 2011; Roweis & Ghahramani, 1999). A generative model is simply a probabilistic description of how causes (i.e., latent states) generate consequences (i.e., data or sensations).

The basic idea behind the FEP is that everything entailed by an artefact, agent or autonomous system optimises variational free energy. This kind of free energy has various names. For example, in machine learning it is known as an evidence lower bound, where evidence is also known as the marginal likelihood (Winn & Bishop, 2005). This means that extremising free energy maximises model evidence, namely, the evidence for a generative model inherent in any observable data or sensations. This means that both action and perception can be neatly summarised as self-evidencing (Hohwy, 2016), under a variational principle of stationary action (Friston, 2019). The action in this instance is a time or path integral of free energy.

Crucially, free energy is a functional (i.e., a function of a function) of two quantities. First, sensory data and a probability distribution over the unobservable states generating those data. This variational density is taken to be encoded, represented, or parameterised by the internal states of any system, ranging from a particle to a person. On this view, perception corresponds to changing internal states to minimise the divergence between the variational density and the posterior density over latent states, given observations. Conversely, action can change the way that data or sensations are sampled—to ensure that they provide the greatest evidence for the generative model entailed by an agent. This dual aspect—to optimising free energy—gracefully accounts for action and perception, where both are in the service of maximising (a variational bound on) marginal likelihood.

This may sound complicated; however, it is just a generalisation of variational procedures that underwrite approximate Bayesian inference to include action. In other words, it subsumes perception and planning into the same Bayesian mechanics by treating planning as inference (Attias, 2003; Botvinick & Toussaint, 2012; Kaplan & Friston, 2018). Although not our focus here, there is a back story to the free energy principle that

comes from the physics of self-organisation; namely, systems that maintain some form of non-equilibrium steady-state (Friston, 2019). In this setting, the action (time or path integral) above corresponds to the system's entropy. This follows because, mathematically, the negative logarithm of model evidence is also known as self-information in information theory—and the average self-information is entropy. This means that maximising model evidence, on average and over time, is just a statement that certain systems resist an increase in the entropy of their constituent states (England, 2015; Friston, 2013; Jeffery, Pollack, & Rovelli, 2019). So how does this help us frame inference and learning?

In general, there are three levels of optimisation under the free energy principle. These correspond to the unknowns (i.e., latent causes) in the generative model. These unknowns comprise (i) latent states generating outcomes, (ii) model parameters encoding contingencies and statistical regularities and, finally (iii) the form or structure of the generative model. Each is equipped with variational density (i.e., a Bayesian belief) that is parameterised by the (i) states, (ii) weights, and (iii) structure of the agent at hand.

2.1. Inference

At the fastest timescale, inference can then be read as optimising the states (e.g., synaptic activity) to optimise variational free energy. This is usually cast in terms of a gradient flow on free energy. Crucially, the gradients of free energy can almost universally be cast as prediction errors. This provides a straightforward and principled way to articulate neuronal dynamics. Furthermore, it leads to particular schemes for free energy optimisation. For example, with generative models of continuous states, we end up with predictive coding schemes (Rao & Ballard, 1999; Srinivasan, Laughlin, & Dubs, 1982) that, in engineering, correspond to Bayesian filters, e.g., extended Kalman and particle filters (Lee & Mumford, 2003; Loeliger, 2002). For generative models of discrete states and time, the equivalent message passing becomes belief propagation or variational message passing (Dauwels, 2007; Winn & Bishop, 2005; Yedidia, Freeman, & Weiss, 2005). All of these schemes have some degree biological plausibility when applied in a neurobiological setting (Friston, Parr et al., 2017).

2.2. Learning

The second set of unknowns are the parameters of the generative model, encoded in slowly changing weights (e.g., synaptic efficacy). Again, learning can be construed as a free energy optimisation process that, in the biological setting, reduces to experience-dependent plasticity via associative or Hebbian schemes (Friston, FitzGerald, Rigoli, Schwartenbeck and Pezzulo, 2017). Because the gradients of free energy can be cast as prediction errors, this also gracefully accommodates back propagation of errors in deep learning and neuronal networks (Amari, 1998; George & Hawkins, 2009; Hinton, 2007; LeCun, Bengio, & Hinton, 2015; Whittington & Bogacz, 2017).

2.3. Model selection

Finally, we have the structure or form of the model, e.g., cortical hierarchies in the brain (Mumford, 1992). The structure of the model can be regarded as being optimised with respect to free energy or model evidence via a process of Bayesian model selection; namely, selecting those models with the greatest marginal likelihood, as assessed over an extended period of time. This level of optimisation manifests at different scales. For example, one can construe natural selection as nature's way of performing




Fig. 1. Bayesian mechanics and active inference. This graphic summarises the belief updating implicit in the minimisation of variational and expected free energy. It provides a generic (active) inference scheme that has been used in a wide variety of applications and simulations; ranging from games in behavioural economics (FitzGerald, Schwartenbeck, Moutoussis, Dolan and Friston, 2015) and reinforcement learning (Schwartenbeck et al., 2015) through to language (Friston, Rosch, Parr, Price and Bowman, 2017) and scene construction (Mirza, Adams, Mathys, & Friston, 2016). In this setup, discrete actions solicit a sensory outcome that informs approximate posterior beliefs about hidden or external states of the world – via minimisation of variational free energy under a set of plausible policies (i.e., *perceptual inference*). The approximate posterior beliefs are then used to evaluate expected free energy and subsequent beliefs about action (i.e., *policy selection*). Note a subtle but important move in this construction: the expected free energy furnishes prior beliefs about policies. This is interesting from several perspectives. For example, it means that agents infer policies and, implicitly, active states. In other words, beliefs about policies – encoded by internal states – are distinct from the active states of the agent’s Markov blanket. In more sophisticated schemes, agents infer hidden states under plausible policies with a generative model based on a Markov decision process. This means the agent predicts how it will behave and then verifies those predictions based on sensory samples. In other words, agents garner evidence for their own behaviour and actively self-evidence. In this setting, variational free energy reflects the surprisal or evidence that a particular policy is being pursued. In sum, this means the agent (will appear to) have elemental beliefs about its enactive self – beliefs that endow it with a sense of purpose, in virtue of the prior preferences that constitute risk. A key insight from simulations is that the form of the generative model can be quite different from the process by which external states generate sensory states. In effect, this enables agents (i.e., particles) to author their own sensorium in a fashion that has close connections with econiche construction (Bruineberg & Rietveld, 2014). Please see Friston, Parr and de Vries (2017) for technical details and Friston, Parr et al. (2017) for a discussion of how the implicit belief updating might be implemented in the brain.

Bayesian model selection—i.e., accumulating evidence about an econiche by selecting phenotypes that have high adaptive fitness or marginal likelihood (Campbell, 2016; Frank, 2012). At a somatic timescale, in biology, this could be regarded as neurodevelopment with (epigenetic) hyperpriors over model structure. In cognitive science, this kind of optimisation process is often referred to as structure learning (Tenenbaum, Kemp, Griffiths, & Goodman, 2011; Tervo, Tenenbaum, & Gershman, 2016). In machine learning, structure learning is closely related to algorithmic learning and, perhaps, meta-reinforcement learning (Ishii, Yoshida, & Yoshimoto, 2002). In statistics, the exploration of different model structures is often cast in terms of nonparametric Bayes (Goldwater, 2007). In all instances, the way in which a new structure or model is scored corresponds to the model evidence.

Generally speaking, the optimisation in terms of inference, learning and model selection go hand-in-hand and contextualise each other. In the variational setting of the free energy principle, this is necessarily so—because optimisation rests upon a factorisation of the variational density over the three different levels of unknowns, which means that each level provides empirical priors for the level below. For example, Bayesian model selection determines which parameters are in play, while learning some model parameters optimises inference about hidden or latent states. With this integrative framework in place, we will now

consider three cardinal issues that will emerge in various guises throughout the remainder of our treatment.

2.4. Inference and precision

Above, we divided optimisation into inference, learning and model selection. However, a finer grained analysis of inference calls for a consideration of the representation of uncertainty. If one subscribes to the free energy principle, then optimisation corresponds to optimising posterior or Bayesian beliefs (or their sufficient statistics). This means that it is not sufficient to use point estimates of various quantities, the precision or inverse dispersion (i.e., negentropy) of these beliefs also has to be optimised. Sometimes this is a more difficult problem that estimating the average or expectation of an unknown (Clark, 2013a; Hohwy, 2013).

In engineering, this would be like optimising the Kalman gain; namely, the precision afforded prediction errors in updating state estimation. In neurobiology, this is often construed in terms of attentional selection; affording more or less precision to various Bayesian beliefs at different levels in a hierarchical world model (Ainley, Apps, Fotopoulou, & Tsakiris, 2016; Aukstulewicz & Friston, 2015; Brown, Adams, Parees, Edwards, & Friston, 2013; Kok, Rahnev, Jehee, Lau, & de Lange, 2012; Limanowski, 2017).

We highlight the importance of precision given its central role in balancing the influence of sensory evidence and prior beliefs during Bayesian belief updating. This has found a particularly powerful explanatory role in many areas of cognitive science and computational psychiatry (Nagai, 2019). For example, a large range of neurological and psychiatric syndromes can be cast in terms of a failure to attenuate sensory precision, and its consequences. This is a particularly important perspective because the neurobiological encoding of precision may lie in the postsynaptic gain or various neuronal populations encoding prediction or prediction errors. In turn, the biological instantiation of precision control may implicate neuromodulatory transmitter systems associated with many neuropsychiatric conditions (e.g., autism, schizophrenia, depression, and Parkinson's disease). Attention as precision should be distinguished from salience (Parr & Friston, 2019), in the sense that salience is a kind of affordance that speaks—not to pure sensory attention and attenuation—but the implications for how to act. This brings us to our second point.

2.5. Active inference

Above, we have considered optimising the generative model to best explain observed sensory inputs or data. However, from the point of view of active inference, under the free energy principle, these optimisation processes are just there to infer what the artefact should do next. This leads to a rather subtle extension of the free energy principle, where there is an additional set of unknowns; namely, the sequences of actions or policies to be pursued. In turn, Bayesian beliefs over policies are optimised with respect to the free energy expected under each action sequence. Actions can then be selected from these particular posterior beliefs in the usual way (Attias, 2003; Baker, Saxe, & Tenenbaum, 2009; Botvinick & Toussaint, 2012; Millidge, 2019).

The interesting twist here is the nature of this expected free energy, conditioned upon a policy or plan. Free energy (i.e., log model evidence) per se, can always be written down as accuracy minus complexity. This decomposition has an interesting interpretation when we consider the expected free energy under the predictive posterior over outcomes, given a particular plan. In this setting, inaccuracy becomes ambiguity and complexity becomes risk. In brief, this means that optimising expected free energy corresponds to minimising the risk of deviating from expected or preferred outcomes while, at the same time, reducing the ambiguity (i.e., conditional uncertainty about outcomes given their causes). Intuitively, this means that there are both pragmatic (extrinsic) and epistemic (intrinsic) imperatives for good plans that resolve the exploration–exploitation dilemma (Cohen, McClure, & Yu, 2007; Friston et al., 2015; Schmidhuber, 2006; Still & Precup, 2012; Sun, Gomez, & Schmidhuber, 2011; Tschantz, Baltieri, Seth, & Buckley, 2020).

When one considers inference as planning from this quintessentially enactive perspective, one moves away from conventional machine learning in two senses. First, we need generative models of the future that entertain counterfactual outcomes under various policies. This means that there is a temporal depth to the generative models that are required to explain the behaviour (Friston, Rosch et al., 2017; Rikhye, Guntupalli, Gothoskar, Lázaro-Gredilla, & George, 2019). Second, we bring of overt action—such as selecting which data to mine or sample—to the table. Additionally, one could regard the deployment of precision as a form of covert action that can be construed in terms of attention (Parr & Friston, 2019). There are many fascinating aspects to this perspective on active inference or perception that we will touch upon in later sections.

One might ask: what is the relationship between free energy minimisation—as an existential imperative—and the formulation

of planning as a minimisation of expected free energy? In one sense, the latter is a natural consequence of the former: heuristically, it has been argued—via a *reductio ad absurdum* argument—that agents who do not act to minimise the surprise expected following an action cannot exist, if existence is a minimisation of surprise. A more formal argument would appeal to a variational principle of least action, in which the trajectories of action minimise a path integral of expected free energy in the future. This begs the question, is this a necessary feature of any self-organising system? To a certain extent, this is an outstanding question; however, there is a direct relationship between the (log) probability of an action and expected free energy, which depends on the degree to which actions minimise the ambiguity of outcomes, given their causes (Friston, Da Costa, Hafner, Hesp, & Parr, 2021). This admits the possibility that certain systems (like people) engage in a precise and unambiguous exchange with their environment—and look as if they are planning deep into the future. Conversely, other simpler systems (like thermostats and viruses) minimise surprise in the short-term without actively minimising ambiguity. In turn, this raises interesting questions about how one might quantify the difference between these kinds of systems, in terms of their information geometry and density dynamics.

