


Federated inference and belief sharing

Karl J Friston, Thomas Parr, Conor Heins, Axel Constant, Daniel Friedman, Takuya Isomura, Chris Fields, Tim Verbelen, Maxwell Ramstead, John Clippinger, Christopher D Frith



PII: S0149-7634(23)00469-4

DOI: <https://doi.org/10.1016/j.neubiorev.2023.105500>

Reference: NBR105500

To appear in: *Neuroscience and Biobehavioral Reviews*

Received date: 4 August 2023

Revised date: 8 November 2023

Accepted date: 1 December 2023

Please cite this article as: Karl J Friston, Thomas Parr, Conor Heins, Axel Constant, Daniel Friedman, Takuya Isomura, Chris Fields, Tim Verbelen, Maxwell Ramstead, John Clippinger and Christopher D Frith, Federated inference and belief sharing, *Neuroscience and Biobehavioral Reviews*, (2023) doi:<https://doi.org/10.1016/j.neubiorev.2023.105500>

This is a PDF file of an article that has undergone enhancements after acceptance, such as the addition of a cover page and metadata, and formatting for readability, but it is not yet the definitive version of record. This version will undergo additional copyediting, typesetting and review before it is published in its final form, but we are providing this version to give early visibility of the article. Please note that, during the production process, errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

© 2023 Published by Elsevier.

# Federated inference and belief sharing

Karl J Friston<sup>1,2</sup>, Thomas Parr<sup>1</sup>, Conor Heins<sup>2,3,4,5</sup>, Axel Constant<sup>6,2</sup>, Daniel Friedman<sup>7,8</sup>, Takuya Isomura<sup>9</sup>, Chris Fields<sup>10</sup>, Tim Verbelen<sup>2</sup>, Maxwell Ramstead<sup>1,2</sup>, John Clippinger<sup>11</sup>, Christopher D Frith<sup>12</sup>

*k.friston@ucl.ac.uk; thomas.parr.12@ucl.ac.uk; conor.heins@gmail.com; axel.constant.pruvost@gmail.com; danielariefriedman@gmail.com; takuya.isomura@riken.jp; fieldsres@gmail.com; tim.verbelen@verses.ai; maxwell.ramstead@verses.ai; jhclipman@gmail.com; c.frith@ucl.ac.uk*

<sup>1</sup> Wellcome Trust Centre for Neuroimaging, Institute of Neurology, University College London, UK.

<sup>2</sup> VERSES Research Lab, Los Angeles, California, 90016, USA

<sup>3</sup> Department of Collective Behaviour, Max Planck Institute of Animal Behavior, 78457 Konstanz, Germany

<sup>4</sup> Centre for the Advanced Study of Collective Behaviour, 78457 Konstanz, Germany

<sup>5</sup> Department of Biology, University of Konstanz, 78457 Konstanz, Germany

<sup>6</sup> School of Engineering and Informatics, The University of Sussex, Brighton, UK

<sup>7</sup> Department of Entomology and Nematology, University of California, Davis, Davis, CA, USA

<sup>8</sup> Active Inference Institute, Davis, CA, 95616, USA

<sup>9</sup> Brain Intelligence Theory Unit, RIKEN Center for Brain Science, Wako, Saitama 351-0198, Japan

<sup>10</sup> Allen Discovery Center at Tufts University, Medford, MA 02155, USA

<sup>11</sup> Bioform Labs, and MIT Media Lab, Boston, USA

<sup>12</sup> Institute of Philosophy, School of Advanced Studies, University of London, UK

**Correspondence:** Karl Friston

The Wellcome Trust Centre for Neuroimaging

Queen Square Institute of Neurology

12 Queen Square, London, UK WC1N 3AR

k.friston@ucl.ac.uk

## Highlights

- This paper offers a proof of concept that communication in general—and language in particular—is an emergent property of conspecifics that seek evidence for generative models of their shared, co-constructed world.
- It uses numerical studies to show that nested free energy minimising—or evidence maximising—processes are sufficient to explain (i) the (horizontal) emergence of language among agents that can broadcast their beliefs and (ii) its (vertical) transmission over generations.
- By grounding these processes in formal (Bayesian) terms, one can join various perspectives on communication; ranging from generalised synchrony in dynamical systems to cultural niche construction in evolutionary psychology.

## Abstract

This paper concerns the distributed intelligence or federated inference that emerges under belief-sharing among agents who share a common world—and world model. Imagine, for example, several animals keeping a lookout for predators. Their collective surveillance rests upon being able to communicate their beliefs—about what they see—among themselves. But, how is this possible? Here, we show how all the necessary components arise from minimising free energy. We use numerical studies to simulate the generation, acquisition and emergence of language in synthetic agents. Specifically, we consider inference, learning and selection as minimising the variational free energy of posterior (i.e., Bayesian) beliefs about the states, parameters and structure of generative models, respectively. The common theme—that attends these optimisation processes—is the selection of actions that minimise expected free energy, leading to active inference, learning and model selection (a.k.a., structure learning). We first illustrate the role of communication in resolving uncertainty about the latent states of a partially observed world, on which agents have complementary perspectives. We then consider the acquisition of the requisite language—entailed by a likelihood mapping from an agent's beliefs to their overt expression (e.g., speech)—showing that language can be transmitted across generations by active learning. Finally, we show that language is an emergent property of free energy minimisation, when agents operate within the same econiche. We conclude with a discussion of various perspectives on these phenomena; ranging from cultural niche construction, through federated learning, to the emergence of complexity in ensembles of self-organising systems.

**Keywords:** *active inference · distributed cognition · federated learning · structure learning · message passing*

## 1 INTRODUCTION

---

This paper concerns the genesis of communication and distributed cognition in ensembles of agents who share the same world—and internal or generative model of that world (Bahrami et al., 2010; Constant et al., 2019; Friston and Frith, 2015; Friston et al., 2022b; Frith and Wentzer, 2013; Kairouz et al., 2021; Levin, 2019; Vasil et al., 2020). It uses simulations of agents who broadcast their beliefs about inferred states of the world to other agents, enabling them to engage in joint inference and learning. Consider, for example, three birds monitoring their environment for predators. Each bird will have more or less precise beliefs about the location and disposition of potential predators. If they broadcast these beliefs to each other, a shared belief about the world could emerge that would be more precise than any individual's belief. We simulate the requisite belief-sharing using a series of free

energy minimising processes to reproduce active inference, learning and selection, where selection refers to the selection of structures or functional forms for generative models of the sensed world.

By equipping agents with a shared generative model—of a shared world—one can effectively distribute or federate Bayesian belief updating over multiple agents—where this federation rests upon the sharing of posterior beliefs about hidden (a.k.a., latent) states of affairs (Kairouz et al., 2021). Clearly, for this to work, collocutors must share a common ground or frame of reference (Adank et al., 2010; Allan, 2013; Barsalou, 2008; Chomsky, 2006; Garrod and Pickering, 2009). In other words, there has to be an isomorphism between the beliefs of one agent and those of another. This necessitates generative models in which some belief structures, or representations, are conserved over agents. Beliefs over this common ground can, in principle, be shared among agents—so that every agent inherits and assimilates the perspective of others. Clearly, to realise this kind of belief-sharing, it is necessary to generate observable outcomes or messages that can be recognised by other agents. In turn, this necessitates a common mapping between shared beliefs and their overt expression (e.g., language); namely, a shared likelihood model responsible for language generation and recognition. In what follows, we will focus on the *deployment*, *acquisition* and *emergence* of this likelihood model, foregrounding the role of *active inference*, *active learning* and *active selection*, respectively. This emergence recapitulates numerical studies of collective intelligence; in which the alignment between agents’ free energy minima “emerges endogenously from the dynamics of interacting AIF [active inference] agents themselves” (Kaufmann et al., 2021).

There are many perspectives that one could take on these nested processes; ranging from peer-to-peer message passing and belief-sharing in computer science to the emergence of language in developmental and evolutionary psychology (Constant et al., 2018; Frith, 2010; Ghazanfar and Takahashi, 2014; Hauser et al., 2002; Heyes, 2018; Kastel and Hesp, 2021; Laland et al., 2016; Steels, 2011; Tomasello, 2016; Veissiere et al., 2019). We will pick up these perspectives in the discussion but first rehearse their theoretical foundations through the lens of the free energy principle (Friston et al., 2022a; Ramstead et al., 2022).

The simulations used in this study are realisations of free energy minimising processes at distinct temporal scales that can be interpreted in terms of inference, learning and selection. Free energy minimisation is just a way of describing self organisation in open—and therefore nonequilibrium—systems that are coupled to each other. Here, we consider several agents that are in exchange with a common environment, and each other. Each agent can be regarded as a (Bayesian) belief updating process, where each agent entails a generative model of her environment<sup>1</sup>. The beliefs in question

---

<sup>1</sup> For clarity, we will use the pronoun ‘she’ for agents and ‘it’ for things (e.g., subjects) that agents have beliefs about.

pertain to latent states, parameters and structures: all of which change to minimise a free energy functional; namely, a variational free energy that incorporates expected free energy (Parr and Friston, 2019). Expected free energy could be regarded as definitive of active inference, in the sense that it scores the likelihood of various actions, while variational free energy scores the marginal likelihood of observations, under some (Bayesian) beliefs about how observations were caused.

In the current application of the free energy principle, we generalise the notion of action to anything that has consequences. At the level of inference, action corresponds to sampling or selecting things that are expected to minimise free energy (Friston et al., 2011). At the level of learning, action involves updating beliefs about the parameters of the generative model: e.g., active learning (Mackay, 1992; Schmidhuber, 2010; Vigorito and Barto, 2010). Similarly, in model selection or structure learning (Gershman and Niv, 2010; Pellet and Elisseff, 2008; Smith et al., 2020; Tervo et al., 2016) action ‘selects’ certain priors over parameters. We will see later, that—from the point of view of model parameters and structure—selecting updates that minimise expected free energy is equivalent to maximising the mutual information of the likelihood mapping between observable consequences and latent or unobservable causes, under certain constraints.

Minimising variational free energy—with respect to beliefs over states and parameters—is equivalent to maximising the model evidence (a.k.a., marginal likelihood) of observations, under the agent's generative model (Winn and Bishop, 2005). This is sometimes referred to as self-evidencing (Fields et al., 2021a; Hohwy, 2016). Put simply, to self-evidence means to get a better grip on the world—a grip that may be tighter when agents pull together by sharing their beliefs (Bruineberg et al., 2018; Constant et al., 2019). Interesting, this rests on getting a grip on the common ground that enables agents to make sense of each other.

In what follows, we unpack the emergence of self-evidencing and common ground in three steps. We first demonstrate the role of belief-sharing, in resolving uncertainty about hidden states, when their observable consequences can only be seen by one agent at a time. We then illustrate the acquisition of language (i.e., likelihood mappings) using active learning or accumulation of Dirichlet counts, where a child learns from its parents and conspecifics. Finally, we illustrate the emergence of language (i.e., precise and shared likelihood mappings) in an ensemble of agents. This emergence rests upon structure learning, in which priors over the Dirichlet counts of likelihood tensors are updated to minimise expected free energy. The ensuing active (Bayesian) model selection uses Bayesian model reduction (Friston et al., 2018).

This paper comprises four key sections. The first provides a summary of active inference under generative models of discrete state spaces; in other words, where the hidden or latent states, causing observations, are in one discrete state or another. This section offers a technical preamble for readers

who want to understand the mechanics of belief updating used in subsequent sections (and other applications of active inference to discrete models). The second section focuses on the special issue of how beliefs are shared among agents, where each agent is equipped with her own generative model of the world. We will use a simple setup to illustrate belief sharing, in which three agents broadcast their beliefs about the causes of their sensations. All three agents share a generative model of what they would ‘see’—from their unique perspectives—given states of affairs in their shared world. Crucially, they also share a generative model of what they would ‘hear’ if beliefs about those states were articulated or broadcast. The implicit communication enables the updating of beliefs based upon what all three agents can see; effectively, accumulating evidence through ‘many pairs of eyes’. The augmented efficiency of belief updating—and its neuronal correlates—rest upon assuming a shared generative model, in which the mapping between each agent’s beliefs and what she says (and hears) is conserved over agents. In the second section, we ask whether this shared mapping could be learned by a naïve agent (e.g., an infant) exposed to the same visual and auditory scenes as the communicating agents (e.g., parents). We will see that the requisite mapping emerges as a consequence of minimising free energy, rendering the sensorium as predictable as possible. These simulations assume the existence of a language that can be learned. In the final simulations, we ask whether free energy minimising processes are sufficient to account for the *de novo* emergence of likelihood mappings—that underwrite linguistic exchange—by exposing three naïve agents to shared visual scenes. Again, we will see the emergence of precise likelihood mappings that are conserved over agents. In short, communication appears to be an emergent property of agents who broadcast their beliefs about a shared world.

In summary, we hoped to show—using numerical analyses—that federated inference and learning emerges from minimising (variational, expected and reduced) free energy. Formally, this tuple of free energy minimising processes provides a first principles account of (state) inference, (parameter) learning and (model) selection that can be applied to any generative model and implicit agent. Our focus here is on its application to an ensemble of agents to evince self-evidencing through belief-sharing. This could be read as a move towards a formal account of consciousness in the pre-Cartesian sense of sharable knowledge (i.e., *con* – ‘together’ and *scire* – ‘to know’). More practically, it might speak to the design principles for ecosystems of intelligent agents (Friston and Frith, 2015; Friston et al., 2022b; Frith and Wentzer, 2013)<sup>2</sup>.

---

<sup>2</sup> To avoid disrupting the narrative, we will put technical details—concerning variational message passing and generative models—in figure legends. This allows us to focus on the conceptual foundations and key functional forms of requisite (free energy minimising) processes.

