



ELSEVIER

Contents lists available at ScienceDirect

Neuroscience and Biobehavioral Reviews

journal homepage: www.elsevier.com/locate/neubiorev

Deep temporal models and active inference

Karl J. Friston^{a,*}, Richard Rosch^a, Thomas Parr^a, Cathy Price^a, Howard Bowman^{b,c}^a Wellcome Trust Centre for Neuroimaging, Institute of Neurology, University College London, WC1N 3BG, United Kingdom^b Centre for Cognitive Neuroscience and Cognitive Systems and the School of Computing, University of Kent at Canterbury, Canterbury, Kent, CT2 7NF, United Kingdom^c School of Psychology, University of Birmingham, Edgbaston, Birmingham, B15 2TT, United Kingdom

ARTICLE INFO

Keywords:

Active inference
Bayesian
Hierarchical
Reading
Violation
Free energy
P300
MMN

ABSTRACT

How do we navigate a deeply structured world? Why are you reading this sentence first – and did you actually look at the fifth word? This review offers some answers by appealing to active inference based on deep temporal models. It builds on previous formulations of active inference to simulate behavioural and electrophysiological responses under hierarchical generative models of state transitions. Inverting these models corresponds to sequential inference, such that the state at any hierarchical level entails a sequence of transitions in the level below. The deep temporal aspect of these models means that evidence is accumulated over nested time scales, enabling inferences about narratives (i.e., temporal scenes). We illustrate this behaviour with Bayesian belief updating – and neuronal process theories – to simulate the epistemic foraging seen in reading. These simulations reproduce perisaccadic delay period activity and local field potentials seen empirically. Finally, we exploit the deep structure of these models to simulate responses to local (e.g., font type) and global (e.g., semantic) violations; reproducing mismatch negativity and P300 responses respectively.

1. Introduction

In recent years, we have applied the free energy principle to generative models of worlds that can be described in terms of discrete states in an attempt to understand the embodied Bayesian brain. The resulting active inference scheme (for Markov decision processes) has been applied in a variety of domains (see Table 1). This paper takes active inference to the next level and considers hierarchical models with deep temporal structure (George and Hawkins, 2009; Kiebel et al., 2009; LeCun et al., 2015). This structure follows from generative models that entertain state transitions or sequences over time. The resulting model enables inference about narratives with deep temporal structure (c.f., sequential scene construction) of the sort seen in reading. In short, equipping an agent or simulated subject with deep temporal models allows them to accumulate evidence over different temporal scales to find the best explanation for their sensations.

This paper has two agendas: to introduce hierarchical (deep) generative models for active inference under Markov decision processes (or hidden Markov models) and to show how their belief updating can be understood in terms of neuronal processes. The problem we focus on is how subjects deploy active vision to disambiguate the causes of their sensations. In other words, we ask how people choose where to look

next, when resolving uncertainty about the underlying conceptual, semantic or lexical causes of sensory input. This means that we are not concerned with computational linguistics *per se* but the more general problem of *epistemic foraging*, while using reading as an example.

Epistemics is at the heart of active inference, which is all about reducing surprise or uncertainty, where uncertainty is expected surprise. Technically, this means that one can describe both inference (perception) and behaviour (action) in terms of minimising a free energy functional of probabilistic or Bayesian beliefs. In this setting, variational free energy approximates surprise and expected free energy approximates uncertainty (a.k.a. entropy). This single imperative provides an inclusive account of established (normative) approaches to perception and action; for example, the principle of maximum mutual information, the principle of minimum redundancy, formulations of saliency as Bayesian surprise, risk sensitive or KL control, expected utility theory, and so on (Barlow, 1974; Itti and Baldi, 2009; Kappen et al., 2012; Ortega and Braun, 2013). Our focus here is on how subjects use accumulated beliefs about the hidden states of the world to prescribe active sampling of new information to resolve their uncertainty quickly and efficiently (Ferro et al., 2010).

Our second agenda is to translate these normative (variational) principles into neurobiology by trying to establish the construct validity

DOI of original article: <http://dx.doi.org/10.1016/j.neubiorev.2017.04.009>

* Corresponding author at: The Wellcome Trust Centre for Neuroimaging, Institute of Neurology, 12 Queen Square, London, WC1N 3BG, United Kingdom.

E-mail addresses: k.friston@ucl.ac.uk (K.J. Friston), r.rosch@ucl.ac.uk (R. Rosch), thomas.parr.12@ucl.ac.uk (T. Parr), c.j.price@ucl.ac.uk (C. Price), H.Bowman@kent.ac.uk (H. Bowman).<https://doi.org/10.1016/j.neubiorev.2018.04.004>0149-7634/© 2018 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Table 1
Applications of active inference for Markov decision processes.