2.6. Structure learning and complexity

As noted above, log model evidence can be decomposed into accuracy and complexity. This is an important consideration that provides a formal link with things like algorithmic complexity and universal computation (Hutter, 2005). In brief, the complexity of a generative model corresponds to the Kullback–Leibler divergence between the posterior and prior. In other words, the effective number of parameters or degrees of freedom that are required to accurately account for some data—and its sampling. Optimising free energy, therefore, puts pressure on finding the simplest explanations and models (Schmidhuber, 2010). This is exactly the same idea that underwrites the minimisation of algorithmic complexity in the setting of minimum description or message length schemes (MacKay, 1995; Wallace & Dowe, 1999). Crucially this imperative applies to inference, learning and model selection. This means that the best models—that generalise and have a high cross validation accuracy—are the simplest models.

One might ask how does one simplify a model? In neurobiology, this would be seen as a form of synaptic regression or pruning to remove redundant model parameters (Tononi & Cirelli, 2006). More generally, the removal of model parameters and associations defines the structure of the model. Perhaps one of the most important examples here is at the heart of deep learning and hierarchical generative models in the brain. This is meant in the sense that a hierarchy is defined in terms of which connections or conditional dependencies are absent. In other words, a hierarchy is the simply the best explanation of how certain data or sensory streams are generated.

Another cardinal aspect of model structure—that conforms to the minimisation of complexity—emerges from factorisation. For example, complexity can be reduced enormously if one can identify conditional dependencies; such that only the marginal distributions need to be encoded or parameterised. A clear example of this is the separation of ‘what’ and ‘where’ in the brain into two hierarchical streams (Friston & Buzsáki, 2016; Ungerleider, 1994). This may follow from the fact that knowing where something is does not tell you what it is and vice versa. Clearly, one has to integrate marginal beliefs of this sort, when generating predictions of sensory input, which speaks to nonlinear interactions of various sorts in the brain—or nonlinearities in deep learning schemes (Lin, Tegmark, & Rolnick, 2017). Another key

perspective on simplicity of structure is the notion of functional segregation, modularity and modular neuronal codes (O’Keefe & Recce, 1993; Zeki & Shipp, 1988). On one view, modular or segregated functional specialisation is just a statement of encoding beliefs about the causes of sensations, using a set of well-chosen marginals (Parr, Sajid and Friston, 2020). In this sense, many questions about the neural code—and attending architectures—can be cast as finding the most parsimonious, minimally complex form of generative model that is apt to explain our sensations.

2.7. Summary of theoretical framing of world modelling

On a final note, it will be apparent that the story above only holds if we commit to optimising beliefs or probability distributions, as opposed to point estimators or expectations. This suggests that in terms of artificial intelligence, only schemes that explicitly represent uncertainty and beliefs (such as variational autoencoders) will enjoy the full benefits of being able to engage in active inference, planning and, more possibly, artificial consciousness. Most importantly, these causes must include the agent herself, which speaks to a special kind of active inference that may be necessary for understanding others, the self or, indeed, building conscious artefacts.

3. Theoretical and computational interpretations of human actions

Action, in the context of the Free Energy Principle falls out as a corollary, known as Active Inference, as simply another, albeit perhaps most important way, to minimise long term surprise. In many machine learning approaches to action selection, in the context of ‘world models’ (Ha & Schmidhuber, 2018) typically ‘the world’ or environment in which an agent sits (for example a car racing game or game of doom), the artificial agent designed to mimic human behaviour has two computational goals with distinct algorithmic architectures (Ha & Schmidhuber, 2018; Hafner et al., 2019). The first computational unit aims to learn and represent the statistical structure of the sensory world (e.g., how pixels representing a racetrack can change over time with bends in the road), while the second policy-making unit—or engine—abuts the world model and aims to perform actions that maximise some cost function, typically a reward structure describing preferred outcomes.

Under active inference the normative model of action in a given environment comprises, effectively, only one unit that acts, learns and infers in concert (Friston, FitzGerald, Rigoli, Schwartenbeck, & Pezzulo, 2016). Estimates of the free energy are computable given a generative model over different sorts of time considerations, as a process of comparing predictions from the generative model with the outcomes of actions in the given environment, for example, by checking whether an action produces the expected sensory feedback (Friston, Samothrakis and Montague, 2012). In other words, actions serve the use of information from the real world, whereby models of human behaviour and action only serve to minimise free energy (Adams, Shipp, & Friston, 2013): by both adjusting parameters of the world model, and by choice of a set of actions that it believes will be associated with the lower prediction error.

3.1. Action under hierarchical dynamic and nonlinear world models

Two specific forms of generative models—for the lived environment—are generally considered and imbue different sorts of actions. Early work focused on nonlinear hierarchical dynamical systems models of environments, where actions reacted online, in continuous time, to the immediate sensory data and

represented reflex arcs in central nervous system (Adams et al., 2013). This sort ‘predictive coding’ account of action dynamics has been used to model mainly proprioceptive (Adams, Perrinet, & Friston, 2012) human behaviours including the discharge of motor sequences such as those applied to hand writing (Friston et al., 2012). Here, actions are driven by gradients of a free energy functional over latent states. These hierarchical dynamic ‘scaffold’ models (Friston, Daunizeau, & Kiebel, 2009) comprised dynamic states and causes in the environment, as well as slower evolving states representing the certainty or precision of those state estimates and finally even slower evolving parameters that governed the form of the dynamical state space (Friston, 2008). These models highlighted that action resulted directly from inference and did not operate in a distinct ‘action selection’ computational unit. Specifically, it was shown that the tripartition of states (a mean field partition) could directly impose specific types of actions linked to particular neurobiological substrates. Importantly, in an agent that both optimises its model and changes its action plans, one overarching cost function can reveal unintuitive consequences and explanations for human behaviour. These may be particularly useful in the context of understanding neurological and psychiatric disorders.

In mimicking human motor control, for example, Parkinsonian-like movements emerged as a failure of inference on the precision of sensory signals (Friston, Shiner et al., 2012). Taking the view of Dopamine as a neuromodulator responsible for precision weighting on sensory contexts, and action as the minimisation of proprioceptive prediction errors, a cued, motor sequence task simulation revealed that sensory signals with affordance (could be acted upon) could be tuned to different levels of precision and lead to different types of motor patterns. In this theoretical work (Adams et al., 2012) as a counterpoint to dopamine as a reward prediction error (FitzGerald, Dolan and Friston, 2015; Pessiglione, Seymour, Flandin, Dolan, & Frith, 2006), tonic dopamine levels were simulated as a controller of bottom-up sensory information—enhancing the sensory prediction errors associated with afforded states of the environment by enhancing the precision of expected sequences. Interestingly, to simulate Parkinsonian-like movements and depleted dopamine, low level sensory precision was diminished, resulting in an over-reliance on higher order model dynamics that predicted (incorrect) sequence transitions. Moreover, specific deficits emerged — namely perseveration to an old sequence of motor commands, once a new cued set of actions had been presented (Friston, Shiner et al., 2012). In a study of patients with Parkinson’s disease on and off levodopa medication, these specific deficits were later observed (Galea, Bestmann, Beigi, Jahanshahi, & Rothwell, 2012). In trials where action reprogramming was required, patients off levodopa displayed a reaction time deficit, similar to the perseverative effects observed in silico in their motor patterns. However, in the context of unpredictable sequences, rare events that required action reprogramming did not result in such a deficit, in line with model predictions (Galea et al., 2012).

3.2. Explore or exploit, action under probabilistic Markovian world models

The second type of generative model that has been used to represent the world in decision-making contexts, where actions have a longer temporal horizon, are partially observable Markov decision processes (Mirza et al., 2016). Here, actions are discrete and policy selection (a series of actions into the future) is, again, adjudicated based upon the Free Energy functional such that uncertainty in the environment is resolved through actions while goals are also attained through actions. This ‘future-looking’ Free Energy can be computed as the Expected Free Energy and

comprises epistemic value and extrinsic value as an emergent feature.

In terms of human behaviour in cognitive neuroscience, this process model (the variational updates) has been theorised to represent neural architectures more associated with abstract planning including hippocampal and prefrontal interactions with the basal ganglia (Parr, Rikhye, Halassa and Friston, 2020). The decision-making schema also encapsulates, for free, the explore–exploit trade off (Sales, Friston, Jones, Pickering, & Moran, 2019). In a study of human players of the arcade game ‘doom’ (Cullen, Davey, Friston, & Moran, 2018), the explore/exploit trade-off, implicit under active inference, was exposed. In aiming to match empirical human behaviours to computational agents, two agents were simulated with identical learning capacities but differing policy selection criteria (Cullen et al., 2018). One agent computed the full expected free energy which comprised the expected log model evidence and the Kullback Leibler divergence from the posterior of the state given an outcome to the expected state. Thus, if the states of the game were all known, without having to observe outcomes, this KL term would be zero and no exploration would be required. The model evidence, in turn, was determined by the prior belief that the agent would ‘win the game’. Here winning was constructed in a simple state–space model of the doom environment as the state of being in front of the monster and firing, and self-evidencing comprises actions that attain this goal. It was demonstrated that, like human performers, the full Free Energy agent tends to remain in the game for longer in initial trials. This is because the agents first explored to learn the structure of the world as illustrated in Fig. 2. In the human players, this exploration phase was shorter however, given that only button press response mappings were only amenable to uncertainty. Crucially, the agent that just maximised extrinsic value—the expected log evidence or priors of winning, poorer performance, farther away from human play was observed (Cullen et al., 2018).

Unlike simple move–point–score arcade games, realistic human behaviour typically requires both continuous and discrete states. When considering these ‘worlds’ both generative models can be combined, and have been built (for example in the context of reading (Friston, Rosch et al., 2017) and speech comprehension) to demonstrate action in the context of understanding worlds. Interestingly, formulating these mixed models requires a supposition of hierarchical levels—which goes first—the discrete or continuous domains? Considerations of cortical neuroanatomy and sensory architectures can reveal a straightforward world model through Free Energy stick-a-brick.

3.3. Action and its emergent timing in mixed world models

Mixed-models of continuous and discrete states of the world have been formulated such that sequences of discrete states of the world predict specific dynamic trajectories in continuous time (Friston, Parr et al., 2017). Under this scheme, reading has been simulated as a prototypical human behaviour. The process of reading requires, simply, the action of our eyes to evince proprioceptive outcomes (position of the eyes) as well as exteroceptive outcomes (what is seen) (Mirza, Adams, Mathys, & Friston, 2018). In previous Free Energy accounts, saccadic eye-movements were simulated under hierarchical dynamic models to evince evidence for a particular visual object, e.g. to ‘see’ a face where priors on contrast at particular locations in the object (eyes, nose, etc.) could be supplied (Friston, Adams, Perrinet and Breakspear, 2012). Later a discrete formulation using MDPs was applied to visual search such that categorisation of visual scenes also required eye movements but applied them in sequence to disambiguate ‘words’ that comprised multiple scene components (Mirza et al., 2016).

By combining the MDP and the continuous dynamical model of visual inputs a deep temporal structure emerges such that ‘reading’ involves a high-level policy of semantic categorisation of distinct sentences, where a policy at a level below prescribes eye movements through words, and the actions below words require the final continuous domain eye-movement that recognises characters in the word. The simulation (Friston, Parr et al., 2017) accounts for reading phenomena such as word skipping (Himmelstoss, Schuster, Hutzler, Moran, & Hawelka, 2020; Rayner, Slattery, Drieghe, & Liversedge, 2011), for example once evidence has been accumulated for a subset of sentences, to disambiguate those sentences, only the remaining disambiguating word needs a fixation event. Importantly, the Free Energy of the whole mixed scheme is internally consistent by using the predicted outcomes of the higher discrete level as prior beliefs or hypothesis on the pixel-based continuous representations below. Then, feeding back from the continuous level are a Bayesian Model Comparison of the visual input predicted by each of these prior models – with evidence accumulating for one or another hypothesis. This requires the higher levels to wait for evidence from lower levels – i.e., the highest decision-making level waits until words are selected at the level below, which then waits to test whether the characters are predicted with high posterior probability. Thus, emerges a natural temporal structure (Friston, Rosch, Parr, Price, & Bowman, 2018) where depth represents both high level semantics and higher orders of time. In turn the process scheme returns firing patterns that resemble pre-saccadic delay period activity in the prefrontal cortex, as well as event-related potentials akin to those observed in inferotemporal cortex.

3.4. Summary of theoretical and computational interpretations of human behaviour

In this section we have reviewed action under both discrete and continuous schemes and in mixed models. Overall, the Active Inference framework provides one overarching normative goal of behavioural outputs, while providing for distinct process models—depending on ‘the world’ at hand. This involves, in turn, distinct computational instantiations of Free Energy minimisation in predictive codes and variational message passing. Whether the demonstrations by these artificial agents truly mimic human behavioural control may require both more specific experimental settings and large-scale environments that test at scale.

4. Theoretical and computational interpretations of cognitive development

Development is a long-term process that involves learning of various cognitive functions. Especially in the first few years of life, infants and toddlers undergo significant changes in their perception, action, and social capabilities. An open question is what neural mechanisms drive cognitive development.