## 2 ACTIVE INFERENCE AND FREE ENERGY

---

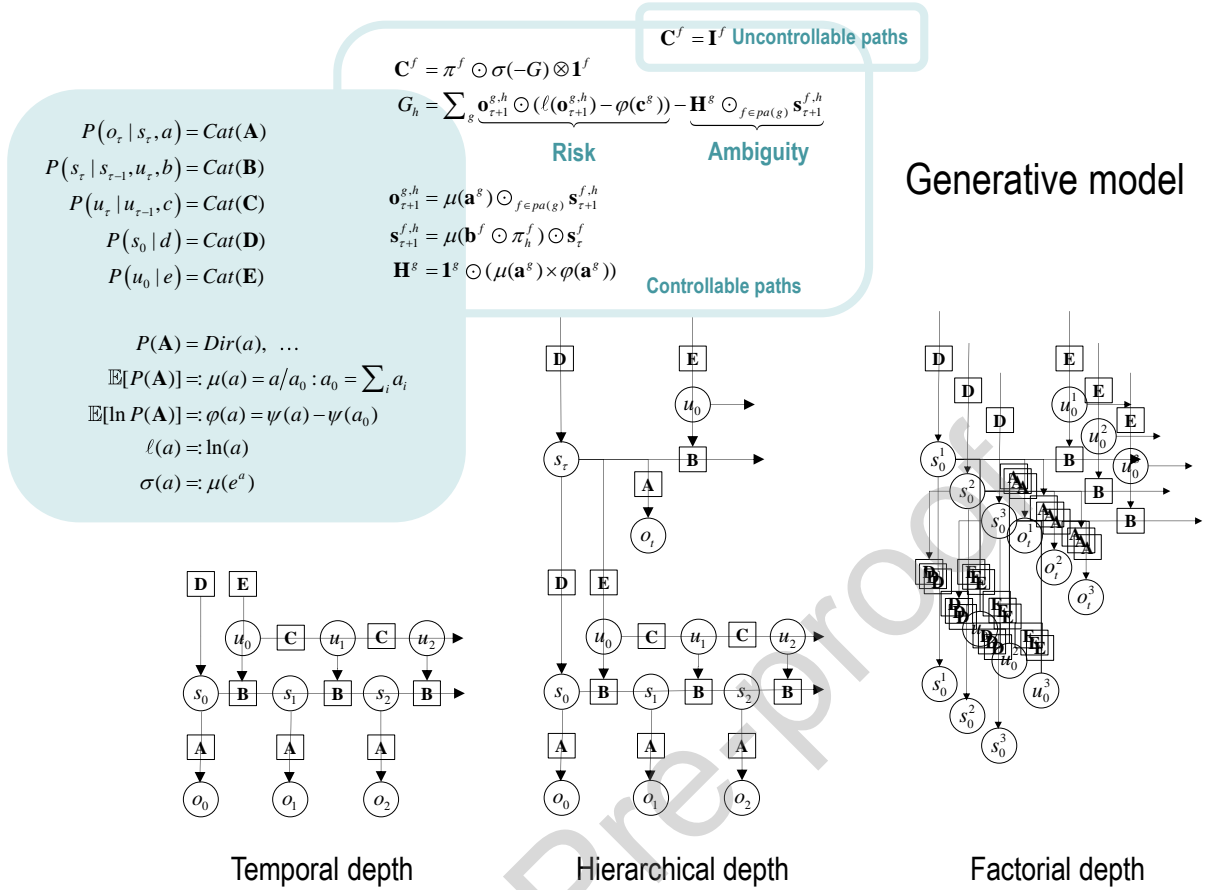
Active inference rests upon a *generative model* of observable outcomes. This model is used to infer the most likely causes of outcomes in terms of expected states of the world. These states (and paths) are latent or *hidden* because they can only be inferred through observations. Some paths are considered controllable in the sense they can be changed by acting. Crucially, certain observations depend upon action (e.g., where one is looking), which requires the generative model to entertain expectations about outcomes under different combinations of actions (i.e., policies)<sup>3</sup>. These expectations are optimised by minimising *variational free energy*. Crucially, the prior probability of a policy depends upon its *expected free energy*. Expected free energy has a number of familiar special cases; including, expected utility, intrinsic value, Bayesian surprise, mutual information, *etc.* Having evaluated the expected free energy of each policy—and implicitly their prior likelihood—the most likely action can be selected. This action generates a new outcome and the (perception-action) cycle starts again (Parr et al., 2022).

### 2.1 THE GENERATIVE MODEL

Figure 1 provides a schematic specification of the generative model used for the sorts of problems considered in this paper. In brief, outcomes at any particular time depend upon hidden *states*, while transitions among hidden states depend upon *paths*. Note that paths are random variables in the sense that a particle can have both a position (i.e., a state) and momentum (i.e., a path). Paths may or may not depend upon action. The resulting partially observed Markov decision process (POMDP) is specified by a set of tensors. The first set **A**, maps from hidden states to outcome modalities; for example, exteroceptive (e.g., visual) or proprioceptive (e.g., eye position) modalities. These parameters encode the likelihood of an outcome given their hidden causes. The second set **B** prescribes transitions among the *factors* of hidden states, under a particular path. These factors correspond to different states of the world, like the location or nature of an object. The remaining tensors encode prior beliefs about paths **C**, and initial states **D**. The tensors—encoding probabilistic mappings or contingencies—are generally parameterised as Dirichlet distributions, whose sufficient statistics are concentration parameters or *Dirichlet counts*. These count the number of times a particular combination of states or outcomes has been observed. We will focus on learning the likelihood model, encoded by Dirichlet counts, **a**.

---

<sup>3</sup> Note that in this setting, a policy is not a sequence of actions but simply a combination of paths, where each hidden factor has an associated state and path. This means there are, potentially, as many policies as there are combinations of paths.



**Figure 1 – Generative models as agents.** A generative model specifies the joint probability of outcomes or consequences and their causes; namely, hidden states. This joint distribution can be expressed in terms of a *likelihood* (the probability of consequences given their causes) and *priors* (over causes). When a prior depends upon a random variable it is called an *empirical prior*. Here, the likelihood is specified by a tensor  $\mathbf{A}$ , encoding the probability of an outcome under every combination of *states* ( $s$ ). The empirical priors in this instance pertain to transitions among hidden states  $\mathbf{B}$  that depend upon *paths* ( $u$ ), whose transition probabilities are encoded in  $\mathbf{C}$ . The key aspect of this generative model is that certain (controllable) paths are more probable *a priori* if they minimise their expected free energy ( $\mathbf{G}$ ), expressed in terms of *risk* and *ambiguity* (lower white panel). If the path is not controllable, it remains unchanged during the epoch in question (upper white panel), where  $\mathbf{E}$  specifies the initial probability of each path. The left panel provides the functional form of the generative model in terms of categorical (*Cat*) distributions. The lower equalities list the various operators required for the variational message-passing—and updating Bayesian beliefs about hidden states and paths—detailed in Figure 2. These functions are taken to operate on each column of their tensor arguments. The graph on the lower left depicts the generative model as a probabilistic graphical model that foregrounds the implicit *temporal depth* implied by priors over state transitions and paths. This example only shows dependencies for uncontrollable paths. When equipped with *hierarchical depth* the POMDP acquires a separation of temporal scales. This follows from a construction in which the succession of states at a higher level generates the initial states (via the  $\mathbf{D}$  tensor) and paths (via the  $\mathbf{E}$  tensor) at the lower level. This means higher levels unfold more slowly than lower levels, thereby furnishing empirical priors (c.f., inductive biases) that contextualise the dynamics of lower level states. At each hierarchical level, the hidden states and accompanying paths are factored to endow the model with *factorial depth*. In other words, the model or agent ‘carves nature at its joints’ into factors that conspire or interact to generate outcomes (or the initial states and paths at lower levels). This means context-sensitive contingencies are mediated by the tensors mapping from one level to the next ( $\mathbf{D}$  and  $\mathbf{E}$ ) or outcomes ( $\mathbf{A}$ ). The subscripts in this figure pertain to



time, while the superscripts denote different factors ( $f$ ), outcome modalities ( $g$ ) and combinations of paths over factors ( $h$ ). Tensors and matrices are denoted by uppercase bold, while posterior expectations are in lowercase bold. The matrix  $\pi$  encodes the probability over paths, under each policy (where  $\pi_h$  denotes the probability over paths for policy  $h$ ). The  $\odot$  notation implies a generalised inner (i.e., dot) product or tensor contraction, while  $\times$  denotes the Hadamard (element by element) product.  $ch(\cdot)$  and  $pa(\cdot)$  return the children and parents of any node; namely, the co-domain and domain, respectively. Finally,  $\psi$  denotes the digamma function (the logarithmic derivative of the gamma function).

The generative model in Figure 1 means that outcomes are generated as follows: first, a policy is selected using a softmax function of expected free energy. Sequences of hidden states are generated using the probability transitions specified by the selected combination of paths (i.e., policy). Finally, these hidden states generate outcomes in one or more modalities. Perception or inference about hidden states (i.e., state estimation) corresponds to inverting a generative model, given a sequence of outcomes, while learning corresponds to updating model parameters. Perception therefore corresponds to accumulating evidence for beliefs about hidden states and paths, while learning corresponds to accumulating knowledge in the form of Dirichlet counts. The requisite expectations constitute the sufficient statistics  $(\mathbf{s}, \mathbf{u}, \mathbf{a})$  of posterior beliefs  $Q(s, u, a) = Q_s(s)Q_u(u)Q_a(a)$ . Because we are dealing with discrete states,  $(\mathbf{s}, \mathbf{u}, \mathbf{a})$  are just the expected probability of states, paths and likelihood parameters, respectively. The implicit factorisation of this approximate posterior effectively partitions model inversion into inference, planning and learning.

## 2.2 VARIATIONAL FREE ENERGY AND INFERENCE

In variational Bayesian inference (a.k.a., approximate Bayesian inference), model inversion entails the minimisation of variational free energy with respect to the sufficient statistics of approximate posterior beliefs. This can be expressed as follows, where—for clarity—we will deal with a single factor, such that the policy (i.e., combination of paths) becomes the path,  $\pi = u$  and omit dependencies on previous states:

$$\begin{aligned}
Q(s_\tau, u_\tau, a) &= \arg \min_Q F \\
F &= \mathbb{E}_Q [\underbrace{\ln Q(s_\tau, u_\tau, a)}_{\text{posterior}} - \underbrace{\ln P(o_\tau | s_\tau, u_\tau, a)}_{\text{likelihood}} - \underbrace{\ln P(s_\tau, u_\tau, a)}_{\text{prior}}] \\
&= \underbrace{D_{KL}[Q(s_\tau, u_\tau, a) || P(s_\tau, u_\tau, a | o_\tau)]}_{\text{divergence}} - \underbrace{\ln P(o_\tau)}_{\text{evidence}} \\
&= \underbrace{D_{KL}[Q(s_\tau, u_\tau, a) || P(s_\tau, u_\tau, a)]}_{\text{complexity}} - \underbrace{\mathbb{E}_Q[\ln P(o_\tau | s_\tau, u_\tau, a)]}_{\text{accuracy}}
\end{aligned} \tag{1}$$

This equation specifies the (approximate posterior) beliefs  $Q(s_\tau, u_\tau, a)$ —about states, paths and likelihoods—as those beliefs that minimise variational free energy, where variational free energy has been expressed in three equivalent functional forms; each affording a complementary interpretation. Here,  $P(o_\tau, s_\tau, u_\tau, a)$  is the generative model; namely, the probability distribution over causes,  $(s_\tau, u_\tau, a)$  and observable consequences,  $(o_\tau)$  at time  $\tau$ .

Because the (KL) divergences cannot be less than zero, the penultimate equality means that free energy is zero when the (approximate) posterior is the true posterior. At this point, the free energy becomes the negative log evidence for the generative model (Beal, 2003). This means minimising free energy is equivalent to maximising model evidence, which is equivalent to minimising the complexity of accurate explanations for observed outcomes.

Minimising free energy therefore ensures expectations encode posterior beliefs, given observed outcomes. This is inference. Planning emerges under active inference by placing priors over (combinations of) paths to minimise expected free energy (Friston et al., 2015):

$$\begin{aligned}
G(u) &= \mathbb{E}_{Q_u} [\ln Q(s_{\tau+1}, a | u) - \ln Q(s_{\tau+1}, a | o_{\tau+1}, u) - \ln P(o_{\tau+1} | c)] \\
&= - \underbrace{\mathbb{E}_{Q_u} [\ln Q(a | s_{\tau+1}, o_{\tau+1}, u) - \ln Q(a | u)]}_{\text{expected information gain (learning)}} - \\
&\quad \underbrace{\mathbb{E}_{Q_u} [\ln Q(s_{\tau+1} | o_{\tau+1}, u) - \ln Q(s_{\tau+1} | u)]}_{\text{expected information gain (inference)}} - \underbrace{\mathbb{E}_{Q_u} [\ln P(o_{\tau+1} | c)]}_{\text{expected cost}} \\
&= - \underbrace{\mathbb{E}_{Q_u} [D_{KL}[Q(a | s_{\tau+1}, o_{\tau+1}, u) \| Q(a | u)]]}_{\text{novelty}} + \\
&\quad \underbrace{D_{KL}[Q(o_{\tau+1} | u) \| P(o_{\tau+1} | c)]}_{\text{risk}} - \underbrace{\mathbb{E}_{Q_u} [\ln Q(o_{\tau+1} | s_{\tau+1}, u)]}_{\text{ambiguity}}
\end{aligned} \tag{2}$$

Here,  $Q_u = Q(o_{\tau+1}, s_{\tau+1}, a | u) = P(o_{\tau+1}, s_{\tau+1}, a | u, o_0, \dots, o_\tau) = P(o_{\tau+1} | s_{\tau+1}, a)Q(s_{\tau+1}, a | u)$  is the posterior predictive distribution over parameters, hidden states and outcomes at the next time step, under a particular path. Note that the expectation is over *observations in the future* that become random variables; hence, *expected* free energy. This means that preferred outcomes—that subtend expected cost and risk—are prior beliefs, which constrain the implicit planning as inference (Attias, 2003; Botvinick and Toussaint, 2012; Van Dijk and Polani, 2013).