Application	Comment	References
Decision making under uncertainty	Initial formulation of active inference for <i>Markov decision processes</i> and <i>sequential policy optimisation</i>	(Friston et al., 2012b)
Optimal control (the mountain car problem) Evidence accumulation: Urns task	Illustration of <i>risk sensitive or KL control</i> in an engineering benchmark Demonstration of how beliefs states are absorbed into a generative model	(Friston et al., 2012a) (FitzGerald et al., 2015b,c)
Addiction Dopaminergic responses	Application to psychopathology Associating dopamine with the encoding of (expected) precision provides a plausible account of dopaminergic discharges	(Schwartenbeck et al., 2015c) (Friston et al., 2014; FitzGerald et al., 2015a)
Computational fMRI	Using Bayes optimal precision to predict activity in dopaminergic areas	(Schwartenbeck et al., 2015a)
Choice preferences and epistemics	Empirical testing of the hypothesis that people prefer to keep options open	(Schwartenbeck et al., 2015b)
Behavioural economics and trust games Foraging and two step mazes	Examining the effects of prior beliefs about self and others Formulation of epistemic and pragmatic value in terms of <i>expected free energy</i>	(Moutoussis et al., 2014) (Friston et al., 2015)
Habit learning, reversal learning and devaluation	Learning as minimising variational free energy with respect to model parameters – and action selection as <i>Bayesian model averaging</i>	(FitzGerald et al., 2014; Friston et al., 2016)
Saccadic searches and scene construction	<i>Mean field approximation</i> for multifactorial hidden states, enabling high dimensional beliefs and outcomes: c.f., functional segregation	(Friston and Buzsaki, 2016; Mirza et al., 2016)
Electrophysiological responses: <i>place-cell activity, omission related responses, mismatch negativity, P300, phase-progression, theta-gamma coupling</i>	Simulating neuronal processing with a gradient descent on variational free energy; c.f., dynamic <i>Bayesian belief propagation</i> based on marginal free energy	In press
Structure learning, sleep and insight	Inclusion of parameters into expected free energy to enable structure learning via <i>Bayesian model reduction</i>	Under review
Narrative construction and reading	Hierarchical generalisation of generative model with <i>deep temporal structure</i>	Current paper

of active inference in terms of behaviour and electrophysiological responses. We do this at three levels: first, by highlighting the similarity between the message passing implied by minimising variational free energy and the neurobiology of neuronal circuits. Specifically, we try to associate the dynamics of a gradient descent on variational free energy with neuronal dynamics based upon neural mass models (Lopes da Silva, 1991). Furthermore, the exchange of sufficient statistics implicit in belief propagation is compared with the known characteristics of extrinsic (between cortical area) and intrinsic (within cortical area) neuronal connectivity. Second, we try to reproduce reading-like behaviour – in which epistemically rich information is sampled by sparse, judicious saccadic eye movements. This enables us to associate perisaccadic updating with empirical phenomena, such as delay period activity and perisaccadic local field potentials (Kojima and Goldman-Rakic, 1982; Purpura et al., 2003; Pastalkova et al., 2008). Finally, in terms of the non-invasive electrophysiology, we try to reproduce the well-known violation responses indexed by phenomena like the mismatch negativity (MMN) and P300 waveforms in event related potential research (Strauss et al., 2015).

This paper comprises four sections. The first (Active inference and free energy) briefly reviews active inference, establishing the normative principles that underlie action and perception. The second section (Belief propagation and neuronal networks) considers action and perception, paying special attention to hierarchical generative models and how the minimisation of free energy could be implemented in the brain. The third section (Simulations of reading) introduces a particular generative model used to simulate reading and provides an illustration of the ensuing behaviour – and simulated electrophysiological responses. The final section (Simulations of classical violation responses) rehearses the reading simulations using different prior beliefs to simulate responses to violations at different hierarchical levels in the model.

2. Active inference and free energy

Active inference rests upon a generative model that is used to infer the most likely causes of observable outcomes in terms of expected states of the world. A generative model is just a probabilistic specification of how consequences (outcomes) follow from causes (states).

These states are called latent or *hidden* because they can only be inferred through observations. Clearly, observations depend upon action (e.g., where you are looking). This requires the generative model to represent outcomes under different actions or policies. Technically, expectations about (future) outcomes and their hidden causes are optimised by minimising variational free energy, which renders them the most likely (posterior) expectations about the (future) states of the world, given (past) observations. This follows because the variational free energy is an upper bound on (negative) log Bayesian model evidence; also known as surprise, surprisal or self-information (Dayan et al., 1995). Crucially, the prior probability of each policy (i.e., action or plan) is the free energy expected under that policy (Friston et al., 2015). This means that policies are more probable if they minimise expected surprise or resolve uncertainty.

Evaluating the expected free energy of plausible policies – and implicitly their posterior probabilities – enables the most likely action to be selected. This action generates a new outcome and the cycle of perception and action starts again. The resulting behaviour represents a principled sampling of sensory cues that has epistemic, uncertainty reducing and pragmatic, surprise reducing aspects. The pragmatic aspect follows from prior beliefs or preferences about future outcomes that makes some outcomes more surprising than others. For example, I would not expect to find myself dismembered or humiliated – and would therefore avoid these surprising state of affairs. On this view, behaviour is dominated by epistemic imperatives until there is no further uncertainty to resolve. At this point pragmatic (prior) preferences predominate, such that explorative behaviour gives way to exploitative behaviour. In this paper, we focus on epistemic behaviour and only use prior preferences to establish a task or instruction set. Namely, to report a categorical decision when sufficiently confident; i.e., under the prior belief one does not make mistakes.

2.1. Hierarchical generative models

We are concerned here with hierarchical generative models in which the outcomes of one level generate the hidden states at a lower level. Fig. 1 provides a schematic of this sort of model. Outcomes depend upon hidden states, while hidden states unfold in a way that

Download English Version:

<https://daneshyari.com/en/article/7302002>

Download Persian Version:

<https://daneshyari.com/article/7302002>

[Daneshyari.com](https://daneshyari.com)