Neuroscience and computational studies have suggested that the free energy principle (Friston, 2010; Friston, Kilner, & Harrison, 2006) and predictive coding or processing (Clark, 2013b; Rao & Ballard, 1999) provide a unified account for cognitive development (Nagai, 2019). Infants and toddlers learn to acquire world models (or internal models) to perceive and act on the world through the minimisation of prediction errors. The theory of predictive coding can address two important aspects of cognitive development: temporal continuity and individual diversity (Nagai, 2019). Temporal continuity refers to developmental dependencies between cognitive abilities. For example, social cognition requires non-social sensorimotor abilities to interact with the environment and other persons. Individual diversity, in contrast, refers to qualitative and quantitative differences

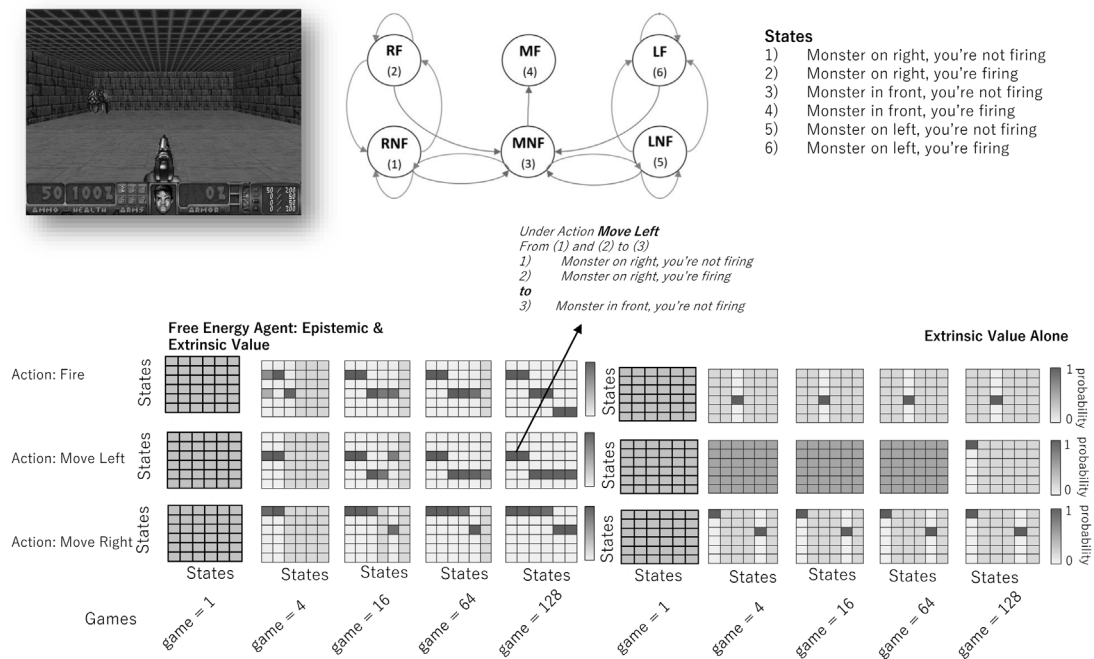


Fig. 2. Learning how actions affect outcomes. State transition matrices shown here corresponding to the fire, move left, move right actions at game trials $t = 4$, $t = 16$, $t = 64$, and $t = 128$ under the free energy–minimisation (left) and extrinsic value only (right) based cost function. Each matrix represents the agent's belief about how the environment will change after making the respective action. The uniformity of the matrix shows initial uncertainty (uniform probability) over state transitions. After 128 epochs of learning, the transition matrices of the free energy agent have converged to those of the optimised agent presented in the MDP. The transition matrix of the value-alone agent is much sparser by comparison, reflecting a lack of knowledge about the environmental contingencies.

in cognitive abilities between individuals. For example, persons with developmental disorders show different characteristics and temporal dynamics in cognitive capabilities compared to those of typically developing individuals. The following sections provide theoretical and computational interpretations of cognitive development based on predictive coding.

4.1. Continuous development based on prediction error minimisation

How do infants learn to communicate with others? What is the origin of social intelligence? Neuroscience studies suggest that internal models for sensorimotor behaviours play an important role in social interaction. Mirror neuron systems (Iacoboni & Dapretto, 2006; Rizzolatti, Fadiga, Gallese, & Fogassi, 1996; Rizzolatti, Fogassi, & Gallese, 2001) are found as a neural basis for social intelligence, which links the abilities of action perception and action production through internal models. The discovery of mirror neuron systems has inspired researchers to model them and to reveal the developmental origin of mirror neuron systems. A computational theory of cognitive development suggests that two ways of minimising prediction errors lead to the emergence of mirror neuron systems and social cognition: updating internal models and generating actions through active inference (Nagai, 2019).

4.1.1. Sensorimotor abilities acquired through internal model learning

In order to develop social abilities, infants first have to learn to control and recognise their body. The process of updating internal models is suggested to enable infants to acquire such basic sensorimotor functions. Humans are born with immature internal models and thus need to learn the models through their sensorimotor experiences. For example, the abilities of self-recognition and self-other discrimination, which are cornerstones of cognitive

development, are considered to develop through the learning of internal models. Self-executed actions are detected as (nearly) perfectly predictable events, whereas actions produced by others are hard to predict. The certainty of predictions by internal models can be utilised to discriminate the self from others. Goal-directed actions such as reaching for and manipulating an object are also acquired by internal models. Exploratory behaviours generated by immature internal models enable infants to experience contingent relationships between multiple sensory modalities. Minimising prediction errors through explorations leads to the acquisition of sensorimotor coordination for goal-directed actions.

4.1.2. Emergence of social cognition based on active inference

Once internal models are trained, they are applied to social contexts. Social cognitive abilities are expected to emerge through the process of generating actions to minimise prediction errors. Internal models learned through infants' sensorimotor experiences are applied to recognise and predict actions generated by other persons. In contrast to infants' own sensorimotor behaviours, certain prediction errors are detected from others' actions because the internal models learned using only infants' sensorimotor experiences cannot perfectly predict others. The mechanism of active inference then triggers infants' actions to minimise errors. These actions alter the state of other persons and the environment, which is regarded as the emergence of social interaction. Note that actions at this stage are not yet socially motivated but are triggered to minimise prediction errors. Additional mechanisms, such as social feedback, are necessary to promote the development from proto-social to social cognition.

4.1.3. Computational interpretations of social cognitive development

The above hypothesis has been examined in several computational experiments. Copete et al. (2016) proposed a deep autoencoder that can acquire the function of mirror neuron systems (see Fig. 3). The network, which was pre-trained through

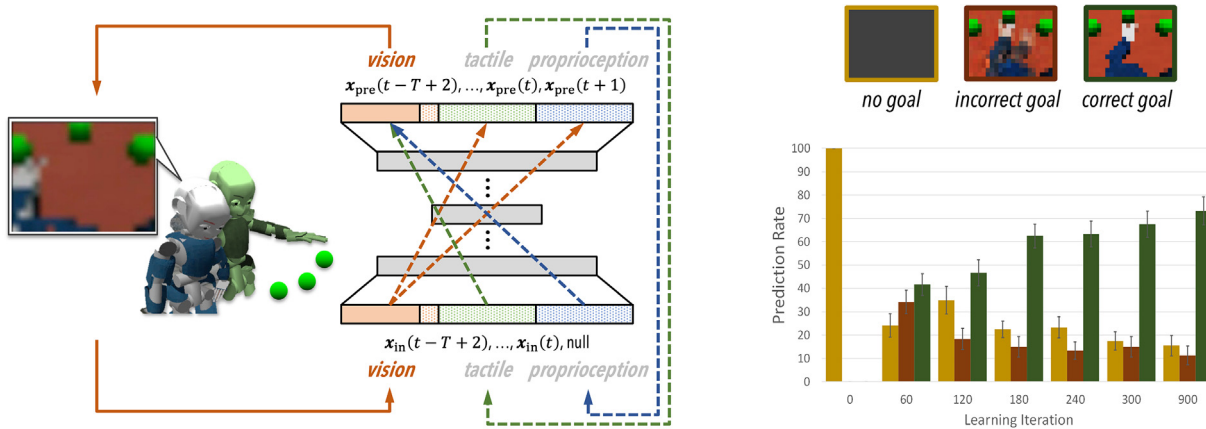


Fig. 3. Development of a mirror neuron system based on predictive coding. Source: Modified from Copete, Nagai, and Asada (2016).

a robot’s motor experiences, was applied to estimate the goal of another robot. Their key idea was that the network can recall the corresponding sensorimotor representations during action observation, where only the visual signal is obtained as input. Their experiment demonstrated that the network trained through action generation outperformed the network trained through action observation. Multimodal representations acquired in the internal model enabled the robot to better estimate the goal of another robot as do mirror neuron systems.

This model was further extended to generate altruistic behaviours. Baraglia et al. (Baraglia, Cakmak, Nagai, Rao, & Asada, 2017; Baraglia, Nagai, & Asada, 2016) developed a robotic system that produced ‘helping actions’ based on active inference. Their robot—on observing a person failing to achieve a task—could accomplish the goal of the observed task as if it was the robot’s own goal. A discrepancy between the observed state (i.e., a failure of the task) and the robot’s predicted state (i.e., a success of the task) motivated the robot to minimise prediction errors by executing a relevant action. Proto-social behaviours emerged without social motivation.

Horii, Nagai, and Asada (2016, 2018) proposed a multimodal deep belief network that enabled a robot to acquire emotional states and to imitate other person’s emotions like infants (see Fig. 4). Their key idea was that the process of minimising prediction errors enables the network to self-organise emotions in the latent space and to produce imitative behaviour using the reconstructed sensory signals. Their experiments demonstrated that the robot developmentally differentiated emotions as observed in infants and that the function of mirror neuron systems acquired in the network improved the estimation of other’s emotions. Similarly, to the previous experiments, perceptual and active inference played an important role in social development.

4.2. Individual diversity caused by aberrant predictive processing

Some children follow different developmental trajectories. What neural mechanisms cause their atypical development? Autism spectrum disorder (ASD) is a type of neurodevelopmental disorder characterised by atypical social communication and interaction and a preference for restricted and repetitive patterns of behaviours and interests (Association, 2013; Baron-Cohen, 1995).

Neuroscience and first-person studies suggest that weakness or difficulties in integrating information might be a core disease of ASD (Happé & Frith, 2006; Kumagaya, 2015). They argue that difficulties in social communication appear as a secondary problem of atypical perception and action. The theory of predictive coding provides further insight into the neural mechanisms of

ASD. Imprecise prior predictions (i.e., hypo-priors) are hypothesised to account for cognitive characteristics of ASD (Brock, 2012; Pellicano & Burr, 2012; Van de Cruys et al., 2014).

4.2.1. Hypo-priors and hyper-priors as potential causes of ASD

It is known that individuals with ASD often suffer from hyperesthesia (Brown, Tollefson, Dunn, Cromwell, & Filion, 2001; O’Neill & Jones, 1997). They show increased sensitivity to sensory stimuli such as vision, audio, and touch. The hypo-priors hypothesis suggests that inaccurate prior predictions make people with ASD highly sensitive to sensory signals (Brock, 2012; Pellicano & Burr, 2012; Van de Cruys et al., 2014). In contrast, some individuals with ASD show hypoesthesia (Brown et al., 2001; O’Neill & Jones, 1997). Their sensations seem to be reduced, and no or weak behavioural responses to stimuli are observed. Notably, even the same person exhibits both hyperesthesia and hypoesthesia depending on sensations and/or contexts.

The predictive coding account for ASD has been extended to cover the diversity within ASD as well as differences between ASD and typical development. Nagai (2019) suggested that hypo-priors and hyper-priors explain seemingly inconsistent or contradicting cognitive characteristics of ASD. On one hand, lower precision in prior predictions (i.e., hypo-priors) makes the brain strongly rely on sensory signals. Hypo-priors indicate that the brain has immature or inaccurate internal models and thus relies strongly on sensory signals. Too strong a reliance on sensory signals results in hyperesthesia and poor generalisation capabilities as observed in ASD. On the other hand, a higher precision of prior predictions (i.e., hyper-priors) is hypothesised to produce different types of ASD. Hyper-priors intimate that the brain is highly confident in predictions. Overly precise prior predictions attenuate the influence of sensory signals and bias active inference towards top-down predictions and intentions. Hypoesthesia and repetitive behaviours might be caused by predictive processing with aberrant precision encoding of this sort.

The above hypothesis provides a new view of the spectrum of cognitive development. The phenomenology of people with ASD are explained by two extremes of aberrant predictive processing (i.e., hypo-priors and hyper-priors). In contrast to the traditional view of the ASD spectrum, typical development is considered to be located in the middle of the spectrum, where people have properly balanced predictive processing. This view implies that even if people with ASD exhibit different difficulties in cognition, the underlying neural mechanisms might share the common characteristics of aberrant predictions. The contrary is true for their similarities in cognitive behaviour caused by distinct neural processing. The theory of predictive coding sheds light on the diagnosis of developmental disorders and their treatment.