One can also express the prior over the parameters in terms of an expected free energy, where, marginalising over paths:

$$\begin{aligned}
P(a) &= \sigma(-G) & (3) \\
G(a) &= \mathbb{E}_{Q_a} [\ln P(s | a) - \ln P(s | o, a) - \ln P(o | c)] \\
&= \underbrace{\mathbb{E}_{Q_a} [\ln P(s | a) - \ln P(s | o, a)]}_{\text{expected information gain}} - \underbrace{\mathbb{E}_{Q_a} [\ln P(o | c)]}_{\text{expected cost}} \\
&= -\underbrace{\mathbb{E}_{Q_a} [D_{KL}[P(o, s | a) || P(o | a)P(s | a)]]}_{\text{mutual information}} - \underbrace{\mathbb{E}_{Q_a} [\ln P(o | c)]}_{\text{expected cost}}
\end{aligned}$$

where  $Q_a = P(o | s, a)P(s | a) = P(o, s | a)$  is the joint distribution over outcomes and hidden states, encoded by the Dirichlet parameters,  $a$ . Note that the Dirichlet parameters encode the mutual information, in the sense that they implicitly encode the joint distribution over outcomes and their hidden causes. When normalising each column of the  $a$  tensor, we recover the likelihood distribution (as in Figure 1); however, we could normalise over every element, to recover a joint distribution [we will use this later in Equation (9)].

Expected free energy can be regarded as a universal objective function that augments mutual information with expected costs or constraints. Constraints—parameterised by  $c$ —reflect the fact that we are dealing with open, nonequilibrium, systems with characteristic outcomes,  $o$ . This can be read as an expression of the constrained maximum entropy principle that is dual to the free energy principle (Ramstead et al., 2022). Alternatively, it can be read as a constrained principle of maximum mutual information or minimum redundancy (Ay et al., 2008; Barlow, 1961; Linsker, 1990a; Olshausen and Field, 1996b). In machine learning, this kind of objective function underwrites disentanglement (Higgins et al., 2021; Sanchez et al., 2019), and generally leads to sparse representations (Gros, 2009; Olshausen and Field, 1996b; Sakthivadivel, 2022b; Tipping, 2001).

When comparing the expressions for expected free energy in (2) with variational free energy in (1), the expected divergence becomes expected information gain. Expected information gain about the parameters and states are sometimes associated with distinct epistemic affordances; namely, novelty and salience, respectively (Schwartenbeck et al., 2019). Similarly, expected log evidence becomes expected value, where value is the logarithm of prior preferences. The last equality provides a complementary interpretation; in which the expected complexity becomes risk, while expected inaccuracy becomes ambiguity.

There are many special cases of minimising expected free energy. For example, maximising expected information gain maximises (expected) Bayesian surprise (Itti and Baldi, 2009), in accord with the

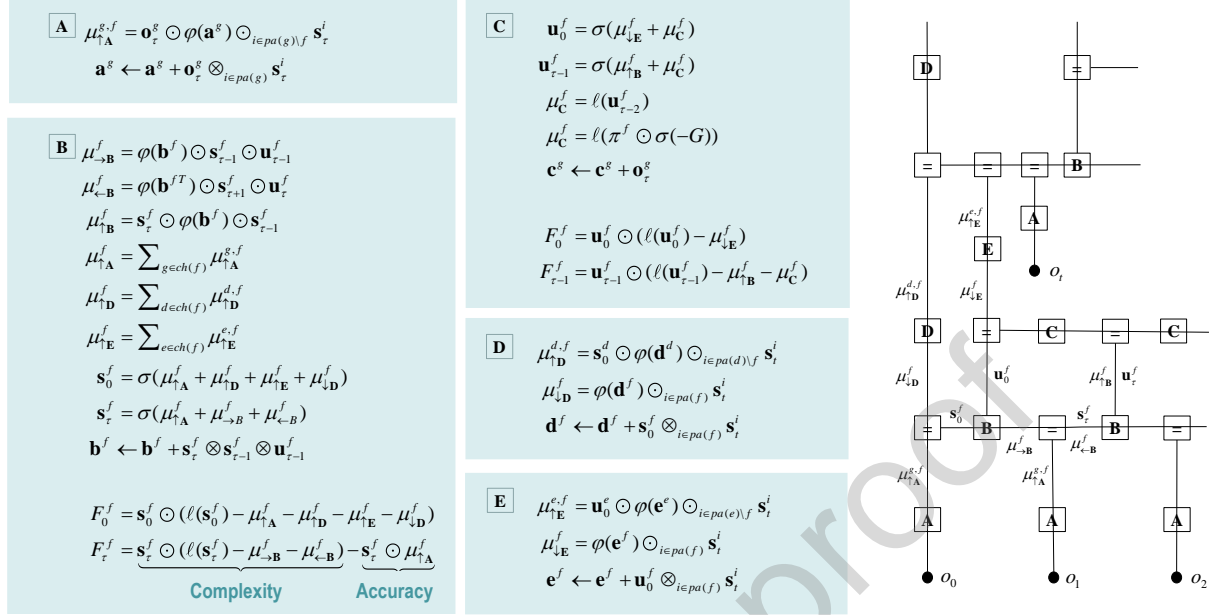
principles of optimal experimental design (Lindley, 1956). This can also be interpreted in terms of the principle of maximum mutual information or minimum redundancy (Barlow, 1961; Laughlin, 2001; Linsker, 1990b; Olshausen and Field, 1996a). This resolution of uncertainty is closely related to artificial curiosity (Schmidhuber, 1991; Still and Precup, 2012) and speaks to the value of information (Howard, 1966), particularly in the context of evincing information necessary to realise preferred outcomes. See (Meder and Nelson, 2012; Nelson et al., 2010), who compare different models of information gain in perceptual decision-making.

Expected complexity or risk is the same quantity minimised in risk sensitive or KL control (Klyubin et al., 2005; van den Broek et al., 2010), and underpins (free energy) formulations of bounded rationality based on complexity costs (Braun et al., 2011; Ortega and Braun, 2013) and related schemes in machine learning; e.g., Bayesian reinforcement learning (Ghavamzadeh et al., 2016). More generally, minimising expected cost subsumes Bayesian decision theory (Berger, 2011).

### 2.3 BELIEF UPDATING

In variational treatments, the sufficient statistics encoding posterior expectations are updated by minimising variational free energy. Figure 2 illustrates these updates in the form of variational message passing (Dauwels, 2007; Friston et al., 2017b; Winn and Bishop, 2005). Although the updates look complicated, they lend themselves to neurobiological implementation—at a certain level of analysis—in a straightforward fashion (Friston et al., 2016; Friston et al., 2014). This is because the updates only require nonlinear mappings and sum-product (tensor) operations.

## Message passing on factor graphs



**Figure 2 – Belief updating and variational message passing:** the right panel presents the generative model as a factor graph, where the nodes (square boxes) correspond to the factors of the generative model (labelled with the associated tensors). The edges connect factors that share dependencies on random variables. The leaves of (filled circles) correspond to known variables, such as observations ( $o$ ). This representation is useful because it scaffolds the message passing—over the edges of the factor graph—that underwrite inference and planning. The functional forms of these messages are shown in the left-hand panels. For example, the expected path—in the first equality of panel C—is a softmax function of two messages. The first is a descending message  $\mu_{\downarrow E}^f$  from E that inherits from expectations about hidden states at the hierarchical level above. The second is the log-likelihood of the path based upon expected free energy. This message depends upon Dirichlet counts scoring preferred outcomes—in modality  $g$ —encoded in  $\mathbf{c}^s$  (see Figure 1). The two expressions for  $\mu_C^f$  correspond to uncontrolled and controlled paths, respectively. Note that the posterior expectation of the initial path is learned, in a context sensitive fashion, by accumulation in  $\mathbf{e}^f$ , which describes how policies can become habits. The updates in the lighter panels correspond to learning; i.e., updating Bayesian beliefs about parameters (adopting the Einstein summation with respect to  $\tau$ ). Similar functional forms for other messages can be derived, by direct calculation. These furnish the fixed points (i.e., free energy minimisers), which render this kind of variational message passing a fixed-point iteration scheme. The requisite functions are defined in Figure 1. The  $\odot$  notation implies a generalised inner product or tensor contraction, while  $\otimes$  denotes an outer product.  $ch(\cdot)$  and  $pa(\cdot)$  return the children and parents of any node; namely, the domain and co-domain, respectively. There is a link (i.e., edge) between a parent and child, if the co-domain of the parent factor (i.e., node) constitutes a domain of a child. This means that each edge—and (the co-domain of) each factor—is uniquely associated with a state, path or outcome.

For example, expectations about hidden states are a softmax function of messages that are linear combinations of other expectations and observations.

$$\begin{aligned}
\mathbf{s}_\tau^f &= \sigma(\mu_{\uparrow\mathbf{A}}^f + \mu_{\rightarrow B}^f + \mu_{\leftarrow B}^f) \\
\mu_{\uparrow\mathbf{A}}^f &= \sum_{g \in \text{ch}(f)} \mu_{\uparrow\mathbf{A}}^{g,f} \\
\mu_{\uparrow\mathbf{A}}^{g,f} &= \mathbf{o}_\tau^g \odot \varphi(\mathbf{a}^g) \odot_{i \in \text{pa}(g) \setminus f} \mathbf{s}_\tau^i
\end{aligned} \tag{4}$$

In this example, conditional expectations about hidden states (of factor  $f$ ) are a normalised exponential of several messages that can be read as log probabilities. Here, these messages correspond to a log likelihood due to observations and log priors from the past and future beliefs about hidden states (that entail dynamics). In turn, the log likelihood message is itself a mixture of ascending messages from those outcome modalities that are the children of the factor in question. The final equality means that each ascending messages is a linear mixture<sup>4</sup> of expected states and observations, weighted by (digamma) functions of the Dirichlet counts, which correspond to the parameters of the likelihood model (c.f., connections weights). In practice, one can use an upper bound for the likelihood messages. From Jensen's inequality (assuming a single parent for clarity)<sup>5</sup>:

$$\begin{aligned}
\mu_{\uparrow\mathbf{A}}^g &= \mathbf{o}_\tau^g \odot \varphi(\mathbf{a}^g) \\
&= \mathbb{E}_{Q(o,a)}[\ln P(\mathbf{A}^g)] \leq \ln \mathbb{E}_{Q(o,a)}[P(\mathbf{A}^g)] \\
&= \ell(\mathbf{o}_\tau^g \odot \mu(\mathbf{a}^g))
\end{aligned} \tag{5}$$

This functional form can exploit the sparsity of likelihood tensors to finesse the von Neumann bottleneck. From a biological perspective, this just means that there is no contribution to a message if the connection is absent or has been removed via structure learning. In fact, these are the messages that are used in belief propagation, during online belief updating (that eschews backward message passing or variational approximations) (Kschischang et al., 2001; Parr et al., 2019).

---

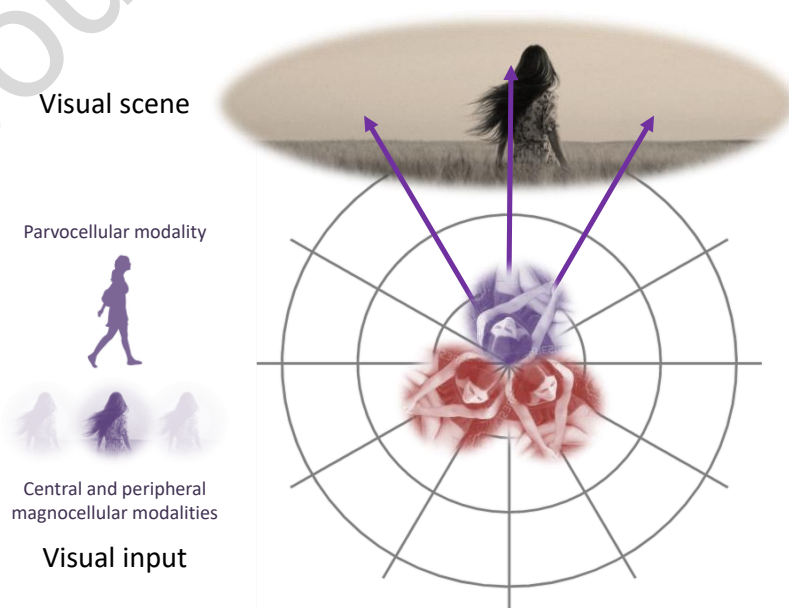
<sup>4</sup> The  $\odot$  notation implies a sum product operator; i.e., the dot or inner product that sums over one dimension of a numeric array or tensor. In this paper, these sum product operators are applied to a vector  $\mathbf{a}$  and a tensor  $\mathbf{A}$  where,  $\mathbf{a} \odot \mathbf{A}$  implies the sum of products is taken over the leading dimension, while  $\mathbf{A} \odot \mathbf{a}$  implies the sum is taken over a trailing dimension. For example,  $\mathbf{1} \odot \mathbf{A}$  is the sum over columns and  $\mathbf{A} \odot \mathbf{1}$  is the sum over rows, where  $\mathbf{1}$  is a vector of ones. This notation replaces the Einstein summation notation to avoid visual clutter.

<sup>5</sup> For deterministic outcomes one can read  $Q(o, a) = \delta(o)Q(a)$ , where the Kronecker delta function,  $\delta(o)$  is a one-hot vector encoding the observation in question. Otherwise, for outcomes that are provided probabilistically, we have  $Q(o, a) = P(o)Q(a) : P(o) = \text{Cat}(\mathbf{o})$ .

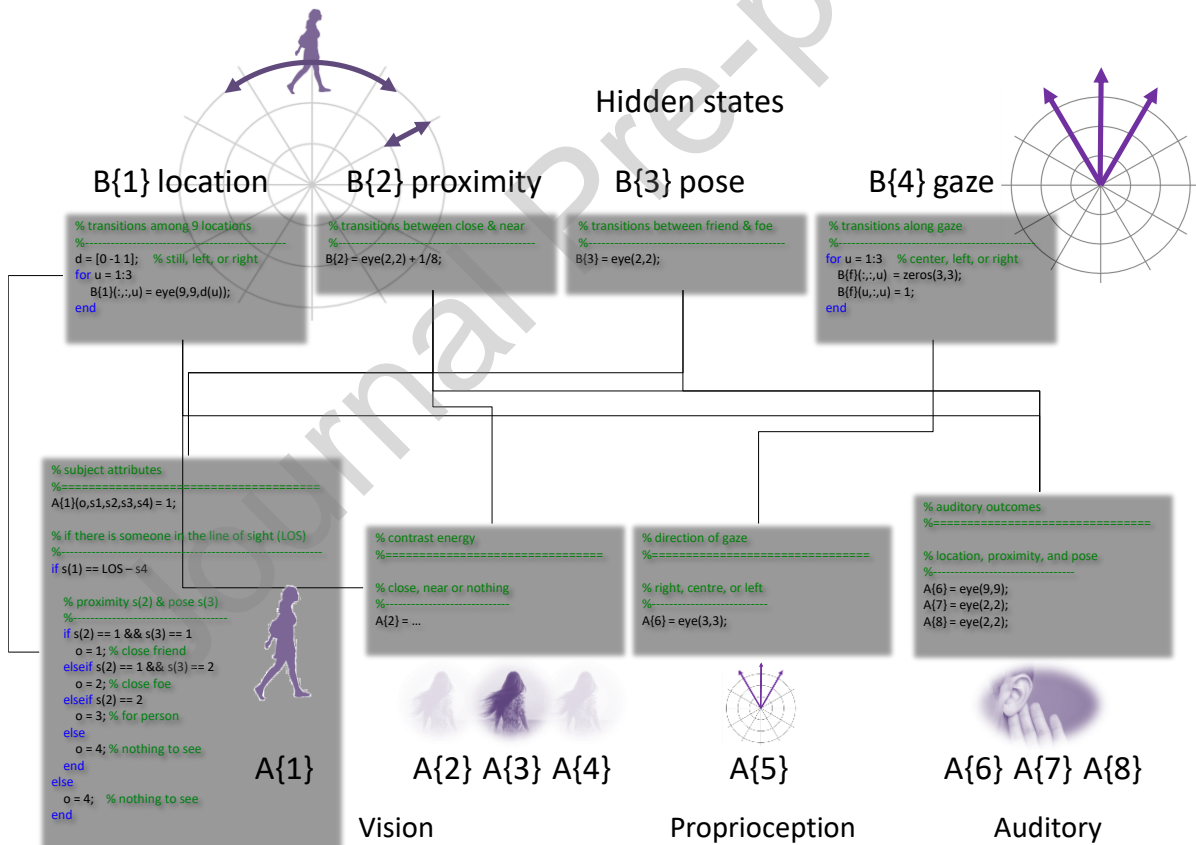
### 3 ACTIVE INFERENCE AND BELIEF SHARING

The setup used to illustrate belief-sharing considers three agents that can be thought of as birds (i.e., sentinels) maintaining surveillance for a predator—or sisters hiding from their mother in the garden. Crucially, each agent has a restricted field of view, covering about a third of the horizon. In addition to visual and proprioceptive (gaze-related) observation modalities, each agent can also hear the others (but not herself). To realise belief-sharing, agents have a simple (identity) likelihood mapping between posterior beliefs—about latent states that are shareable among agents—and output modalities or communication channels (e.g., talking). In subsequent sections we will consider the emergence of the requisite likelihood mappings. In this section, we will focus on the mechanisms and benefits of communication afforded by belief sharing among three agents, who have complementary views of the world.

Here, agents report the location of a subject (e.g., potential predator or mother) in terms of its radial location in an allocentric frame of reference and its proximity (*near* or *far*). In addition, the agents can report whether the subject's disposition is friendly or not. The requisite inference is difficult (in addition to the limited field of view afforded each agent): first, the agents cannot see motion. This means any movement of the subject has to be inferred from successive inferences about their location. Inferred motion is crucial because it underwrites predictions of location at the next time step, which are broadcast to other agents. Second, the disambiguation between friend and foe is only possible when the subject is *close* to an agent. Figure 3 describes this setup in terms of what each agent can see, while Figure 4 describes the hidden states generating observations.



**Figure 3 – The agents' worldview.** This graphic illustrates how the world generates (visual) observations for our agents. In this setup, there are three agents, sitting back-to-back, who survey a scene in which a subject (e.g., predator or person) is nearby. The subject can be standing *still* or encircling the agents in a clockwise (*right*) or anticlockwise (*left*) direction; either *close* or *near* to the agents. When the subject is *close*, agents can discern whether it is a *friend* or a *foe* (e.g., conspecific or a predator). The requisite visual observations have four modalities: a central foveal input—we associate with the parvocellular stream—that classifies the subject as either a proximate *friend* or *foe* or *someone* in the distance (c.f., the product of an image classification algorithm). The remaining three (magnocellular) modalities are low spatial frequency reports of contrast energy, in central and peripheral regions of the visual field (c.f., LIDAR detectors). Each contrast energy has three levels; *close*, *near* or *none*). Each agent can only see what is in front of her but can look to the *right* or *left*—to extend her field of view (purple arrows). In the example of visual input above, the purple agent is looking straight ahead at someone, and her foveal (parvocellular) observation is someone in the distance. The remaining visual modalities (lower circles) register *near* contrast energy in the centre of the visual field with *none* on the right and left. The latent states of the environment generating these observations are described in the subsequent figure.



**Figure 4 – states of affairs in the world.** This figure summarises the hidden states generating outcomes in terms of factors and their accompanying levels. It is presented in a way that will be familiar to people specifying generative models for active inference; specifically, using (Matlab) pseudocode to illustrate how one specifies likelihood mappings and transition priors in terms of simple matrices and logical expressions. In this generative model, there are four factors corresponding to the allocentric position (*location*) of the subject, whether the subject



is close or near (*proximity*), whether the subject is friendly or not (*pose*) and where the agent is currently looking (*gaze*). The transitions among the states of these four factors are encoded by prior  $\mathbf{B}$  tensors. These tensors are probability transition matrices with multiple slices; where each slice corresponds to a particular path, trajectory or action. Only the location and gaze factors have multiple paths, modelling trajectories taken by the subject (standing still, walking to the left or right) and the agent (looking straight-ahead or to the left or right). Crucially, only the *gaze* factor is controllable; in other words, the agent can choose where to look based upon her prior beliefs about policies, supplied by their respective expected free energies. The factors generate a discrete outcome in each modality; where each modality is equipped with a likelihood  $\mathbf{A}$  tensor, whose dimensions correspond to the number of outcomes in each modality times the number of states in each factor. Here, there are eight modalities. The four visual modalities are described in Figure 3. The remaining four modalities comprise a proprioceptive modality that reports the current direction of gaze, and three auditory modalities with identity mappings from the hidden states that are shared by all agents; namely, *location*, *proximity* and *pose*. The  $\mathbf{A}$  and  $\mathbf{B}$  tensors collectively specify the transition priors and likelihoods of the generative model. In hierarchical models there are further tensors linking states at higher levels to the initial states and paths of lower levels. Here, we set these to be uninformative vectors; i.e., no hierarchical constraints. The pseudocode in the panels illustrates the way that these tensors can be specified in terms of straightforward (Boolean) logic.

### 3.1 BELIEF SHARING

In simulations of belief-sharing or communication, there is a crucial difference between the way outcomes are generated by the environment and by other agents. Outcomes generated by the environment are caused by hidden states, which are the same for all agents. However, outcomes that underwrite belief-sharing are caused by an agent's *beliefs* or predictions about hidden states. If belief sharing were implemented by message passing on a factor graph, it would just involve the exchange of (log) posteriors among agents that share beliefs about factor  $f$ . Expressing this in terms of message passing, just entails supplementing likelihood and prior messages with the corresponding messages from other agents. For agent  $n$ , this translates into:

$$\mathbf{s}_\tau^{n,f} = \sigma(\mu_{\uparrow\mathbf{A}}^{n,f} + \mu_{\rightarrow\mathbf{B}}^{n,f} + \mu_{\leftarrow\mathbf{B}}^{n,f} + \sum_{m \in pa(n)} (\mu_{\uparrow\mathbf{A}}^{m,f} + \mu_{\rightarrow\mathbf{B}}^{m,f} + \mu_{\leftarrow\mathbf{B}}^{m,f})) \quad (6)$$

Federated inference of this sort ensures that agents come to share posterior beliefs that minimise the joint free energy over agents (at which point free energy gradients vanish):

$$\begin{aligned} \frac{\partial F}{\partial \mathbf{s}_\tau^{n,f}} &= \frac{\partial}{\partial \mathbf{s}_\tau^{n,f}} [\mathbf{s}_\tau^{n,f} \odot \ell(\mathbf{s}_\tau^{n,f}) - \mathbf{s}_\tau^{n,f} \odot \sum_m (\mu_{\uparrow\mathbf{A}}^{m,f} + \mu_{\rightarrow\mathbf{B}}^{m,f} + \mu_{\leftarrow\mathbf{B}}^{m,f})] = 0 \\ &\Rightarrow \\ \ell(\mathbf{s}_\tau^{n,f}) &= \sum_m (\mu_{\uparrow\mathbf{A}}^{m,f} + \mu_{\rightarrow\mathbf{B}}^{m,f} + \mu_{\leftarrow\mathbf{B}}^{m,f}) \Rightarrow \\ \mathbf{s}_\tau^{n,f} &= \sigma(\mu_{\uparrow\mathbf{A}}^{n,f} + \mu_{\rightarrow\mathbf{B}}^{n,f} + \mu_{\leftarrow\mathbf{B}}^{n,f} + \sum_{m \in pa(n)} (\mu_{\uparrow\mathbf{A}}^{m,f} + \mu_{\rightarrow\mathbf{B}}^{m,f} + \mu_{\leftarrow\mathbf{B}}^{m,f})) \end{aligned} \quad (7)$$

This is formally distinct from Bayesian belief updating, in which the likelihood of conditionally independent observations are assimilated. *Belief updating* would result in agents with different posterior beliefs, depending upon their priors. *Belief sharing* assimilates posteriors to evince a consensus that can

be likened, anecdotally, to a ‘hive mind’; in which agents inherit both the likelihood *and priors* from other agents. Note that beliefs are only shared when they pertain to the same states of the world, under a shared frame of reference (Fields et al., 2022). This means that only some beliefs are shared (e.g., about the location of a subject in the environment), while others are not (e.g., where each agent is looking).

The direct (peer-to-peer) belief-sharing in (6) may be apt for federated inference, where federated inference can be read as the assimilation of messages from multiple agents during inference or belief updating. However, federated inference—so defined—dissolves the notion of an individual agent. This follows because each agent is defined by its Markov blanket, which precludes reciprocal message passing with other agents (Heins et al., 2023; Palacios et al., 2017; Parr et al., 2020; Pellet and Elisseeff, 2008). In other words, for one agent to be individuated from another agent, they have to be separated by a Markov blanket. This means one agent cannot be a parent (or child) of another; that is, they cannot pass messages to each other. To preserve the conditional independence of agents they have to communicate through a shared Markov blanket; namely, their observations. In sum, to individuate one agent from another requires a (likelihood) mapping from beliefs to exchangeable observations. We will associate the implicit exchange of sufficient statistics with *communication* in general, and language in particular (Isomura et al., 2019).

By equipping each agent with a likelihood mapping from her beliefs to observable outcomes, one effectively creates agents that broadcast their beliefs and recognise what other agents believe. Endowing each agent with an outcome modality  $\mathbf{o}_\tau^f$ —for each shareable factor—we have, from (5):

$$\begin{aligned}
\mu_{\uparrow\mathbf{A}}^{n,f} &= \ell(\mathbf{o}_\tau^{n,f} \odot \mu(\mathbf{a}^f)) \\
\ell(\mathbf{o}_\tau^{n,f}) &= \sum_{m \in pa(n)} \ell(\mu(\mathbf{a}^f) \odot \mathbf{s}_\tau^{m,f}) \\
&\Rightarrow \\
\mu_{\uparrow\mathbf{A}}^{n,f} &= \sum_{m \in pa(n)} \ell(\mathbf{s}_\tau^{m,f}) : \mu(\mathbf{a}^f) = \mathbf{I}^f \\
&= \sum_{m \in pa(n)} (\mu_{\uparrow\mathbf{A}}^{m,f} + \mu_{\rightarrow\mathbf{B}}^{m,f} + \mu_{\leftarrow\mathbf{B}}^{m,f})
\end{aligned} \tag{8}$$

Comparison with (6) shows that the ascending likelihood message is formally identical to the log posteriors required for belief-sharing. This holds if the likelihood mappings are identity matrices—and each agent ‘hears’ a probabilistic mixture of outputs generated by other agents. In fact, the likelihood mappings can be any permutation matrix, as we will see later. Notice that the likelihood tensors for broadcasting beliefs are just matrices that uniquely associate an outcome modality with a hidden factor.