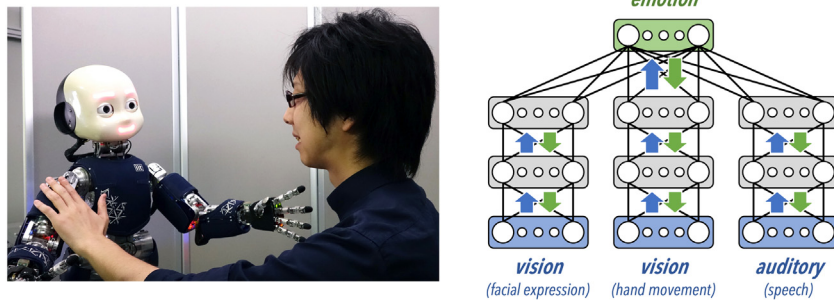


Fig. 4. Development of emotion and its imitation through perceptual and active inference.
Source: Modified from Horii et al. (2016, 2018).

4.2.2. Computational interpretations of developmental individuality

Philippsen and Nagai (2020, in press), investigated how modifications of prior predictions affect the learning and behaviour performance of neural networks (see Fig. 5). They designed a recurrent neural network based on predictive coding with Bayesian inference and trained it to perform the representational drawing of multiple objects. Their experiments, where a parameter of prior precision was controlled, revealed different effects of hyper- and hypo-priors. The networks with hyper-priors often misinterpreted the intended object and stuck to preferred patterns regardless of the input, whereas the networks with hypo-priors scribbled and failed in completing drawings. A close analysis of the internal representations of the networks revealed that learning with hyper- and hypo-priors resulted in undifferentiated stronger attractors and no/weak attractors, respectively. Only the networks with normal priors succeeded in acquiring attractors differentiated for different objects and thus completing the representational drawing of objects.

The above experiment was replicated with human children to verify the hypothesis. Philippsen, Tsuji, and Nagai (2020) conducted the same drawing experiment with children aged two to eight years old. They found that younger children were too immature for representational drawing, whereas older children were better at drawing. The ability to complete missing parts significantly improved with age. Furthermore, younger children showed significant differences between individuals. Some young children performed scribbling and tracing as observed in hypo-prior networks, while other young children repetitively drew preferred patterns similar to hyper-prior networks. This experiment empirically supports the predictive coding hypothesis for individuality.

4.3. Toward open-ended development of world models

This section presented theoretical and computational interpretations of cognitive development. The important hypothesis is that the theory of predictive coding accounts for both temporal continuity and individual diversity of development. Infants and toddlers learn to acquire world models through sensorimotor experiences and to communicate with others using the models. Perceptual and active inference plays a key role in perception, action, and social interaction, where aberrant processes of inference could lead to developmental disorders.

To advance our understanding—and open-ended development of world models in humans and robots—the following issues could be addressed: First, hierarchical predictive processing is necessary for sophisticated cognitive functions. Lower levels to process local and short-term information must be integrated with global and long-term information processed in higher levels. Predictions from higher-levels control activities in lower-levels, while prediction errors from lower levels are used to

update higher levels. Open-ended cognitive development can be achieved by acquiring hierarchical processes of predictive coding. Second, intrinsic motivations that drive continuous development should be integrated. Current robot experiments are still limited to specific tasks, although developmental dependencies between non-social and social cognition were demonstrated. A criterion to minimise prediction errors can be used for designing motivational mechanisms for open-ended learning (Oudeyer, 2007). Third, the question as to why and how the precision of prior predictions differs should also be addressed. Although computational studies demonstrated that modifications in the prior precision account for individual and developmental diversities, it is not fully understood yet what neural mechanisms control prior precisions. Recent neuroscience and psychiatry studies suggest that neurotransmitters such as dopamine and serotonin might be related to the precision control (Friston, Shiner et al., 2012; Sterzer et al., 2018). Further computational studies that examine the roles of neurotransmitters are necessary to better understand the predictive processing and its disorder in cognitive development.

5. World model learning and inference in robotics

5.1. Cognitive systems and world models

Natural Intelligence is not a single-purpose and task-oriented cognitive module developed by an external designer (e.g., a pattern recognition system), but a phenomenon emerging through organising sensorimotor information flow. A brain of an embodied system interacting with its environment (i.e., the world) codes the information to predict the observations. Therefore, world model learning is essential for understanding cognitive systems in the real world. A world model is about a subjective world for an agent; autonomous agents (e.g., a human and an autonomous robot) cannot directly observe their state but should infer its latent state based on its history of sensorimotor information (i.e., observations of its subjective viewpoint). Von Uexküll, who established biosemiotics, carefully distinguished *umwelt* from the objective environment (i.e., the world; Von Uexküll, 1992). Biological systems have cognitive limitations, and they can only observe information obtained through their sensorimotor systems. Therefore, a world model is an *umwelt* model.

From an engineering viewpoint, advancement in the field of artificial intelligence has provided numerous intelligent functions, such as visual recognition, speech recognition, and machine translation systems. However, the functions have been trained separately from an embodied cognitive system living in a physical and social environment. The human brain evolves and develops through sensorimotor interaction with its environment for survival. Developing a cognitive system that can behave autonomously and learn a wide range of cognitive capabilities in the real-world environment is an actual challenge in




Fig. 5. Diverse drawing abilities generated by modifications of precision of prior predictions.
Source: Modified from Philippsen and Nagai (in press).

the field of artificial intelligence. This argument results in the use of robots. Possessing a body as that of a robot allows a computational cognitive system to interact with its environment using its sensorimotor system.

The extent to which the world model learning approach can explain human cognitive development is a challenge. Human cognition involves low-level cognitive capabilities, such as sensory perception and motor action, which are mainly related to physical interactions, and high-level cognitive capabilities, such as planning and language, related to social interactions. However, they are intertwined in many cases. Considering the child language acquisition process, we notice that the process is also performed through interaction with their real-world environment. Moreover, such a learning process inevitably involves a multimodal sensorimotor experience. For example, when a child learns the lexicon of a fruit (e.g., apple), the cognitive process of the lexical acquisition involves sensorimotor interactions, including observing, grasping, and biting the object. From the viewpoint of a probabilistic generative model, learning a world model corresponds to modelling sensorimotor observations by inferring local and global latent variables. A variety of integrative, multimodal machine learning systems—representing concept formation (i.e., representation learning and language acquisition)—has been developed in the field of symbol emergence in robotics (Taniguchi et al., 2016, 2019).

5.2. Symbol emergence in robotics

Symbol emergence in robotics is regarded as a part of cognitive and developmental robotics, and it is a research field in which an integrative cognitive system that enables a robot to learn motor skills, form multimodal representations, and acquire language from real-world experience has been explored. For example, SpCoSLAM is a probabilistic generative model that integrates mapping, localisation, multimodal spatial concept formation, and lexical acquisition (Taniguchi, Hagiwara, Taniguchi, & Inamura, 2017, 2020a). Without a pre-existing dictionary, a robot with SpCoSLAM can discover words in an online manner. The discovered words are always grounded in relation to spatial concepts that are organised through the sensorimotor information that the robot experiences; that is, the symbols are grounded. Fig. 6 shows the world or generative (probabilistic graphical) model that underwrites SpCoSLAM. This shows that SpCoSLAM is trained to predict multimodal sensorimotor information, including speech, visual, distance sensor, and odometer information. All these sensations convey sensorimotor information that a robot can obtain single-handedly. This demonstrates that even lexical acquisition in the real-world environment, which is the primary task in language development, can be modelled from the predictive coding viewpoint. SpCoSLAM enables a robot to acquire spatial concepts and

lexicons and recognise the current status through inference of local and global latent variables in the probabilistic generative model.

A series of studies on symbol emergence in robotics has used probabilistic generative models to represent the dynamics of an internal representation of a cognitive system (Taniguchi et al., 2019, 2016). The dynamics have often been referred to as multimodal concept formation. Nakamura et al. proposed a multimodal latent Dirichlet allocation (MLDA) of a machine learning model for representing object concept formation (Nakamura, Araki, Nagai, & Iwahashi, 2011). MLDA and its variants integrate haptic, visual, auditory, and linguistic information to form object categories. Similar phenomena can be reproduced using multimodal variational autoencoder (MVAE) and multimodal bidirectional transformers, which are deep generative models and self-supervised learning architectures, respectively (Miyazawa, Aoki, Horii, & Nagai, 2020; Suzuki, Nakayama, & Matsuo, 2016). Latent variables corresponding to internal representations of multimodal sensorimotor observations of a real-world object or event can be inferred by maximising the marginal likelihood or by maximising the evidence lower bound, that is, by increasing the predictability of its subjective world. This concept is in line with the idea of the free energy principle and the world model.

5.3. Probabilistic generative models for robots

Learning a world model as a probabilistic generative model and inferring its latent variables are not limited to representation learning. Probabilistic generative models are a general framework, not only as descriptive models of human and animal cognitive processes, but also as engineering models for robotics. The idea of “control as probabilistic inference” showed that many reinforcement algorithms on Markov decision process (MDP) are equivalent to probabilistic inference problems on an extended MDP (Levine, 2018). The partially observable MDP (POMDP) is a general probabilistic generative model for representing an autonomous agent. A world model is mostly regarded as a forward and observation model of POMDP. Reinforcement learning on POMDP has been regarded as a Bayesian inference on POMDP (Okada, Kosaka, & Taniguchi, 2020). As shown in Fig. 6, SpCoSLAM is a type of structured POMDP. Therefore, following the concept of control as probabilistic inference, path planning algorithms based on spatial concepts that are formed in a bottom-up manner can be derived. The probabilistic path planning method is called SpCoNavi (Taniguchi, Hagiwara, Taniguchi, & Inamura, 2020b). Accordingly, SpCoSLAM is also considered a computational model for learning world models and their inference.

If a robot acquires a rich world model based on its multimodal sensorimotor experience, it performs many tasks using

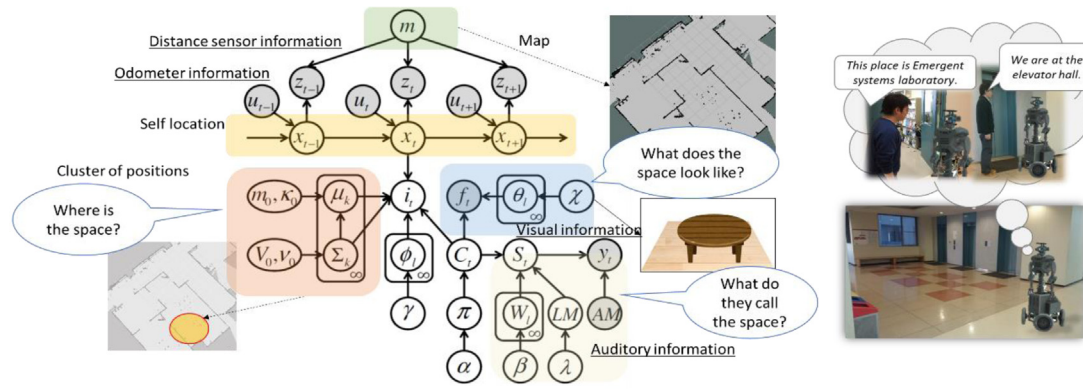


Fig. 6. Probabilistic generative model of SpCoSLAM and overview of a human–robot interaction experiment (Taniguchi et al., 2017).

the world model and its inference. For example, a robot can infer the haptic sensation of a target object through cross-modal interference using visual information with MLDA. A robot can infer the position represented by a word discovered from a user utterance (e.g., “copy room”) and plan the path to the position through inference on SpCoSLAM. Intelligence is not a single purpose and task-oriented function. Thus, multiple functions should be identified as a variety of inferences in a world model.

The perspective of the world model is not limited to physical action learning. However, most studies related to world models in artificial intelligence are physical sensorimotor learning. However, studies on symbol emergence in robotics have proposed probabilistic generative models, which are trained to increase the predictability of the world and can also achieve language acquisition. Therefore, the world model and its inference have quite a broad range in artificial intelligence and robotics.

5.4. Cognitive architecture for real-world robots

Most studies on world models are still limited to abstract models in a simulation environment and simple physical tasks (e.g., driving a car in a simulation environment and playing a simple game). The extent to which the idea of world model can explain human and animal intelligence should be tested in a real-world environment. Human and animal intelligence has evolved to adapt to the real-world environment. This implies that artificial intelligence should be evaluated to clarify if the system can behave appropriately in a physically and socially real environment. In addition, a world model should be able to achieve a variety of tasks, including language acquisition and understanding.

If we intend to test if an artificial intelligence system is an appropriate model of human and animal intelligence, which has evolved to adapt to the real-world environment, we should place it in the real world. This indicates that the validity of the world model should be tested on an artificial embodied system, that is, a robot. However, the number of studies of world models using real robots remains limited. Thus, more robotic studies on world models should be conducted.

Note that creating such robots is insufficient to test the accompanying model as a suitable human or animal cognition model. If we have several models to explain the target behaviours, we should choose the most apt model by considering other sources of evidence, e.g., biological and physiological constraints. We can divide the constructive approach to cognition exhibiting target behaviours into two parts. First, we need to create models that can exhibit target behaviours in the real-world environment, i.e., making reasonable hypotheses from the viewpoint of behaviours. Second, we need to validate these models considering more detailed evidence with appropriate biological constraints. In this section, we focused on the first challenge.

The fusion of studies on cognitive architecture and the world model is also expected. The events and tasks of a robot in the daily environment should deal with diversity and uncertainty. In many theoretical studies of world models, a relatively simple and general graphical model (e.g., POMDP) have been considered. However, if we intend to develop a robot that can perform a variety of manipulation and navigation tasks, recognise objects and events, and acquire and understand language, such a simple model cannot make a robot work in a practical environment. Similar to the human and animal brains with some structure, an explicit design of the latent structure of the cognitive system is required. Cognitive architecture refers to the entire cognitive system that operates a robot, including visual and auditory, decision making, and even emotion (Gonzalez-Billandon, Sciutti, Sandini, & Rea, 2020; Vernon, Metta, & Sandini, 2007). Structured multimodal probabilistic generative models involving many functions (e.g., SpCoSLAM) can be regarded as a type of cognitive architecture; SpCoSLAM is considered as a world model as well. Conversely, most studies on cognitive architecture in cognitive and developmental robotics have limited online learning capability, although they modelled an integrative cognitive system. Furthermore, studies on world models have not captured a variety of cognitive functions that are sufficient for making a robot work in a real-world environment. Fusing the two ideas and creating an integrative cognitive architecture is a future challenge.

To develop an integrative cognitive architecture based on the idea of the world model, deep generative models are a promising approach. However, suppose that we attempt to develop an integrative cognitive architecture; the deep generative model’s size becomes large as it involves speech and visual recognition, planning, manipulation, reasoning, language understanding, among others, in the real-world environment. This development incurs high engineering costs. Therefore, a framework for the distributed development of integrative cognitive architecture is necessary. SERKET is a framework that enables us to compose and decompose probabilistic models to develop a cognitive architecture efficiently (Nakamura, Nagai, & Taniguchi, 2018; Taniguchi et al., 2020).

5.5. Towards world models for robotics

Implementing world model learning and inference in a robot and enabling it to perform lifelong learning in our daily environment is a critical challenge. In contrast to a simulation environment, the real daily environment is full of unexpected uncertainty and complexity. Through the process of implementation and application, we will face new essential problems. This will be crucial for examining the potential of a world model learning-based approach. This section provides three challenges.

The application to service robotics is a challenge. Service robots (e.g., robots in a shop and a domestic environment) must perform physical tasks and social tasks (e.g., communication with a customer). The concept of a world model-based approach is tested in a practical manner.

Multimodal language learning and understanding is another challenge. In the real-world environment, the meaning of language is interpreted in relation to multimodal sensorimotor information. Thus, physical experience is crucial in language acquisition by infants as well. However, most natural language processing studies ignore such real-world information. Language learning and understanding should be studied from a robot's subjective viewpoint. We believe that language learning and understanding are also within the target domain of world model learning.

Studying the world model learning of soft robots is also essential. Unlike the objective world, the subjective world model involves the agent's body as a part of the forward dynamics. In general, humans and animals take advantage of the physical and dynamic characteristics of soft bodies. However, a soft body is considered a source of uncertainty. Hence, world model learning in soft robotics will be a challenge.

6. Building machines that learn, see, and think like people

6.1. Human intelligence and general-purpose AI

Many scientific and engineering researchers have been involved in a joint enterprise between the science and the engineering of intelligence to understand how intelligence arises in the human mind and brain and to create more human-like forms of machine intelligence. One of the best ways to understand how the brain works and to build more human-like forms of artificial intelligence is to build models of intelligence by a kind of reverse engineering. Even though the current AI technologies, alternative ways to realise the human-like intelligence, perform incredibly well in some situations, like in the case of AlphaGo (Silver et al., 2016) beating the world's best players at the world's oldest game, they are still far from the flexible general-purpose intelligences that can use to learn many tasks.

Recent machine-learning technologies, such as deep learning and reinforcement learning, can be used for pattern recognition in the real world. However, human intelligence is not just about finding patterns (Lake, Salakhutdinov, & Tenenbaum, 2015; Lake, Ullman, Tenenbaum, & Gershman, 2017). Intelligence is about all the ways that we build models of the world through our perceptual systems for learning and reasoning. For example, by glancing at a photo of a basketball game, you can easily understand the scene even though it was likely taken from a viewpoint that is new to you. Scene understanding would be quite robust and could be generalised even if the image quality is quite low, such as old video game characters. On the other hand, a system which has only been trained on a courtside view would find it difficult to generalise its scene understanding. The development of a generalisation capability is a grand challenge that we should all try to pursue in this research field. How can we build systems that understand scenes as well as goals, and that can generalise their activities across different tasks without retraining?

6.2. Children's learning of physical world

To tackle the problem outlined above, Tenenbaum's group (Tenenbaum et al., 2011; Ullman & Tenenbaum, 2020) has been trying to reverse engineer the basic kinds of common-sense present even in young children. For instance, when children play with blocks or stacking up caps, they are using intuitive physics,

a basic understanding of objects. In addition, intuitive psychology was demonstrated in a famous experiment (Warneken & Tomasello, 2009), in which infants (14~18 months of age) helped others to attain goals, such as opening cabinets without any instruction; the results suggested natural altruism in infants. To build machines that have the kinds of common senses present even in infants, we should study complex multi-object scenes because general-purpose intuitive physics is one of the earliest building blocks of the kind of common sense for understanding of the world.

How can we develop learning programs for the earliest building blocks? We would not have to build them from scratch because the brain does not start from scratch, thanks to evolution. That is, biological learning over one's lifetime starts with a lot of built-in evolutionarily acquired structure. Researchers in developmental psychology (Saxe, Tenenbaum, & Carey, 2005; Spelke & Kinzler, 2007) have shown that much of the high-level architecture of cognition seems to be in place, structurally and even functionally in babies who are only a few months old. Their work has inspired us to create learning algorithms that start with having a lot of understanding about the physical world. Children from a very young age may be able to understand concepts such as object space, object persistence, and object-contact relationships, but they would acquire a deep understanding of gravity, weight, and force dynamics only through experience at later stages of the developmental process. Therefore, apart from the problem how the brain finds an appropriate structure and connections (Lake et al., 2017), the abovementioned studies imply a possibility to explain how the brain builds an intuitive physics engine in the first several months of life by starting with a fully rich proto-physics engine and by creating a program that modifies the program according to experience.

6.3. Probabilistic programming language (PPL)

Probabilistic programs have been studied since the early 2000s. They are a set of tools, mathematical formalisms, programming languages, implemented systems, and platforms that allow us to combine the best ideas from multiple areas of artificial intelligence for both reverse engineering and engineering human-like forms of intelligence. The probabilistic program concept includes the neural networks for pattern recognition and function approximation, and the original good idea of the field of AI symbolic languages for representing and reasoning with abstract knowledge, as well as probabilistic models and Bayesian inference for reasoning about unobserved causes from sparse uncertain data. There have been several generations of probabilistic programming languages; here, we are especially interested in what we call modern PPL, which are built on top of modern platforms such as PyTorch and TensorFlow, for deep learning, and Gen that gives you all the tools that you needed for symbolic abstract knowledge representation, reasoning, and probabilistic inference.

Another technological idea is the 'game engine in the head', i.e., the use of very fast approximate simulators for graphics physics and planning, which can simulate complex physical situations in realistic ways and use those as a prototype for the model in the agent's head. It is something like an approximation to the common-sense systems of understanding the world, that evolution has built into our brains, and that even young babies can use to explore the world. To model a range of different cognitive capacities in intuitive physics and intuitive psychology, Tenenbaum's group developed an intuitive physics engine almost ten years ago. They have used these models to capture a range of intuitions in a block world setting (Battaglia, Hamrick, & Tenenbaum, 2013; Ullman & Tenenbaum, 2020), and showed, for instance, that they can capture reasonably well people's sense of

stability of a block tower. This idea has been embodied in a rapid trial-and-error learning algorithm that enables us to understand the cognitive foundations of how people come to use tools and understand tools (Allen, Smith, & Tenenbaum, 2020). With this probabilistic algorithm in which the prior is updated using a simple policy gradient algorithm, the number of trials for learning is pretty close to a human's.

6.4. Intuitive physics engine in the brain and its neural network models

Several fMRI studies (Fischer, Mikhael, Tenenbaum, & Kanwisher, 2016; Schwettmann, Tenenbaum, & Kanwisher, 2019) have been conducted to determine where these computations occur in the brain. They revealed the involvement of the premotor and parietal regions in physical reasoning and for representing basic physical quantities. The studies suggest that these areas are the basis of the brain's physics engine. However, these areas have been studied before in different contexts. In particular, they overlap substantially with regions that have been thought to be action planning and tool-use networks, which is reasonable from an evolutionary perspective. Our brain circuits, originally developed to understand how to move our bodies through space, and then to move objects, can generalise their abilities to understanding physical dynamics, even when we are not moving or not moving anything. Recently, there have been several attempts to develop neural network models of intuitive physics. Mrowca et al. (2018) developed a hierarchical graph neural network that learns the representations of the physics of different materials. Bear and colleagues (Bear et al., 2020) have tried to couple a sort of hierarchical physical scene graph to a visual network. Their work combines a convolution neural network with special networks that dynamically translate between visual features with a flexible hierarchical object based on a three-dimensional description of the physical scene.

An example of neural networks that approximate inference in a probabilistic program is the efficient inverse graphics in biological face processing (Tian, Ellis, Kryven, & Tenenbaum, 2020). Since a deep neural network was trained to invert a 3D generative graphics program that generates faces, efficient inverse graphics network presented a good account of the representational similarity space at each of the last three stages of infratemporal cortex processing in the primate brain for faces (Freiwald & Tsao, 2010), and was much better than those generated by a traditional computer vision model, such as the VGG face model.

6.5. Program-learning program

In addition to the generative models in the form of probabilistic programs, we also need to have some learning algorithm that is a program-learning program. Development of such algorithm is a hard problem, because unlike the case of learning in a traditional neural network, there is no nice smooth wait space that we can just follow down to the local minimum with stochastic gradient descent. The space of programs is much rougher. One attempt at understanding how you can learn programs is a Bayesian program learning (Lake et al., 2015) in which a learning program gives a generative model for writing and drawing with strokes and then uses Bayesian inference to invert this relationship for inferring a drawing program. This program in this case could draw a new character after learning one instance and was able to generalise its ability to other characters.

To get this idea to work for something like an intuitive physics engine, either to build one or to understand how you could learn one and modify one from experience, you need learning outcomes that are like program-writing programs. This idea is

described in a recent opinion 'the child as hacker' by Rule and colleagues (Rule, Tenenbaum, & Piantadosi, 2020). The hacking activities include more than tuning parameters of existing code that is like what researchers performing gradient descent in the neural network; for instance, they include writing new functions, writing new libraries, refactoring your code or even writing whole new languages. All of these hacking activities have analogies with children's learning. This is how we build our models of the world.

This concept has been recently embedded in a model of Bayesian program learning, named DreamCoder (Ellis et al., 2020) shown in Fig. 7, by which knowledge is interpreted and generalised. This method was inspired by the wake-sleep algorithm (Dayan et al., 1995; Hinton, Dayan, Frey, & Neal, 1995). The DreamCoder has two different sleep phases, one that grows its library to discover best compressed programs found during waking (named Abstraction in Fig. 7) and another which trains a neural recognition model to find patterns in data that help to guide program search in wake phase (named Dreaming in Fig. 7). Therefore, it is a kind of neuro symbolic learning to program algorithms. For example, when you give it towers, it has to learn how to build or how to draw. Starting from just saying a basic logo like drawing language, it can learn new programming primitives for drawing primitives (See Ellis et al., 2020 for details). Tian et al. (2020) developed a naturalistic drawing task to study how humans rapidly acquire structured prior knowledge, and developed a model that can account for actual human behaviour in which humans are given different kinds of training tasks of learning languages for drawing. The model starts off with a very simple drawing language and it learns to enrich its expressibility. This idea of learning new programming libraries actually seems to describe how humans learn to draw.

In summary, we described recent attempts trying to build machines that learn, see, and think like human beings. By extending the presented approaches and toolkits, we may be able to fulfil AI researchers' oldest dream of building a machine that grows intelligence in the way a human being does, as well as to better understand how our minds are built.

7. General conclusion

Understanding information processing in the brain and creating a general-purpose artificial intelligence are long-standing dreams of many researchers—and are two complementary sides of the same coin. One could argue that a world model is a key construct to progress synergetic studies of natural and artificial intelligence, i.e., the science and engineering of cognition, undertaken collaboratively and harmoniously. The ideas of world model learning and inference provide a formal (i.e., mathematical) calculus for describing cognitive dynamics. In Section 6, we described the challenges to building machines that learn, see, and think like people. Despite the impressive progresses in artificial intelligence, current systems are still limited to a narrow range of cognitive capabilities. In contrast, human cognitive systems can build rich models of the world through our sentience, learning and reasoning.