In what follows, communication was modelled according to (8), generating a probabilistic outcome (e.g., chorus) heard by all agents<sup>6</sup>. To keep things simple, we assume that each kind of belief is shared in a separate (e.g., auditory) modality; e.g., ‘word’ or ‘call’. In this setting, each ‘word’ holds more information than conventional tokens because it contains the sufficient statistics over possible outcomes for this kind of word (e.g., encoded in the amplitude of different frequencies of auditory streams). In other words, communication is taken to include not just the content of beliefs but the confidence in those beliefs (Bahrami et al., 2010). A more elaborate simulation of linguistic communication would linearise a word sequence, by equipping the generative model with a lower hierarchical level: e.g., (Friston et al., 2020; Friston et al., 2021). Another approach would be to sample a token in proportion to the (negative) expected free energy of commutative actions (Albarracín et al., 2022; Heins et al., 2023). However, for simplicity, we will deal with the direct exchange of sufficient statistics, necessary for communicating beliefs.

In this setup, an agent’s communication will be imprecise or quiet, if the agent is uncertain about the state in question. Conversely, if the agent has precise beliefs, her contribution will dominate in a way that can be read as speaking loudly and clearly. Note that agents do not hear themselves. Technically, this precludes double-counting of an agent’s log posterior: c.f., (Jardri and Deneve, 2013). Neurobiologically, this can be viewed as sensory attenuation; namely, attenuating self-generated sensations (Blakemore et al., 1999; Limanowski, 2017). There are many other models of communication that involve turn-taking and attribution of agency (Friston and Frith, 2015; Ghazanfar and Takahashi, 2014; Wilson and Wilson, 2005): however here, we imagine that the sentinels are constantly reporting their beliefs to each other. This has the particular consequence that the communicated beliefs are posterior predictive distributions over latent states. In other words, the environment generates observations and agents broadcast their beliefs at the same time. This means that the beliefs are based upon observations at the previous time step, and are therefore predictive in nature. In turn, the predictive validity of communicated beliefs rests sensitively on inferred state transitions (e.g., motion of the subject of observation). We will see examples of this below.

---

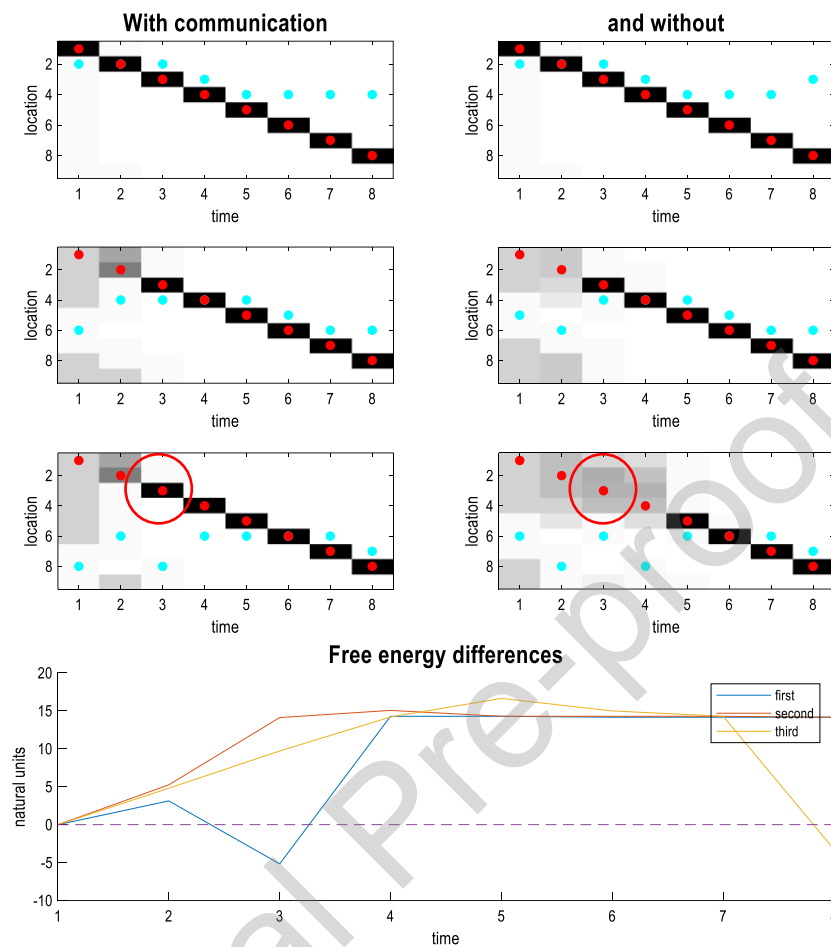
<sup>6</sup> In other words, instead of supplying observations as a one-hot vector, auditory observations are supplied as probability distributions over the discrete outcomes of each modality; i.e., the sufficient statistics of categorical distributions.

### 3.2 BELIEF SHARING AND COMMUNICATION

In what follows, we compare belief updating with and without communication, where communication is suppressed by reducing the precision of the auditory likelihood mappings to zero. This means that the agents can neither generate or recognise auditory cues and are effectively rendered incommunicado.

Figure 5 illustrates the belief updating and ensuing action for three agents, with and without communication (the left and right panels, respectively). In both conditions, the subject started at the first location and moved clockwise around the agents for eight timesteps (about two seconds of real and simulated time). The most prescient difference, when suppressing communication, is a failure of the second and third agents (second and third rows) to resolve uncertainty about the subject's location—and act in a suitably pre-emptive fashion. For example, the third agent only forms precise beliefs about location on time step five. In contrast, the third communicating agent quickly infers the location of the subject on the basis of what the first agent says. The first agent sees the subject at the beginning of the episode and broadcasts her precise predictions on the first time step, to which the remaining agents are able to commit; despite not seeing the subject until the third and fifth times steps, respectively.

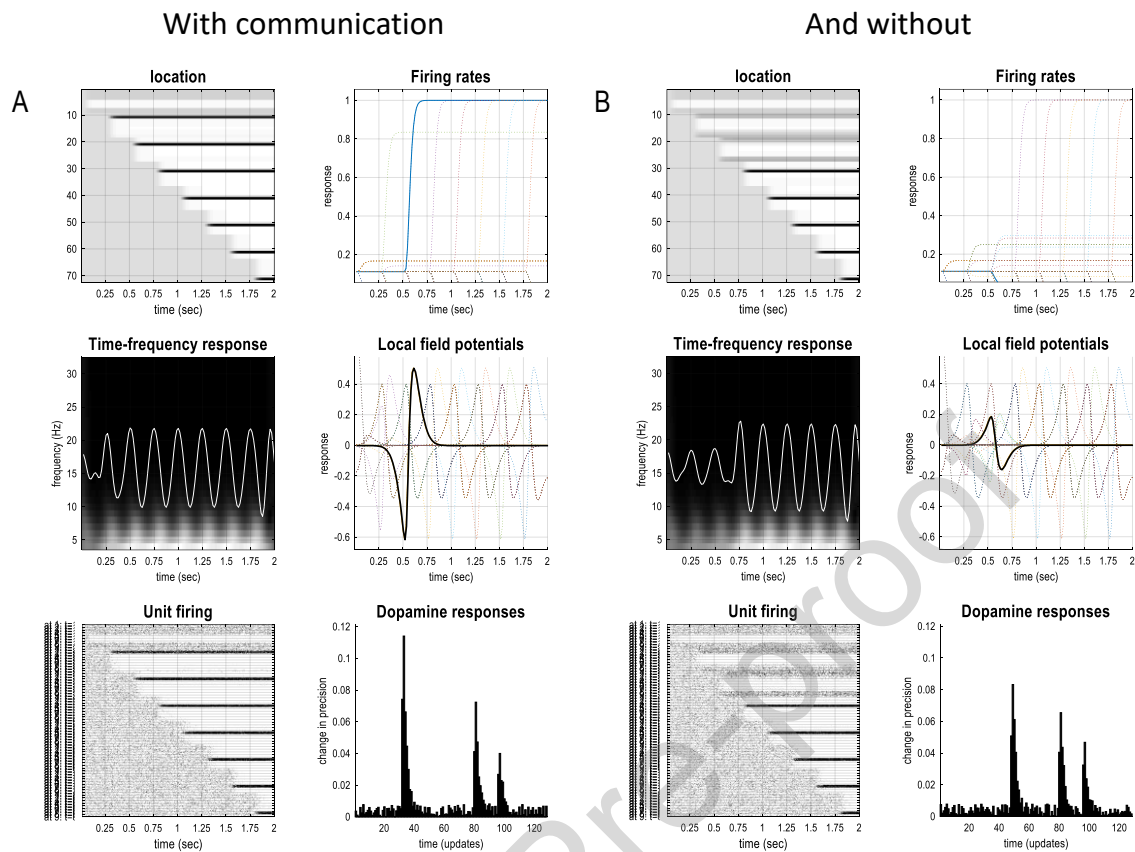
Note that each agent tries to see the subject by looking towards its inferred location, to the extent that they can, given their limited field of view. This visual tracking is driven purely by expected information gain; namely, the imperative to resolve uncertainty by responding to the epistemic affordances scored by expected free energy: i.e., the information gain in (2). The lower panel of Figure 5 shows the differences in free energies when comparing the agents with and without communication. With one exception, these differences suggest that communicating agents have a better grip on the world, and a lower free energy. This is most marked for the third agent (yellow line) who only sees the agent in the periphery of her vision on time step five. Interestingly, the first agent appears to have been a little distracted by the second and third agents, with a slightly higher free energy on the third time step.



**Figure 5 – Simulations of active inference.** The upper panels depict posterior beliefs—about the location of the subject—as posterior distributions over the nine possible locations (columns) at each of the (eight) timesteps. The right-hand panels reproduce the belief updating shown in the left hand panels but in the absence of communication. Each panel reports the posterior beliefs as a function of time along the x-axis. The beliefs in this case are probability distributions over nine locations, shown as a greyscale image. In other words, uncertain, imprecise beliefs correspond to grey regions that cover multiple locations. Over time, uncertainty is resolved, and beliefs becomes very precise with a concentration of probability in the black regions. The true location is shown by the red dots, while the cyan dots indicate where the agent was looking. Note that these locations are restricted to the agent’s line of sight—and one location to the right and left. The three rows of panels report the belief updating for the three agents. The left panels show belief updating with communication. The three panels on the right show the corresponding active inference in the absence of communication, using the same sequence of hidden states. Communication or language was suppressed by rendering the auditory likelihood mappings very imprecise. The lower panel shows differences in free energy as a function of time for the three agents (coloured lines), when comparing belief updating with and without communication. The red circle highlights the epoch of belief updating illustrated in the next figure. This epoch illustrates the benefits of communication in the sense that the third agent has resolved her uncertainty about the location of the subject by the third epoch, whereas the same agent without communication only forms precise beliefs, after seeing the subject at the fifth epoch.

Figure 6 shows the simulated neuronal responses that would accompany the belief updating illustrated in Figure 5. These synthetic results are shown to underscore the fact that the variational message passing described above can be implemented in a neuronally plausible fashion (Friston et al., 2017a; Friston et al., 2017b; Parr and Friston, 2018), producing simulated electrophysiological responses that are similar to those observed empirically. In this example, the key takeaway is the greater degree of belief updating when an agent can assimilate informative communications (left panels), relative to when she has an imprecise auditory likelihood mapping (right panels). An informative likelihood mapping enables auditory input to resolve uncertainty about latent states, producing a greater degree of belief updating and simulated electrophysiological responses.

The accompanying resolution of uncertainty is scored by simulated dopaminergic discharges in the lower right panels that accompany an augmented event-related potential (middle panels). Please see (Friston et al., 2017a; Friston et al., 2014; Friston et al., 2017b; Parr and Friston, 2018) for a fuller discussion. This biomimetic formulation of variational message passing replaces the fixed point iteration scheme in Figure 3 with a gradient descent on variational free energy. Neuronal implementation is not a central part of the current argument; however, it speaks to the sort of predictions that can be made, when testing process theories based upon active inference.



**Figure 6 – Simulated electrophysiological responses.** This figure shows expectations about hidden states for the third agent with (A) and without (B) communication. A: The upper left panel shows the activity (firing rate) of units encoding the location of a subject over time points (i.e., epochs), each corresponding to roughly 250 ms of simulated and computational time. These responses are organised such that the upper rows encode the probability of alternative states in the first epoch, with subsequent epochs in lower rows. In other words, the top row shows the expectations about hidden states at the beginning of the exchange—and how these expectations evolve over time. Conversely, the first column shows expectations about the future. The plot to the right of the image presents the same information to illustrate the implicit evidence accumulation. These values can be interpreted as neuronal firing rates of units encoding expectations. The associated local field potentials (i.e. the rate of change of neuronal firing) are shown in the middle panels in terms of a time frequency decomposition (left middle) and local field potentials (right middle). The local field potential showing the greatest difference between the two conditions is highlighted in black. The lower left panel reproduces the upper left panel but in the form of a raster plot, assuming the expectations have a Poisson rate code. Similarly, the lower right panel shows the simulated dopaminergic responses under a Poisson rate code assumption. Dopamine, in this formulation, scores changes in the confidence afforded policy selection; namely, the negative entropy of posterior beliefs over combinations of actions based upon expected free energy. In this example, the agent is more confident about her action (i.e., where to look) when she is able to communicate with other agents. This is manifest as a short latency, high amplitude, phasic dopaminergic response.

In summary, this section has showcased the role of belief-sharing through communication in resolving uncertainty during active inference. In effect, it enables communicating agents to benefit from the complementary perspectives and observations afforded each agent (Bahrami et al., 2010). In the next

section, we turn to active learning and the acquisition of the likelihood mapping that underwrites the generation and recognition of communicative exchanges.