So—from the perspective of neuroscience and human intelligence—what are the outstanding challenges for artificial intelligence? There are clearly many avenues one could pursue. The particular focus of things like predictive processing and radical constructivism (Glaserfeld, 2002) highlight three key areas. The first is confronting the problem of epistemics in world modelling; namely, building objective functions for inference, learning and action that properly (optimally) balance the imperatives to reach goals, while—at the same time—resolving uncertainty about the context in which those goals are attained. From a normative (optimality) perspective, this rests on combining the

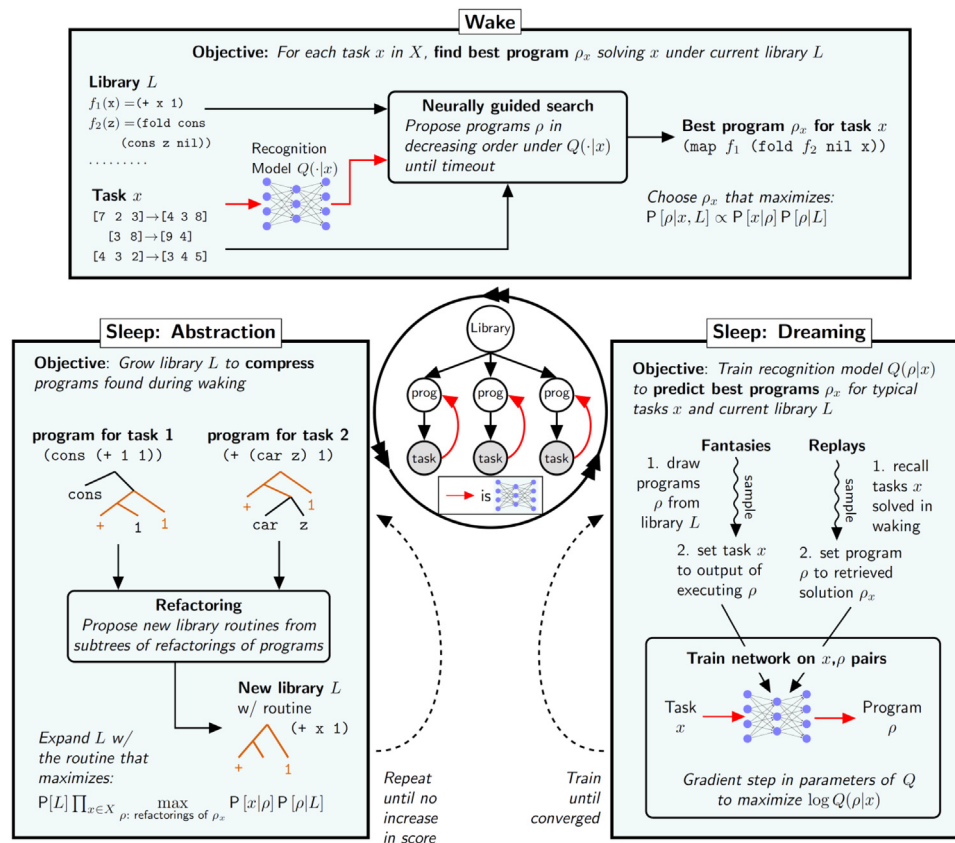


Fig. 7. The architecture of DreamCoder, which serves to perform approximate Bayesian inference for the graphical model shown in the middle. The system observes programming tasks while jointly inferring a latent library capturing cross-program regularities. (Ellis et al., 2020).

principles of optimal Bayesian design (Lindley, 1956) and decision theory (Berger, 2011). In other words, scoring the plausibility of policies in terms of their ability to generate the right kind of outcomes that afford the greatest information gain, under constraints afforded by various loss or reward functions. In the (active inference) neurosciences, this is achieved by absorbing goals into prior preferences and to create objective functions that have both pragmatic and epistemic affordance (Bruineberg & Rietveld, 2014; Cisek, 2007; Gibson, 1977). In machine learning, a similar direction of travel is emerging in the form of generalised KL divergence and free energy functionals; e.g., Hafner et al. (2020).

Another key outstanding challenge is the structure learning problem (TerVO et al., 2016): namely, optimising the structure and form of world models through active engagement with the sensorium. As mentioned in Section 6, the brain does not start from scratch; rather, the seeds of intelligence and the architecture of cognition are inherently embedded in an infant’s brain network. In order to overcome the limitations of current models, we need to further elucidate the learning algorithms, models, and functions that humans have acquired during the evolutionary process and that even infants are born with. Specifically, in order to precisely describe variations in human behaviour and to examine the fitness of the underlying world models, it is necessary to examine their scalability and versatility by integrating hierarchical structures in predictive processing, intrinsic motivations driving continuous development, and factors and parameters altering prior distributions in the brain, etc. In artificial intelligence research, one can imagine two approaches here. The first starts with an overcomplete structure that is weakly parameterised and then prunes redundant model parameters to minimise model complexity, without sacrificing accuracy or predictive performance, e.g., Smith, Schwartenbeck, Parr, and Friston

(2019). The alternative would be to consider growing models through its some principled exploration of model space—of the sort entertained by nonparametric Bayes. And the capacity to model a potentially infinite number of world states (Gershman & Blei, 2012; Goldwater, 2006; Teh, Jordan, Beal, & Blei, 2006). At present, this is a problem confronted by both developmental robotics and computational linguistics. From a neurobiological perspective, the hierarchical co-evolution of world models at evolutionary and developmental timescales may hold important clues here; especially in relation to evolutionary psychology and the important role of culture (Dennett, 2017; Hauser, Chomsky, & Fitch, 2002; Heyes, 2018). It is entertaining to think about self-assembly in robotics, where robots build robots, with ever increasing finesse.

This brings us to a final challenge nicely illustrated by the treatment of self modelling and autism on the one hand, and the special role of language on the other. This challenge speaks to artificial intelligence with minimal selfhood that may be necessary for linguistic (and non-linguistic) communication. From the perspective of world or generative models, this suggests that the generative model should include a hypothesis that “I am an agent”. As exemplified in Sections 3 and 4, various classes of human actions and cognitive development have been interpreted within the theoretical foundations of world model learning based on generative models, called FEP. The probabilistic generative model is also an essential factor of world-model in cognitive robotics, as described in Section 5. Importantly, human agents or advanced robot agents—that evince rich interactions with environments—should be included in any framework, as noted in Sections 3 and 5, because the manner of interaction may define the limits and nature of intelligence. So, what licences the complexity of world models that include a model of selfhood

(Ainley et al., 2016; Fotopoulou & Tsakiris, 2017; Limanowski & Blankenburg, 2013; Seth, 2013). One obvious answer here is to disambiguate between self and other—to provide the necessary context for communication. In turn, this speaks to the potential importance of encultured artificial intelligence, with a special emphasis on the dyadic interactions and the learning of requisite world models.

As we discussed throughout this paper, we again stress that world model learning and inference are crucial concepts in brain and cognitive science, as well as in AI and robotics. World model learning is a fundamental mechanism of human and artificial cognitive systems and contributes to a wide range of cognitive capabilities, e.g., pattern recognition, action selection, social cognition, language learning, and reasoning. In contrast to the distinct development of functional intelligent modules—e.g., visual and speech recognition, machine translation systems, which are trained independently for each functional module—a human brain develops as a whole, while interacting with the body and surrounding environments, i.e., the world. This suggests that we need to explore computational models for world model learning and inference to build both a human-like intelligence and to understand the human brain. By developing models and algorithms and by testing through biological, computational, and robotic experiments, we aspire to a better understanding of the two sides of the same coin; namely, intelligence.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This article was composed based on the talks and discussions at the International Symposium on Artificial Intelligence and Brain Science 2020 held online, supported by KAKENHI Project on Artificial Intelligence and Brain Science. The authors express special thanks to Kenji Doya (symposium organiser) and Mitsuo Kawato (designated discussant). H.G thanks to J De Havas for editing.

Funding

This work was supported by Grants-in-Aid for Scientific Research on Innovation Areas from Japan Society for the Promotion of Science (16H06569 to TT) and (16H06566 to HG), by JST CREST (JPMJCR16E2 to YN).

References

Adams, R. A., Perrinet, L. U., & Friston, K. (2012). Smooth pursuit and visual occlusion: active inference and oculomotor control in schizophrenia. *PLoS One*, 7, Article e47502.

Adams, R. A., Shipp, S., & Friston, K. J. (2013). Predictions not commands: active inference in the motor system. *Brain Structure and Function*, 218, 611–643.

Ainley, V., Apps, M. A. J., Fotopoulou, A., & Tsakiris, M. (2016). 'Bodily precision': a predictive coding account of individual differences in interoceptive accuracy. *Philosophical Transactions of the Royal Society, Series B (Biological Sciences)*, 371, Article 20160003.

Allen, K. R., Smith, K. A., & Tenenbaum, J. B. (2020). Rapid trial-and-error learning with simulation supports flexible tool use and physical reasoning. *Proceedings of the National Academy of Sciences*, 117, 29302–29310.

Amari, S.-i. (1998). Natural gradient works efficiently in learning. *Neural Computation*, 10, 251–276.

Association, A. P. (2013). *Diagnostic and statistical manual of mental disorders (dsm-5®)*. American Psychiatric Pub.

Attias, H. (2003). Planning by probabilistic inference. In *Proc. of the 9th int. workshop on artificial intelligence and statistics*.

Auksztulewicz, R., & Friston, K. (2015). Attentional enhancement of auditory mismatch responses: a DCM/MEG study. *Cerebral Cortex*, 25, 4273–4283.

Baker, C. L., Saxe, R., & Tenenbaum, J. B. (2009). Action understanding as inverse planning. *Cognition*, 113, 329–349.

Baraglia, J., Cakmak, M., Nagai, Y., Rao, R. P. N., & Asada, M. (2017). Efficient human–robot collaboration: When should a robot take initiative? *International Journal of Robotics Research*, 36, 563–579.

Baraglia, J., Nagai, Y., & Asada, M. (2016). Emergence of altruistic behavior through the minimization of prediction error. *IEEE Transactions on Cognitive and Developmental Systems*, 8, 141–151.

Baron-Cohen, S. (1995). *Mindblindness: An essay on autism and theory of mind*. MIT Press.

Barto, A. G., Sutton, R. S., & Anderson, C. W. (1983). Neuronlike adaptive elements that can solve difficult learning control problems. *IEEE Transactions on Systems, Man, and Cybernetics*, 834–846.

Barto, A. G., Sutton, R. S., & Brouwer, P. S. (1981). Associative search network: A reinforcement learning associative memory. *Biological Cybernetics*, 40, 201–211.

Battaglia, P. W., Hamrick, J. B., & Tenenbaum, J. B. (2013). Simulation as an engine of physical scene understanding. *Proceedings of the National Academy of Sciences*, 110, 18327–18332.

Beal, M. J. (2003). *Variational algorithms for approximate bayesian inference* (Ph.D. thesis), University College London, Unpublished.

Bear, D. M., Fan, C., Mrowca, D., Li, Y., Alter, S., Nayeibi, A., et al. (2020). Learning physical graph representations from visual scenes. arXiv:2006.12373 [cs].

Berger, J. O. (2011). *Statistical decision theory and bayesian analysis*. New York, London: Springer.

Botvinick, M., & Toussaint, M. (2012). Planning as inference. *Trends in Cognitive Sciences*, 16, 485–488.

Brock, J. (2012). Alternative Bayesian accounts of autistic perception: comment on Pellicano and Burr. *Trends in Cognitive Sciences*, 16(573).

Brown, H., Adams, R. A., Parees, I., Edwards, M., & Friston, K. (2013). Active inference, sensory attenuation and illusions. *Cognitive Processing*, 14, 411–427.

Brown, C., Tollefson, N., Dunn, W., Cromwell, R., & Fillion, D. (2001). The adult sensory profile: Measuring patterns of sensory processing. *American Journal of Occupational Therapy*, 55, 75–82.

Bruineberg, J., & Rietveld, E. (2014). Self-organization, free energy minimization, and optimal grip on a field of affordances. *Frontiers in Human Neuroscience*, 8(599).

Campbell, J. O. (2016). Universal darwinism as a process of Bayesian inference. *Frontiers in Systems Neuroscience*, 10.

Cisek, P. (2007). Cortical mechanisms of action selection: the affordance competition hypothesis. *Philosophical Transactions of the Royal Society, Series B (Biological Sciences)*, 362, 1585–1599.

Clark, A. (2013a). The many faces of precision (replies to commentaries on whatever next? Neural prediction, situated agents, and the future of cognitive science). *Frontiers in Psychology*, 4.

Clark, A. (2013b). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36, 181–204.

Cohen, J. D., McClure, S. M., & Yu, A. J. (2007). Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society, Series B (Biological Sciences)*, 362, 933–942.

Copete, J. L., Nagai, Y., & Asada, M. (2016). Motor development facilitates the prediction of others' actions through sensorimotor predictive learning. In *2016 joint ieee international conference on development and learning and epigenetic robotics (icdl-epirob)* (pp. 223–229). IEEE.

Cullen, M., Davey, B., Friston, K. J., & Moran, R. J. (2018). Active inference in OpenAI gym: a paradigm for computational investigations into psychiatric illness. *Biological Psychiatry: Cognitive Neuroscience and Neuroimaging*, 3, 809–818.

Dauwels, J. (2007). On variational message passing on factor graphs. In *2007 ieee international symposium on information theory* (pp. 2546–2550).

Dayan, P., Hinton, G. E., Neal, R. M., & Zemel, R. S. (1995). The Helmholtz machine. *Neural Computation*, 7, 889–904.

Dennett, D. C. (2017). *From bacteria to bach and back : the evolution of minds*. Penguin.

Doya, K. (1999). What are the computations of the cerebellum, the basal ganglia and the cerebral cortex? *Neural Networks*, 12, 961–974.

Doya, K., Ishii, S., Pouget, A., & Rao, R. P. N. (2007). *Bayesian brain: probabilistic approaches to neural coding mass*. MIT Press.

Ellis, K., Wong, C., Nye, M., Sable-Meyer, M., Cary, L., Morales, L., et al. (2020). Dreamcoder: Growing generalizable, interpretable knowledge with wake-sleep bayesian program learning. arXiv preprint arXiv:2006.08381.

England, J. L. (2015). Dissipative adaptation in driven self-assembly. *Nature Nanotechnology*, 10, 919–923.

Eslami, S. M. A., Jimenez Rezende, D., Besse, F., Viola, F., Morcos, A. S., Garnelo, M., et al. (2018). Neural scene representation and rendering. *Science*, 360, 1204–1210.

- Fischer, J., Mikhael, J. G., Tenenbaum, J. B., & Kanwisher, N. (2016). Functional neuroanatomy of intuitive physical inference. *Proceedings of the National Academy of Sciences*, 113, E5072–E5081.
- FitzGerald, T. H. B., Dolan, R. J., & Friston, K. (2015). Dopamine, reward learning, and active inference. *Frontiers in Computational Neuroscience*, 9(136).
- FitzGerald, T. H. B., Schwartenbeck, P., Moutoussis, M., Dolan, R. J., & Friston, K. (2015). Active inference, evidence accumulation, and the urn task. *Neural Computation*, 27, 306–328.
- Fotopoulou, A., & Tsakiris, M. (2017). Mentalizing homeostasis: the social origins of interoceptive inference – replies to Commentaries. *Neuropsychanalysis*, 19, 71–76.
- Frank, S. A. (2012). Natural selection. V. How to read the fundamental equations of evolutionary change in terms of information theory. *Journal of Evolutionary Biology*, 25, 2377–2396.
- Freiwald, W. A., & Tsao, D. Y. (2010). Functional compartmentalization and viewpoint generalization within the Macaque Face-Processing system. *Science*, 330, 845–851.
- Friston, K. (2008). Hierarchical models in the brain. *PLoS Computational Biology*, 4, Article e1000211.
- Friston, K. (2010). The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11, 127–138.
- Friston, K. (2013). Life as we know it. *Journal of the Royal Society Interface*, 10, Article 20130475.
- Friston, K. (2019). A free energy principle for a particular physics. (p. 148).
- Friston, K., Adams, R., Perrinet, L., & Breakspear, M. (2012). Perceptions as hypotheses: saccades as experiments. *Frontiers in Psychology*, 3, 151.
- Friston, K., & Buzsáki, G. (2016). The functional anatomy of time: What and when in the brain. *Trends in Cognitive Sciences*, 20, 500–511.
- Friston, K., Da Costa, L., Hafner, D., Hesp, C., & Parr, T. (2021). Sophisticated inference. *Neural Computers*, 33, 713–763.
- Friston, K. J., Daunizeau, J., & Kiebel, S. J. (2009). Reinforcement learning or active inference? *PLoS One*, 4, Article e6421.
- Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., & Pezzulo, G. (2016). Active inference and learning. *Neuroscience & Biobehavioral Reviews*, 68, 862–879.
- Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., & Pezzulo, G. (2017). Active inference: a process theory. *Neural Computation*, 29, 1–49.
- Friston, K., Kilner, J., & Harrison, L. (2006). A free energy principle for the brain. *Journal de Physiologie (Paris)*, 100, 70–87.
- Friston, K. J., Parr, T., & de Vries, B. (2017). The graphical brain: belief propagation and active inference. *Network Neuroscience*, 1, 381–414.
- Friston, K., Rigoli, F., Ognibene, D., Mathys, C., Fitzgerald, T., & Pezzulo, G. (2015). Active inference and epistemic value. *Cognitive Neuroscience*, 6, 187–214.
- Friston, K. J., Rosch, R., Parr, T., Price, C., & Bowman, H. (2017). Deep temporal models and active inference. *Neuroscience & Biobehavioral Reviews*, 77, 388–402.
- Friston, K. J., Rosch, R., Parr, T., Price, C., & Bowman, H. (2018). Deep temporal models and active inference. *Neuroscience & Biobehavioral Reviews*, 90, 486–501.
- Friston, K., Samothrakis, S., & Montague, R. (2012). Active inference and agency: optimal control without cost functions. *Biological Cybernetics*, 106, 523–541.
- Friston, K. J., Shiner, T., FitzGerald, T., Galea, J. M., Adams, R., Brown, H., et al. (2012). Dopamine, affordance and active inference. *PLoS Computational Biology*, 8, Article e1002327.
- Galea, J. M., Bestmann, S., Beigi, M., Jahanshahi, M., & Rothwell, J. C. (2012). Action reprogramming in parkinson's disease: response to prediction error is modulated by levels of dopamine. *Journal of Neuroscience*, 32, 542–550.
- George, D., & Hawkins, J. (2009). Towards a mathematical theory of cortical micro-circuits. *PLoS Computational Biology*, 5, Article e1000532.
- Gershman, S. J., & Blei, D. M. (2012). A tutorial on Bayesian nonparametric models. *Journal of Mathematical Psychology*, 56, 1–12.
- Gibson, J. J. (1977). The theory of affordances. In R. Shaw, & J. Bransford (Eds.), *Perceiving, acting, and knowing: toward an ecological psychology* (pp. 67–82). Hillsdale, NJ: Erlbaum.
- Glaserfeld, E. v. (2002). *Radical constructivism : a way of knowing and learning*. London: RoutledgeFalmer.
- Goldwater, S. (2006). *Nonparametric bayesian models of lexical acquisition*. Brown University.
- Goldwater, S. J. (2007). *Nonparametric bayesian models of lexical acquisition*. Brown University.
- Gonzalez-Billandon, J., Sciutti, A., Sandini, G., & Rea, F. (2020). Towards a cognitive architecture for self-supervised transfer learning for objects detection with a Humanoid Robot. In *2020 joint IEEE 10th international conference on development and learning and epigenetic robotics (icdl-epirob)* (pp. 1–8). IEEE.
- Ha, D., & Schmidhuber, J. (2018). World models. arXiv:1803.10122 [cs, stat].
- Hafner, D., Lillicrap, T., Fischer, I., Villegas, R., Ha, D., Lee, H., et al. (2019). Learning latent dynamics for planning from pixels. In *International conference on machine learning* (pp. 2555–2565). PMLR.
- Hafner, D., Ortega, P. A., Ba, J., Parr, T., Friston, K., & Heess, N. (2020). Action and perception as divergence minimization. arXiv:2009.01791.
- Happé, F., & Frith, U. (2006). The weak coherence account: Detail-focused cognitive style in autism spectrum disorders. *Journal of Autism and Developmental Disorders*, 36, 5–25.
- Hauser, M. D., Chomsky, N., & Fitch, W. T. (2002). The faculty of language: What is it, who has it, and how did it evolve? *Science*, 298, 1569–1579.
- Heyes, C. M. (2018). Cognitive gadgets : the cultural evolution of thinking.
- Himmelstoss, N. A., Schuster, S., Hutzler, F., Moran, R., & Hawelka, S. (2020). Co-registration of eye movements and neuroimaging for studying contextual predictions in natural reading. *Language, Cognition and Neuroscience*, 35, 595–612.
- Hinton, G. E. (2007). Learning multiple layers of representation. *Trends in Cognitive Sciences*, 11, 428–434.
- Hinton, G. E., Dayan, P., Frey, B. J., & Neal, R. M. (1995). The wake-sleep algorithm for unsupervised neural networks. *Science*, 268, 1158–1161.
- Hinton, G. E., Osindero, S., & Teh, Y.-W. (2006). A fast learning algorithm for deep belief nets. *Neural Computation*, 18, 1527–1554.
- Hohwy, J. (2013). *The predictive mind*. Oxford University Press.
- Hohwy, J. (2016). The self-evidencing brain: The self-evidencing brain. *Noûs*, 50, 259–285.
- Horii, T., Nagai, Y., & Asada, M. (2016). Imitation of human expressions based on emotion estimation by mental simulation. *Paladyn Journal of Behavioral Robotics*, 7, 40–54.
- Horii, T., Nagai, Y., & Asada, M. (2018). Modeling development of multi-modal emotion perception guided by tactile dominance and perceptual improvement. *IEEE Transactions on Cognitive and Developmental Systems*, 10, 762–775.
- Hutter, M. (2005). *Universal artificial intelligence: sequential decisions based on algorithmic probability*. Berlin: Springer.
- Iacoboni, M., & Dapretto, M. (2006). The mirror neuron system and the consequences of its dysfunction. *Nature Reviews Neuroscience*, 7, 942–951.
- Imamizu, H., Miyauchi, S., Tamada, T., Sasaki, Y., Takino, R., Putz, B., et al. (2000). Human cerebellar activity reflecting an acquired internal model of a new tool. *Nature*, 403, 192–195.
- Ishii, S., Yoshida, W., & Yoshimoto, J. (2002). Control of exploitation–exploration meta-parameter in reinforcement learning. *Neural Networks*, 15, 665–687.
- Jeffery, K., Pollack, R., & Rovelli, C. (2019). On the statistical mechanics of life: Schrödinger revisited. arXiv:1908.08374 [physics].
- Jordan, M. I., & Rumelhart, D. E. (1992). Forward models: Supervised learning with a distal teacher. *Cognitive Science*, 16, 307–354.
- Kaplan, R., & Friston, K. J. (2018). Planning and navigation as active inference. *Biological Cybernetics*, 112, 323–343.
- Kawato, M. (1999). Internal models for motor control and trajectory planning. *Current Opinion in Neurobiology*, 9, 718–727.
- Kawato, M., Furukawa, K., & Suzuki, R. (1987). A hierarchical neural-network model for control and learning of voluntary movement. *Biological Cybernetics*, 57, 169–185.
- Kawato, M., & Gomi, H. (1992). A computational model of four regions of the cerebellum based on feedback-error learning. *Biological Cybernetics*, 68, 95–103.
- Kingma, D. P., & Welling, M. (2013). Auto-encoding variational bayes. arXiv preprint arXiv:1312.6114.
- Knill, D. C., & Pouget, A. (2004). The Bayesian brain: the role of uncertainty in neural coding and computation. *Trends in Neurosciences*, 27, 712–719.
- Kok, P., Rahnev, D., Jehee, J. F. M., Lau, H. C., & de Lange, F. P. (2012). Attention reverses the effect of prediction in silencing sensory signals. *Cerebral Cortex*, 22, 2197–2206.
- Kumagaya, S. (2015). Tojisha-Kenkyu of autism spectrum disorders. *Advanced Robotics*, 29, 25–34.
- Lake, B. M., Salakhutdinov, R., & Tenenbaum, J. B. (2015). Human-level concept learning through probabilistic program induction. *Science*, 350, 1332–1338.
- Lake, B. M., Ullman, T. D., Tenenbaum, J. B., & Gershman, S. J. (2017). Building machines that learn and think like people. *Behavioral and Brain Sciences*, 40, Article e253.
- Le Roux, N., Heess, N., Shotton, J., & Winn, J. (2011). Learning a generative model of images by factoring appearance and shape. *Neural Computation*, 23, 593–650.
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521, 436–444.
- Lee, T. S., & Mumford, D. (2003). Hierarchical Bayesian inference in the visual cortex. *Journal of the Optical Society of America A*, 20, 1434.
- Levine, S. (2018). Reinforcement learning and control as probabilistic inference: Tutorial and review. arXiv:1805.00909 [cs, stat].
- Limanowski, J. (2017). (Dis-)attending to the body: Action and self-experience in the active inference framework. In *Philosophy and predictive processing: open mind*. Frankfurt am Main: MIND Group.
- Limanowski, J., & Blankenburg, F. (2013). Minimal self-models and the free energy principle. *Frontiers on Human Neurosciences*, 7, 547.
- Lin, H. W., Tegmark, M., & Rolnick, D. (2017). Why does deep and cheap learning work so well? *Journal of Statistical Physics*, 168, 1223–1247.
- Lindley, D. V. (1956). On a measure of the information provided by an experiment. *The Annals of Mathematical Statistics*, 27, 986–1005.