## 4 ACTIVE INFERENCE AND LEARNING

---

"The second question is why, out of the infinite range of knowable items in the universe, certain pieces of knowledge are more ardently sought and more readily retained than others" (Berlyne, 1954) p180.

In this section, we turn to active learning and the acquisition of an apt likelihood model for communication. Active learning has a special meaning in this context. It implies that the action that updates the Dirichlet counts (c.f., experience-dependent plasticity) is selected on the basis of expected free energy; where expected free energy—from the perspective of model parameters—is, effectively, the mutual information encoded by the Dirichlet tensors: see Equation (3). Put simply, this means that an update—to knowledge encoded in likelihood tensors—is only selected in proportion to the expected information gain. Consider two policies: to update or not to update. From Figure 2, we have (dropping the modality superscript for clarity):

$$\Delta \mathbf{a} = \mathbf{o}_\tau \otimes_{i \in pa} \mathbf{s}_\tau^i \quad (9)$$

$$\mathbb{E}_Q[a | u = u_o] = \mathbf{a} | u_o = \mathbf{a}$$

$$\mathbb{E}_Q[a | u = u_1] = \mathbf{a} | u_1 = \mathbf{a} + \Delta \mathbf{a}$$

$$\bar{a} =: \frac{a}{\Sigma(a)} = \mathbb{E}[P(o, s | a)]$$

Here,  $\Sigma(a) =: \mathbf{1} \odot a \odot_{i \in pa} \mathbf{1}$  is just sum of all tensor elements. The prior probability of committing to an update is given by the expected free energy of the respective Dirichlet parameters, which scores the expected information or gain (i.e., mutual information) and cost<sup>7</sup>:

---

<sup>7</sup> For simplicity, we assume tensors have been formatted as matrices, by vectorising the second and subsequent dimensions.



$$\begin{aligned}
P(u) &= \sigma(-\alpha \cdot G(\mathbf{a} | u)) \\
G(\mathbf{a}) &= -\underbrace{\mathbb{E}_{Q_a} [D_{KL}[Q(s, o | \mathbf{a}) \| Q(s | \mathbf{a})P(o | \mathbf{a})]}]_{\text{mutual information}} - \underbrace{\mathbb{E}_{Q_a} [\ln P(o | c)]}_{\text{expected cost}} \\
&= (\mathbf{1} \odot \bar{\mathbf{a}}) \odot \ell(\mathbf{1} \odot \bar{\mathbf{a}}) + (\bar{\mathbf{a}} \odot \mathbf{1}) \odot \ell(\bar{\mathbf{a}} \odot \mathbf{1}) - \mathbf{1} \odot (\bar{\mathbf{a}} \times \ell(\bar{\mathbf{a}})) \odot \mathbf{1} - \varphi(\mathbf{c}) \odot (\bar{\mathbf{a}} \odot \mathbf{1})
\end{aligned} \tag{10}$$

This prior over update policies furnishes a Bayesian model average of the likelihood parameters, effectively marginalising over update policies:

$$\begin{aligned}
\mathbb{E}_Q[a_{\tau+1}^g] &= P(u_0) \cdot \mathbb{E}_Q[a_\tau^g | u = u_0] + P(u_1) \mathbb{E}_Q[a_\tau^g | u = u_1] \\
&\Rightarrow \\
\mathbf{a}_{\tau+1}^g &= P(u_0) \cdot \mathbf{a}_\tau^g + P(u_1)(\mathbf{a}_\tau^g + \Delta \mathbf{a}_\tau^g) \\
&= \mathbf{a}_\tau^g + P(u_1) \Delta \mathbf{a}_\tau^g
\end{aligned} \tag{11}$$

In equation (10),  $\alpha$  plays the role of a hyperprior that determines the sensitivity to expected free energy. When this precision parameter is large, the Bayesian model average above becomes Bayesian model selection; i.e., either the update is selected, or it is not. It may seem odd to impose constraints on updates in this fashion; however, active inference rests on a circular causality, in which actions on the world realise predicted outcomes. This means that committing to a world model, with precise likelihood mappings, can bring about the precise generation of outcomes. Communication—and implicit niche co-construction—is a nice example of this, as we will see later.

The active learning in (11) would be Bayes optimal in a world that never changed; enabling the eternal accumulation of Dirichlet counts, such that more and more evidence would be required to change the expected likelihoods. However, when the world is changing (e.g., through the action of agents that themselves are learning) there exists a particular timescale over which evidence should be retained. One can accommodate this by introducing a hyperprior on the effective number of observations that should be retained as follows; noting that the total Dirichlet counts—that could be accumulated after each observation—sum to one, by construction:

$$\begin{aligned}
\mathbf{a}_{\tau+1}^g &= (\mathbf{a}_\tau^g + P(u_1) \Delta \mathbf{a}_\tau^g) \left( \frac{\eta}{\eta + P(u_1)} \right), \quad \Sigma(\Delta \mathbf{a}_\tau^g) = 1 \\
&\Rightarrow \\
\eta(\Sigma(\mathbf{a}_{\tau+1}^g) - \Sigma(\mathbf{a}_\tau^g)) &= P(u_1)(\eta - \Sigma(\mathbf{a}_{\tau+1}^g)) \\
\Sigma(\mathbf{a}_{\tau+1}^g) &= \eta \Rightarrow \Sigma(\mathbf{a}_{\tau+1}^g) = \Sigma(\mathbf{a}_\tau^g)
\end{aligned} \tag{12}$$

Equation (12) says the total number of Dirichlet counts saturates at  $\eta$ . In other words, the Dirichlet counts acquire an upper bound, via a slight decay at the point of updating. Technically, this can be

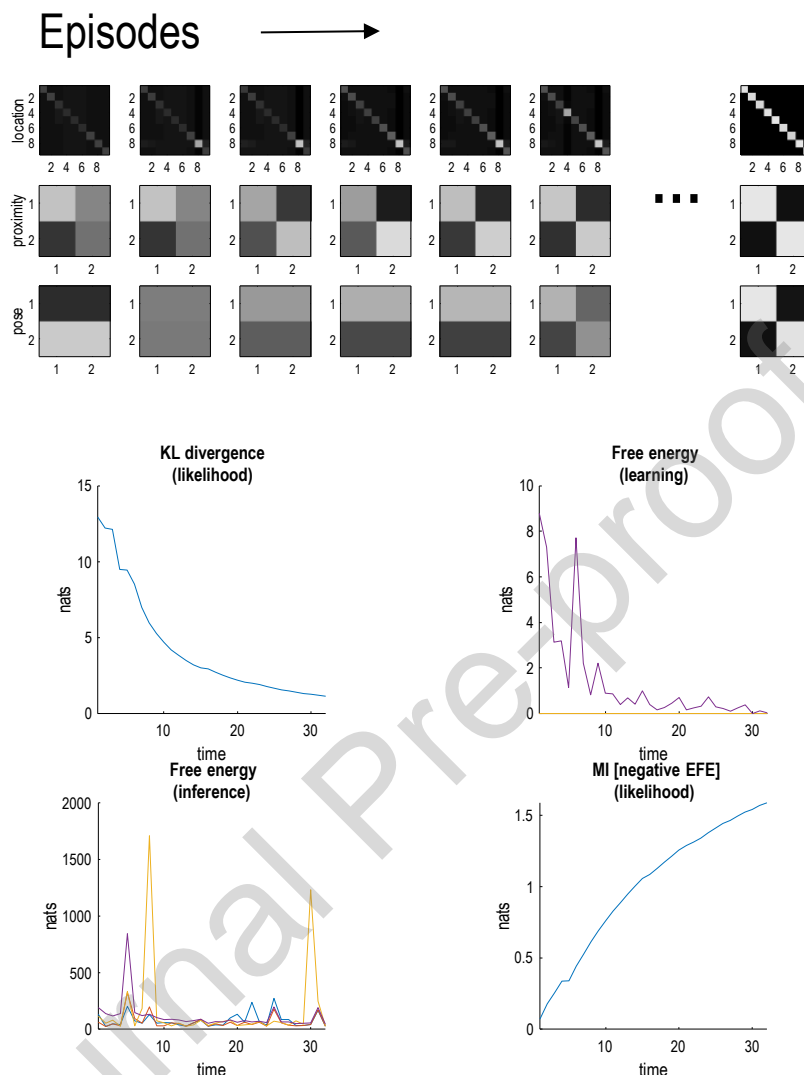
regarded as a hyperprior that implements Bayes-optimal forgetting in a volatile environment (Ishii et al., 2002; Moens and Zenon, 2019). On this view,  $\tau$  is the timescale over which an adiabatic approximation holds. Neurobiologically, this is not unrelated to reconsolidation of memories (Stickgold and Walker, 2007) and synaptic homeostasis (Huber et al., 2004; Toutounji and Pipa, 2014).

Heuristically, the hyperprior determines how impressionable the agent is; in the sense that if the number of Dirichlet counts is small, new observations will have a greater effect on the expected likelihood (because their relative values change more readily). In our simulations, all the agents were effectively young and impressionable with  $\tau$  set to 32. In other words, the last 32 experiences predominate in the accumulation of knowledge or memory encoded in the Dirichlet parameters. Notice how learning has become a key part of inference, with its own dynamics. This reflects the fact that—in active inference—there are probability distributions over all random variables; including hidden states, model parameters and their structure. This enables the application of variational Bayes to simulate or realise (Bayes-optimal) behaviour.

#### 4.1 SYNTHETIC LEARNING

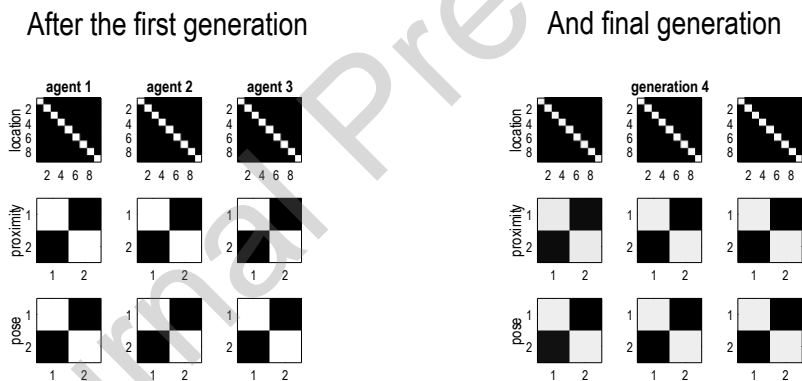
To illustrate active learning in the current setting, we simulated the acquisition of language by a child of one agent, who sees and hears the same things as her parent but has a completely ambiguous or imprecise likelihood mapping. This means the child can neither understand, nor participate, in communication, until she has learned a sufficiently precise auditory mapping. This naïve agent can be thought of as a child in virtue of the fact that her prior Dirichlet counts were small (uniformly one), meaning that she is inherently impressionable.

Figure 7 shows the acquisition of language by the child of the first agent. The upper panels show the emergence of a precise likelihood mapping, as Dirichlet parameters are accumulated by seeing and hearing the same thing as her parent. In this example, there were 32 exposures or episodes—of 16 epochs—by which time the learning of a precise language mapping was almost complete—as reflected in the small KL divergence between the child and her parent, and the high mutual information encoded by the respective auditory likelihood mappings (lower panels in Figure 7). Note that some episodes were informative, and others were not—as reflected in the episodic variations in free energy associated with inference (lower left panel). This reflects the fact that in some trials there was no information about pose (i.e., *friend* or *foe*), unless the subject approached one of the agents. These results illustrate language acquisition and pave the way for a simple simulation of how language can be transmitted over generations.



**Figure 7 – language acquisition.** this figure illustrates the acquisition of language in terms of the (three) auditory likelihood mappings, starting from a completely naïve or uniform tensor (with all prior Dirichlet counts set to 1) to a structured identity matrix, after a sufficient number of exposures; here, 32 episodes. Each episode lasted for 16 timesteps, where each time step corresponds to 1/4 of the second (250 ms). The upper panels show the likelihood mappings for the three auditory modalities (reporting *location*, *proximity* and *pose*) over the first six and last exposures. The lower panels report the subsequent learning in terms of various free energies and relative entropies (a.k.a., KL divergences). The upper right panel shows the KL divergence of the likelihood mapping (summed over hidden states and outcome modalities) between the child and her parent. It can be seen that this falls systematically during language acquisition. The lower left panel shows the accompanying free energy of inference (i.e., due to posterior beliefs) for the three agents over time. The instances of high free energy correspond to episodes in which there is very little (and possibly no information) about one or more latent states for one or more agents. The upper right panel scores the free energy associated with learning. This is the information gain as measured by the KL divergence between the likelihood mappings before and after parametric belief updates. It can be seen that this is greatest during early periods of exposure and then converges to zero as precise likelihood mappings are acquired. The lower right panel shows the increase in the mutual information of the likelihood mappings as a function of acquisition time. This is equivalent to the negative expected free energy due to the model (i.e., Dirichlet) parameters, noting that in these simulations the expected cost was zero.

After 32 episodes, the child replaced her parent and the next parent in line was given a new child. This process was repeated until all parents had been replaced by children. The auditory likelihood mappings encoding the generation and recognition of language are shown in Figure 8. With the exception of slightly imprecise auditory mappings for proximity and pose, these mappings are almost identical to those of the parents. In other words, after four generations, we end up where we started. The remarkable thing here is that language acquisition was mediated entirely through experience-dependent plasticity and active learning. At no point were the prior Dirichlet parameters or structure of any child informed about the language used by her parents. From an evolutionary perspective, cultural transmission—mediated by experience-dependent plasticity of a (neuro) developmental sort—could be read in terms of niche construction and evolutionary developmental biology (Ghazanfar and Takahashi, 2014; Hauser et al., 2002; Heyes, 2018; Laland et al., 1999; Lehmann, 2008; Vasil et al., 2020; Veissiere et al., 2019). Given that the simulations only covered about two minutes of simulated (and real) time, one could regard them as simulating the learning of a new videogame.