- Loeliger, H. A. (2002). Least squares and Kalman filtering on forney graphs. In *Codes, graphs, and systems* (pp. 113–135). Springer.
- MacKay, D. J. (1995). Free-energy minimisation algorithm for decoding and cryptanalysis. *Electronics Letters*, 31, 445–447.
- Miall, R. C., & Wolpert, D. M. (1996). Forward models for physiological motor control. *Neural Networks*, 9, 1265–1279.
- Millidge, B. (2019). Deep active inference as variational policy gradients. arXiv:1907.03876 [cs].
- Mirza, M. B., Adams, R. A., Mathys, C. D., & Friston, K. J. (2016). Scene construction, visual foraging, and active inference. *Frontiers in Computational Neuroscience*, 10, 56.
- Mirza, M. B., Adams, R. A., Mathys, C., & Friston, K. J. (2018). Human visual exploration reduces uncertainty about the sensed world. *PLoS One*, 13, Article e0190429.
- Miyazawa, K., Aoki, T., Horii, T., & Nagai, T. (2020). LamBERT: Language and action learning using multimodal BERT. arXiv:2004.07093 [cs, stat].
- Mrowca, D., Zhuang, C., Wang, E., Haber, N., Fei-Fei, L. F., Tenenbaum, J., et al. (2018). Flexible neural representation for physics prediction. In *Advances in neural information processing systems* (pp. 8799–8810).
- Mumford, D. (1992). On the computational architecture of the neocortex II the role of cortico-cortical loops. *Biological Cybernetics*, 66, 241–251.
- Nagai, Y. (2019). Predictive learning: its key role in early cognitive development. *Philosophical Transactions of the Royal Society, Series B (Biological Sciences)*, 374, Article 20180030.
- Nakamura, T., Araki, T., Nagai, T., & Iwahashi, N. (2011). Grounding of word meanings in latent Dirichlet allocation-based multimodal concepts. *Advanced Robotics*, 25, 2189–2206.
- Nakamura, T., Nagai, T., & Taniguchi, T. (2018). SERKET: An architecture for connecting stochastic models to realize a large-scale cognitive model. *Frontiers in Neurobotics*, 12.
- Okada, M., Kosaka, N., & Taniguchi, T. (2020). Planet of the Bayesians: Reconsidering and improving deep planning network by incorporating Bayesian inference. arXiv:2003.00370 [cs, stat].
- O'Keefe, J., & Recce, M. L. (1993). Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus*, 3, 317–330.
- O'Neill, M., & Jones, R. S. P. (1997). Sensory-perceptual abnormalities in autism: A case for more research? *Journal of Autism and Developmental Disorders*, 28, 3–293.
- Oudeyer, P.-Y. (2007). What is intrinsic motivation? A typology of computational approaches. *Frontiers in Neurobotics*, 1.
- Parr, T., & Friston, K. J. (2019). Attention or salience? *Current Opinion in Psychology*, 29, 1–5.
- Parr, T., Rikhye, R. V., Halassa, M. M., & Friston, K. J. (2020). Prefrontal computation as active inference. *Cerebral Cortex*, 30, 682–695.
- Parr, T., Sajid, N., & Friston, K. J. (2020). Modules or mean-fields? *Entropy*, 22, 552.
- Pellicano, E., & Burr, D. (2012). When the world becomes 'too real': a Bayesian explanation of autistic perception. *Trends in Cognitive Sciences*, 16, 504–510.
- Pessiglione, M., Seymour, B., Flandin, G., Dolan, R. J., & Frith, C. D. (2006). Dopamine-dependent prediction errors underpin reward-seeking behaviour in humans. *Nature*, 442, 1042–1045.
- Philippsen, A., & Nagai, Y. (2020). Deficits in prediction ability trigger asymmetries in behavior and internal representation. *Frontiers in Psychiatry*, 11, Article 564415.
- Philippsen, A., & Nagai, Y. (2021). A predictive coding account for cognition in human children and chimpanzees: A case study of drawing. *IEEE Transactions on Cognitive and Developmental Systems*, (in press).
- Philippsen, A., Tsuji, S., & Nagai, Y. (2020). Picture completion reveals developmental change in representational drawing ability: An analysis using a convolutional neural network. In *2020 joint IEEE 10th international conference on development and learning and epigenetic robotics (icdl-epirob)* (pp. 1–8).
- Rao, R. P. N., & Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Nanotechnology*, 2, 79–87.
- Rayner, K., Slattery, T. J., Drieghe, D., & Liversedge, S. P. (2011). Eye movements and word skipping during reading: Effects of word length and predictability. *Journal of Experimental Psychology: Human Perception and Performance*, 37, 514.
- Rikhye, R. V., Guntupalli, J. S., Gothoskar, N., Lázaro-Gredilla, M., & George, D. (2019). Memorize-generalize: An online algorithm for learning higher-order sequential structure with cloned hidden Markov models. *BioRxiv*.
- Rizzolatti, G., Fadiga, L., Gallese, V., & Fogassi, L. (1996). Premotor cortex and the recognition of motor actions. *Cognitive Brain Research*, 13, 1–141.
- Rizzolatti, G., Fogassi, L., & Gallese, V. (2001). Neurophysiological mechanisms underlying the understanding and imitation of action. *Nature Reviews Neuroscience*, 2, 661–670.
- Roweis, S., & Ghahramani, Z. (1999). A unifying review of linear Gaussian models. *Neural Computation*, 11, 305–345.
- Rule, J. S., Tenenbaum, J. B., & Piantadosi, S. T. (2020). The child as hacker. *Trends in Cognitive Sciences*, 24, 900–915.
- Sales, A. C., Friston, K. J., Jones, M. W., Pickering, A. E., & Moran, R. J. (2019). Locus coeruleus tracking of prediction errors optimises cognitive flexibility: An active inference model. *PLoS Computational Biology*, 15, Article e1006267.
- Saxe, R., Tenenbaum, J. B., & Carey, S. (2005). Secret agents: Inferences about hidden causes by 10-and 12-month-old infants. *Psychological Science*, 16, 995–1001.
- Schmidhuber, J. (2006). Developmental robotics, optimal artificial curiosity, creativity, music, and the fine arts. *Connection Science*, 18, 173–187.
- Schmidhuber, J. (2010). Formal theory of creativity, fun, and intrinsic motivation (1990–2010). *IEEE Transactions on Autonomous Mental Development*, 2, 230–247.
- Schwartenbeck, P., FitzGerald, T. H. B., Mathys, C., Dolan, R., Kronbichler, M., & Friston, K. (2015). Evidence for surprise minimization over value maximization in choice behavior. *Scientific Reports*, 5, 1–14.
- Schwettmann, S., Tenenbaum, J. B., & Kanwisher, N. (2019). Invariant representations of mass in the human brain. *Elife*, 8.
- Seth, A. K. (2013). Interoceptive inference, emotion, and the embodied self. *Trends on Cognitive Sciences*, 17, 565–573.
- Shidara, M., Kawano, K., Gomi, H., & Kawato, M. (1993). Inverse-dynamics model eye movement control by purkinje cells in the cerebellum. *Nature*, 365, 50–52.
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., van den Driessche, G., et al. (2016). Mastering the game of go with deep neural networks and tree search. *Nature*, 529, 484–489.
- Smith, R., Schwartenbeck, P., Parr, T., & Friston, K. J. (2019). An active inference approach to modeling structure learning: concept learning as an example case. *BioRxiv*, Article 633677.
- Spelke, E. S., & Kinzler, K. D. (2007). Core knowledge. *Developmental Science*, 10, 89–96.
- Srinivasan, M. V., Laughlin, S. B., & Dubs, A. (1982). Predictive coding: a fresh view of inhibition in the retina. *Proceedings of the Royal Society of London. Series B. Biological Sciences*, 216, 427–459.
- Sterzer, P., Adams, R. A., Fletcher, P., Frith, C., Lawrie, S. M., Muckli, L., et al. (2018). The predictive coding account of psychosis. *Biological Psychiatry*, 84, 634–643.
- Still, S., & Precup, D. (2012). An information-theoretic approach to curiosity-driven reinforcement learning. *Theory in Biosciences*, 131, 139–148.
- Sun, Y., Gomez, F., & Schmidhuber, J. (2011). Planning to be surprised: Optimal Bayesian exploration in dynamic environments. In *Artificial general intelligence: 4th international conference*.
- Sutton, R. S., & Barto, A. G. (1987). A temporal-difference model of classical conditioning. *Cognitive Science Society*, 1–24.
- Suzuki, M., Nakayama, K., & Matsuo, Y. (2016). Joint multimodal learning with deep generative models. arXiv:1611.01891 [cs, stat].
- Taniguchi, A., Hagiwara, Y., Taniguchi, T., & Inamura, T. (2017). Online spatial concept and lexical acquisition with simultaneous localization and mapping. In *2017 IEEE/RSJ international conference on intelligent robots and systems (iros)* (pp. 811–818).
- Taniguchi, A., Hagiwara, Y., Taniguchi, T., & Inamura, T. (2020a). Improved and scalable online learning of spatial concepts and language models with mapping. *Autonomous Robots*, 44, 927–946.
- Taniguchi, A., Hagiwara, Y., Taniguchi, T., & Inamura, T. (2020b). Spatial concept-based navigation with human speech instructions via probabilistic inference on Bayesian generative model. *Advanced Robotics*, 34, 1213–1228.
- Taniguchi, T., Mochihashi, D., Nagai, T., Uchida, S., Inoue, N., Kobayashi, I., et al. (2019). Survey on frontiers of language and robotics. *Advanced Robotics*, 33, 700–730.
- Taniguchi, T., Nagai, T., Nakamura, T., Iwahashi, N., Ogata, T., & Asoh, H. (2016). Symbol emergence in robotics: a survey. *Advanced Robotics*, 30, 706–728.
- Taniguchi, T., Nakamura, T., Suzuki, M., Kuniyasu, R., Hayashi, K., Taniguchi, A., et al. (2020). Neuro-SERKET: development of integrative cognitive system through the composition of deep probabilistic generative models. *New Generation Computing*, 1–26.
- Taniguchi, T., Piater, J., Worgotter, F., Ugur, E., Hoffmann, M., Jamone, L., et al. (2019). Symbol emergence in cognitive developmental systems: A survey. *IEEE Transactions on Cognitive and Developmental Systems*, 11, 494–516.
- Teh, Y. W., Jordan, M. I., Beal, M. J., & Blei, D. M. (2006). Hierarchical Dirichlet processes. *Journal of the American Statistical Association*, 101, 1566–1581.
- Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to grow a mind: Statistics, structure, and abstraction. *Science*, 331, 1279–1285.
- Tervo, D. G., Tenenbaum, J. B., & Gershman, S. J. (2016). Toward the neural implementation of structure learning. *Current Opinion on Neurobiology*, 37, 99–105.
- Tian, L. Y., Ellis, K., Kryven, M., & Tenenbaum, J. B. (2020). Learning abstract structure for drawing by efficient motor program induction. *Advances in Neural Information Processing Systems*, 33, 12.
- Tononi, G., & Cirelli, C. (2006). Sleep function and synaptic homeostasis. *Sleep Medicine Reviews*, 10, 49–62.
- Tschantz, A., Baltieri, M., Seth, A. K., & Buckley, C. L. (2020). Scaling active inference. In *2020 international joint conference on neural networks (ijcnn)* (pp. 1–8).

- Ullman, T. D., & Tenenbaum, J. B. (2020). Bayesian models of conceptual development: Learning as building models of the world. *Annual Review of Developmental Psychology*, 2, 533–558.
- Ungerleider, L. (1994). 'what' and 'where' in the human brain. *Current Opinion in Neurobiology*, 4, 157–165.
- Van de Cruys, S., Evers, K., Van der Hallen, R., Van Eylen, L., Boets, B., de Wit, L., et al. (2014). Precise minds in uncertain worlds: Predictive coding in autism. *Psychological Review*, 121, 649–675.
- Vernon, D., Metta, G., & Sandini, G. (2007). The icub cognitive architecture: Interactive development in a humanoid robot. In *2007 IEEE 6th international conference on development and learning* (pp. 122–127).
- Von Uexküll, J. (1992). A stroll through the worlds of animals and men: A picture book of invisible worlds. *Semiotica*, 89, 319–391.
- Wallace, C. S., & Dowe, D. L. (1999). Minimum message length and Kolmogorov complexity. *Computer Journal*, 42, 270–283.
- Warneken, F., & Tomasello, M. (2009). The roots of human altruism. *British Journal of Psychology*, 100, 455–471.
- Whittington, J. C. R., & Bogacz, R. (2017). An approximation of the error backpropagation algorithm in a predictive coding network with local hebbian synaptic plasticity. *Neural Computation*, 29, 1229–1262.
- Winn, J., & Bishop, C. M. (2005). Variational message passing. *Journal of Machine Learning Research*, 6, 661–694.
- Wolpert, D., Ghahramani, Z., & Jordan, M. (1995). An internal model for sensorimotor integration. *Science*, 269, 1880–1882.
- Yamins, D. L., Hong, H., Cadieu, C. F., Solomon, E. A., Seibert, D., & DiCarlo, J. J. (2014). Performance-optimized hierarchical models predict neural responses in higher visual cortex. *Proceeding on the National Academy Sciences United States of America*, 111, 8619–8624.
- Yedidia, J. S., Freeman, W. T., & Weiss, Y. (2005). Constructing free-energy approximations and generalized belief propagation algorithms. *IEEE Transactions on Information Theory*, 51, 2282–2312.
- Zeki, S., & Shipp, S. (1988). The functional logic of cortical connections. *Nature*, 335, 311–317.