**Figure 8 – transgenerational transmission.** This figure shows the likelihood mappings pertaining to language or auditory modalities for three parents at the beginning of the first generation (left panel) and after all the parents have been replaced by their children in the fourth generation (right panel). Each column corresponds to an agent and each row corresponds to a likelihood mapping (shown in image format) from hidden states 23 outcome modalities (location, proximity and pose). These likelihood mappings correspond to the final mappings—after learning—shown in Figure 7 (on the far right). The key thing to notice here is that these likelihood mappings are virtually indistinguishable; despite the fact that subsequent generations had to learn these mappings from scratch.

In the previous section, we saw that precise likelihood mappings were necessary for communication. In this section, we saw that agents acquire precise likelihood mappings, because this is what they expect to acquire *a priori*. This begs the question: does communication emerge from prior expectations? The next section addresses this question by starting with three agents devoid of any language capabilities.

## 5 ACTIVE INFERENCE AND SELECTION

---

In this section, we turn to the emergence of language as a consequence of nested, free-energy minimising processes. In these numerical experiments, we rendered all the auditory likelihood mappings imprecise (and learnable) by setting all the Dirichlet parameters to one, plus an unsigned random Gaussian variate. We then exposed the three agents to 512 episodes to see if language mappings—and attending communication—would emerge.

The results are shown in Figure 9 and speak to a fairly rapid and precise emergence of a shared language that, interestingly, was distinct from the language of the agents in the preceding simulations. In other words, the meaning of various ‘words’ was completely different, such that one of the agents in Figure 9 would not be able to communicate with the agents in Figure 8. This convergence to a common ground (Allan, 2013; Tomasello, 2016) or frame of reference (Fields et al., 2021a; Fields et al., 2021b) was mediated by free energy selection processes at the level of the model structure, using Bayesian model reduction (Smith et al., 2020).

### 5.1 BAYESIAN MODEL REDUCTION AND STRUCTURE LEARNING

In contrast to learning—that optimises posteriors over parameters—Bayesian model selection or structure learning (Tenenbaum et al., 2011; Tervo et al., 2016; Tomasello, 2016) can be framed as optimising the *priors* over model parameters. Bayesian model reduction is a top-down approach to this kind of structure learning, which starts with an expressive model and removes redundant parameters to reveal the best sparsity structure. Crucially, Bayesian model reduction can be applied to the posterior beliefs after the data have been assimilated. In other words, Bayesian model reduction is a *post hoc* optimisation that refines current beliefs based upon alternative models that may provide potentially simpler explanations (Friston and Penny, 2011).

Technically, Bayesian model reduction is a generalisation of ubiquitous procedures in statistics, ranging from the Savage-Dickey ratio (Savage, 1954), through to classical F-tests. In our context, it reduces to something remarkably simple: by applying Bayes rules to full and reduced models it is straightforward to show that the change in free energy can be expressed in terms of posterior Dirichlet counts  $\mathbf{a}$ , prior counts  $a$  and the prior counts that define a reduced model  $a'$ . Using  $\mathbf{B}$  to denote the beta function, we have (Friston et al., 2018):

$$\begin{aligned} \Delta F &= \ln P(o | a) - \ln P(o | a') \\ &= \ln \mathbf{B}(\mathbf{a}) + \ln \mathbf{B}(a') - \ln \mathbf{B}(a) - \ln \mathbf{B}(\mathbf{a} + a' - a) \\ \mathbf{a}' &= \mathbf{a} + a' - a \end{aligned} \tag{13}$$

Here,  $\mathbf{a}'$  corresponds to the posterior one would have obtained under the reduced priors.

Clearly, to realise this form of free energy minimisation, one has to have a space of models or reduced priors to evaluate. So, how does one explore the space of priors over parameters? The idea behind active model selection is to consider priors that minimise expected free energy. These models entail a high mutual information and sparse probabilistic mappings: c.f., (Navarro and Perfors, 2011). The active model selection considered here entertains some new log priors<sup>8</sup> that decrease expected free energy by one natural unit (omitting superscripts for clarity):

$$\ell(\hat{a}) = \ell(a) - \frac{\partial G(a)}{\partial \ell(a)} \Rightarrow \hat{a} = a \times \exp\left(-a \times \frac{\partial G}{\partial a}\right) \quad (14)$$

$$\begin{aligned} \frac{\partial G}{\partial a} &= \frac{\partial \bar{a}}{\partial a} \times \frac{\partial G}{\partial \bar{a}} = \left( \frac{1}{\Sigma(a)} - \frac{\bar{a}}{\Sigma(a)} \right) \times \frac{\partial G}{\partial \bar{a}} \\ \frac{\partial G}{\partial \bar{a}} &= 1 - \ell(\bar{a}) + \ell((\bar{a} \odot \mathbf{1}) \otimes (\mathbf{1} \odot \bar{a})) - \varphi(\mathbf{c}) \otimes \mathbf{1} \end{aligned}$$

As in active learning, one can then use Bayesian model averaging to take a weighted mixture of old and new (reduced) priors:

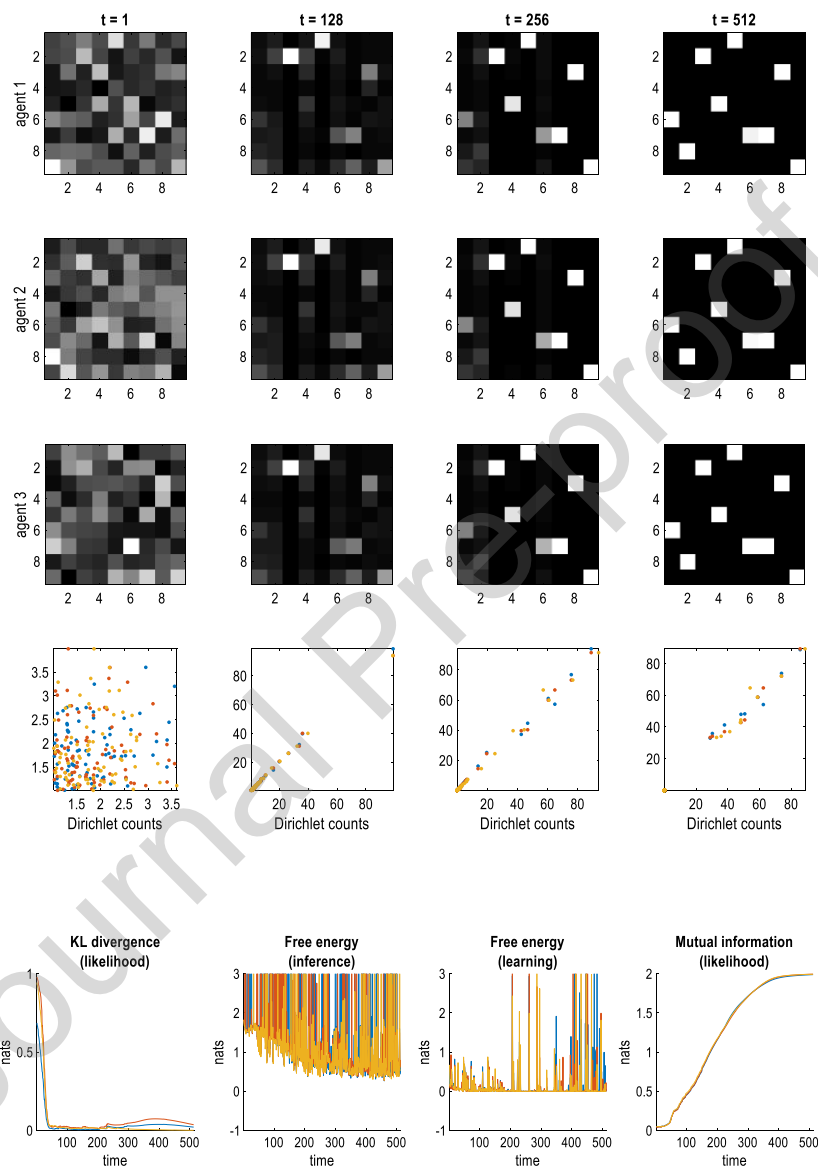
$$\begin{aligned} \mathbb{E}_p[a_\tau^g | u = u_0] &= a \\ \mathbb{E}_p[a_\tau^g | u = u_1] &= \hat{a} \\ \mathbb{E}_p[a_\tau^g] &= P(u_0) \cdot \mathbb{E}_p[a | u = u_0] + P(u_1) \cdot \mathbb{E}_p[a | u = u_1] \\ &\Rightarrow \\ \mathbb{E}_p[a_\tau^g] &= a' = P(u_0) \cdot a + P(u_1) \cdot \hat{a} \\ \mathbb{E}_Q[a_\tau^g] &= \mathbf{a}' = \mathbf{a} + a' - a \end{aligned} \quad (15)$$

$$P(u) = \sigma(-\alpha \cdot F(u)), \quad F(u) = [0, \Delta F]$$

Here, the implicit Bayesian model averaging weights each model by its evidence as scored by the reduction in variational free energy. This kind of model selection is therefore guaranteed to select structures with precise or unambiguous probabilistic mappings. When deployed in the simulations of language emergence, the auditory likelihood mappings converge to the same precise structure over agents; thereby enabling the communication illustrated in Figure 9.

---

<sup>8</sup> Expected free energy gradients are taken with respect to log Dirichlet counts because Dirichlet concentration parameters are nonnegative scale parameters.



**Figure 9 – the emergence of language.** This figure uses a similar format to Figure 7; however, in this simulation we exposed three language-naïve agents to 512 episodes. This means the likelihood mappings start off with a random configuration and autodidactically converge to each other. This contrasts with Figure 7 in which a naïve agent converged to an (identity) likelihood mapping by learning from her teachers. The upper panels show the likelihood mapping from hidden states pertaining to location and the corresponding auditory outcome, for three agents (in each column). The left columns report the initial random connectivity that subsequently resolves into a

sparse structure, with repeated exposure, indicated by the number of exposures,  $t$ . Crucially, a few precise mappings emerge by 128 episodes that are subsequently refined. By the end of the simulation, nearly every hidden state has become associated with a unique ‘word’ in a way that is shared over agents. Interestingly, locations 5 and 6 are reported with the same ‘word’ (outcome 7), speaking to a slight coarsening of the language used by the agents in Figure 7. The plots in the penultimate row show the increasing correlation between the corresponding Dirichlet parameters for all pairs of agents. Initially, there are no correlations because the initial parameters (c.f., connections) were assigned randomly by adding an unsigned random Gaussian variates to 1, for every connection. After 128 exposures, a few of these connections have been selected and acquire high Dirichlet counts of about 80. As time proceeds, other connections are selected, while the remaining connections are eliminated; rendering the final connectivity sparse. The lower row illustrates the correlates of this structure learning or language acquisition in terms of the KL divergence between all pairs of likelihood mappings between the agents; the ensuing variational free energy associated with inference; the free energy associated with learning and, finally, the mutual information of the likelihood mappings for location, for each agent (coloured lines). In this example, the likelihood mappings become nearly identical by about 128 exposures with a reduction in variational free energy. As anticipated, there is a progressive increase in the mutual information afforded by these auditory likelihood mappings, saturating at about two natural units (i.e., about three bits).

Note how active learning and selection complement each other. In active learning, posterior parameters change to minimise variational free energy if, and only if, expected free energy is reduced. Conversely, in active selection, prior parameters change to minimise expected free energy if, and only if, variational free energy is reduced—as scored with Bayesian model reduction. The implicit bootstrapping underwrites self-evidencing towards precise and predictable exchanges with—in the context of communication—a co-constructed world.

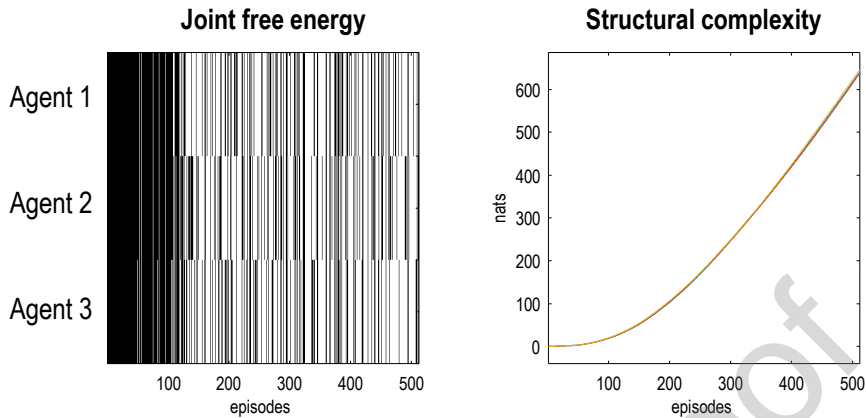
Figure 10 shows the ensuing increase in the ability of agents to resolve uncertainty about the scene. The left panel illustrates episodes when variational free energy, per epoch and modality, was less than one. In other words, when the self-information or surprisal afforded by each observation was negligible<sup>9</sup>. Initially, all three agents remain rather confused about what is going on around them, with high levels of variational free energy. As they acquire the ability to share their beliefs, the number of episodes with a low free energy starts to increase where, in some instances, all three agents have a veridical understanding of the latent states generating their observations. This is accompanied by a progressive sparsification (Spielman and Srivastava, 2011) and convergence of their likelihood mappings onto a shared structure. The complexity of this structure can be scored in terms of the cumulative KL divergence between the posterior over the Dirichlet distributions, in relation to the initial distributions (i.e., an approximation to the information length or distance between initial and current posteriors over parameters). This is the cumulative free energy or information gain (i.e., complexity) due to parameters

---

<sup>9</sup> Roughly speaking, a variational free energy of 3 would be considered surprising in the sense of a Bayes factor, relative to perfect predictions Kass, R.E., Raftery, A.E., 1995. Bayes Factors. Journal of the American Statistical Association 90, 773-795..



and can be read a measure of *structural complexity* that increases progressively with experience (right panel of Figure 10).



**Figure 10 – Schematic.** the emergence of structural complexity. The left panel indicates episodes in which the variational free energy of inference—per modality and observation—fell below one (white bars), towards its lower bound of zero, for each agent (rows). This illustrates the incidence of low free energy (high model evidence) episodes increases as language is acquired. The right panel shows the increase in structural complexity that underwrites this improvement in inference—mediated by belief sharing or communication—over time. The three lines correspond to the three agents. This is the cumulative expected free energy due to learning in Figure 9.

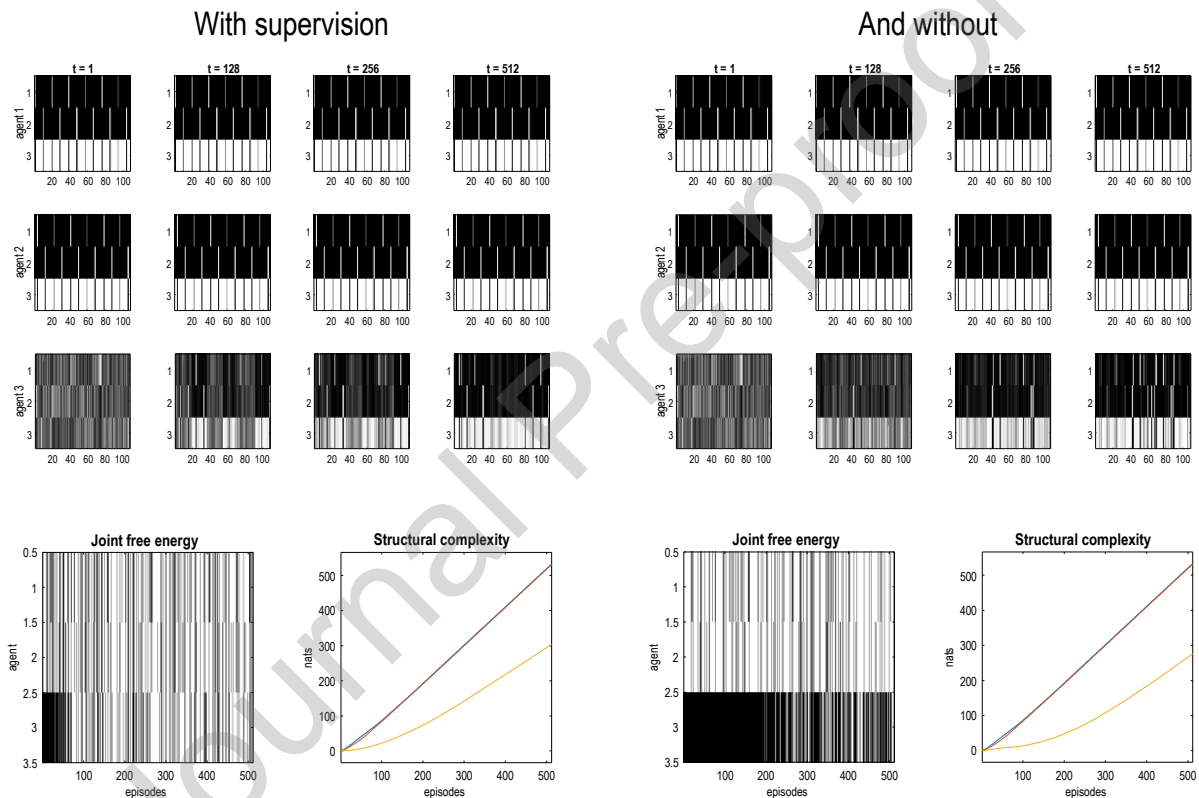
## 5.2 SUPERVISED STRUCTURE LEARNING

When illustrating active inference, we saw that belief-sharing augmented the resolution of uncertainty, enabling a more precise grip on states of affairs (see Figure 5). One can now ask whether belief sharing has a similar synergetic role in learning. In other words, could communication scaffold the learning of likelihood models and thereby disentangle observations to recognise their hidden causes? To address this, we repeated the preceding simulations, equipping all agents with precise language mappings but withholding precise visual mappings from the last agent (by setting all the Dirichlet parameters to one, plus an unsigned random Gaussian variate). Heuristically, this can be regarded as replacing an experienced sentinel with a young novice, who has yet to develop a likelihood model for the causes of her visual observations. However, she can hear her ‘supervisors’ talking about those causes. Heuristically, this can be likened to teaching a child to read by pairing the presentation of pictures and spoken names of an object.

Under this kind of supervised structure learning, the supervisee quickly comes to learn a visual likelihood mapping—that is sufficient to render her inferences indistinguishable from her supervisors—by about 64 exposures. Figure 11 illustrates this in terms of the acquisition of the likelihood tensors mapping from hidden states to visual outcomes (here, the central, contrast energy modality). This supervised structure learning can be contrasted with the learning in the right panels of Figure 11, when

the novice cannot hear her supervisors (implemented by setting the Dirichlet counts of her auditory likelihood mappings to 64 everywhere). In the absence of supervision, learning proceeds slowly and does not support skilled inference, even after 512 exposures.

These numerical analyses speak to a generic point: belief-sharing appears to be a fundamental in realising joint free energy minimisation, distributed over ensembles of free energy minimising processes and over the timescales at which these processes unfold. Technically, this is just a realisation of a variational principle of least action, where action is the path or time integral of distributed—and extensive<sup>10</sup>—variational free energies.



**Figure 11 – supervised structure learning.** This graphic reports numerical experiments asking whether communication enhances the learning of likelihood models. In this example, the last agent was rendered visually naïve by setting the Dirichlet counts mapping from hidden states to visual outcomes to one (plus an unsigned random Gaussian variate). However, proprioceptive and auditory (identity) mappings were preserved for all three agents. The ensuing simulations illustrate learning of the visual likelihood mapping: e.g., the development of visual pathways and subsequent skilled recognition of the causes of visual sensations. The left hand panels illustrate the implicit supervised learning, while the right hand panels reproduce the same stimulation in the absence of supervision (by rendering the third agent impervious to auditory input). The upper panels show the likelihood tensor mapping from hidden states to the central contrast energy modality for the three agents (rows), over selected time points during the 512 simulated episodes (columns). Because these likelihood mappings are tensors, we have placed the slices of the tensor next to each other, to form a matrix. For example, the first row of

<sup>10</sup> Extensive in the sense that the total free energy is the sum of local free energies.

each matrix illustrates the fact that there are six combinations of hidden states generating a *near* outcome. These combinations correspond to the three directions of *gaze* times the two *poses* of the subject (friend or foe). The lower panels report the ensuing inference and cumulative free energies due to learning (i.e., structural complexity) using the format of Figure 10. The key point to take from this figure is that supervision—via belief sharing—enables a grip on the world that is indistinguishable from agents with veridical generative models. Conversely, when denied supervision, learning is based purely upon visual and proprioceptive inputs. In this instance, the unsupervised agent learns slowly, failing to disentangle the causes of her observations and attain the free energy minimisation of her conspecifics. This is reflected in a slight lower structural complexity of her likelihood mappings, which are more disorganised and diffuse with, implicitly, a lower mutual information.

There are clearly many other simulations that we could pursue from this point. For example, we could combine the above simulations to simulate evolution by selecting those agents that have the lowest free energy (i.e., highest marginal likelihood or adaptive fitness) to augment the structure learning and language acquisition evinced above (Friston et al., 2023a). One can also consider the emergence of language under different frames of reference. For example, we have assumed that the generative model of each agent shares a common frame of reference in relation to allocentric locations. This is not necessary. In principle, it should be possible to use hidden states in an egocentric frame of reference and have agent-specific likelihood mappings that support belief sharing. This kind of simulation would become even more interesting if the agents moved around. However, for the purposes of the current paper, we now turn to a discussion of the phenomena illustrated by these numerical studies.

## 6 DISCUSSION

---

The foregoing illustrates the emergence of distributed (i.e., federated) inference and learning—where posterior beliefs are shared among agents—under the imperative to maximise the evidence for (generative) models of a shared world; namely, self-evidencing (Hohwy, 2016). This can be read from a number of perspectives: from a systemic perspective, this emergence can be read as the minimisation of joint free energy that ensues when inference, learning and selection are simulated as nested free energy minimising processes. An anthropomorphic interpretation is in terms of linguistic communication; namely, the acquisition of language, and its transmission over generations, in the spirit of cultural niche construction.

A key technical point illustrated in these simulations is the optimisation of free energy functionals of (Bayesian) *beliefs* about the latent causes of observations over temporal scales. This foregrounds the importance of working with probabilistic representations of latent states, parameters and structures that constitute a generative world model. Active learning and model selection would not be possible without the sufficient statistics of beliefs about model parameters and implicit model structure. When working

with discrete state space models, the requisite variational updates transpire to be straightforward, local and—at a certain level of analysis—biomimetic. In what follows, we consider some key perspectives on the ensuing mechanics.

## 6.1 SELF ORGANISATION: A PERSPECTIVE FROM PHYSICS

It is interesting to reflect upon the behaviours that emerge—under free energy minimisation—in light of questions about emergence in random dynamical systems (Arnold, 2003; Crauel and Flandoli, 1994); namely, is the emergence of sparse coupling and generalised synchrony inevitable? (Bak et al., 1988; Ellis et al., 2011; Ellison et al., 2011; England, 2013, 2015; Gershenson, 2012; Hunt et al., 1997; Jafari et al., 2016; Jeffery et al., 2019; Namikawa, 2005; Sakthivadivel, 2022b). In other words, is the sparse coupling, self-organised criticality and synchronisation of chaos—seen in complex dynamical systems—a necessary property of such systems that exist?

As noted in the introduction, the free energy principle just prescribes a variational principle of least action that can be applied to any random dynamical system that possesses an attracting set (i.e., a pullback attractor) with a Markov blanket (Crauel and Flandoli, 1994; Friston et al., 2022a; Sakthivadivel, 2022b). This means that the accompanying dynamics of such systems conform to certain principles that can either be articulated in terms of the free energy principle or, equivalently, a constrained maximum entropy principle (Sakthivadivel, 2022a). The constraints in question here are furnished by the generative model that describes the pullback attractor in terms of a probability density. In effect, this enables one to simulate a system—as it converges on its pullback attractor—by expressing the dynamics as a functional of a generative model. The numerical analyses above are an example of this, under generative models of discrete states.

This formulation licences teleological interpretations, such as optimisation, inference and learning. Indeed, when variational and expected free energy are read as objective functions, they can be seen as equivalent to Bayes-optimal inference and learning, respectively (Winn and Bishop, 2005). However, one ubiquitous aspect of free energy minimisation does not have a clear teleology. This is the emergence of complexity as free energy is minimised at various timescales. Heuristically, this is unremarkable in the sense that free energy is a bound on log evidence. And log evidence is accuracy minus complexity<sup>11</sup>.

---

<sup>11</sup> Where complexity is read as the relative entropy or KL divergence between posterior and prior beliefs.

This means that if any system or agent learns a more accurate account of its exchanges with an environment (or other agents) its complexity must increase.

What is remarkable is the way in which this increasing complexity is expressed across temporal scales in the numerical studies above. While, on average, the variational free energy associated with beliefs about latent states declines, with learning and language acquisition, the complexity associated with the likelihood parameters increases, as knowledge is accumulated<sup>12</sup>. This means that the agents or observers progressively evince a more complex and sparse internal structure—quantified, for example, in the structural complexity of Figure 10. The question now is whether this kind of structural complexity goes hand in hand with communication. In other words, are communication and structural complexity themselves emergent properties of any loosely-coupled random dynamical system within which subsystems (e.g., agents) can be identified. Put simply, are our generative or world models destined for progressive increases in complexity simply because we have to explain and understand our exchange with conspecifics, colleagues, confederates and conspirators. Clearly, to answer this kind of question with simulations one will have to scale up the numerical studies above. Some (active inference) work in this direction addresses the spread of ideas at scale (Albarracín et al., 2022; Heins et al., 2023; Kastel and Hesp, 2021).

## 6.2 CONVERSELY, COMMUNICATION BEGETS SPACE: A PERSPECTIVE FROM QUANTUM INFORMATION THEORY

The communication considered above unfolds in space and time; indeed, location in space is what it is primarily about. The communicating agents are both embedded in space and separated in space. The way that the agents are embedded in space not only gives them something to talk about; it also assures that they have different points of view, and hence that they only partially share their world: c.f., (Williford et al., 2018). All of this is so obvious that it is seldom thought about. Distinct agents occupying distinct positions in a surrounding spatial ‘container’ is assumed as a matter of course when modeling communication.

The model presented in Figures 1 and 2, however, makes no assumptions about an embedding space: it is constructed entirely out of states (which occupy some state-space) and tensors that act on them (operators on the state-space): c.f., (Fields et al., 2023a; Knill and Laflamme, 1997). The idea of an embedding space that supports a radial distance and a distinction between angles of view is not introduced until Figure 3. We can therefore ask: where does this (projective) embedding space come

---

<sup>12</sup> The free energy due to learning in Figure 7 scores the increase in complexity after each episode.

from? Is it just an assumption of convenience? What role is the embedding space playing in the communication scenario, other than providing something (location) for the observers to talk about?

Figure 3 makes the second role of the embedding space explicit: it serves to distinguish and separate the agents. The embedding space is not, however, strictly needed for this; the model could simply have assumed that the agents all have mutually conditionally independent states (i.e., that each agent has her own Markov blanket). In this case, Figure 3 can be described as depicting a set of distinct agents that share a classical communication channel. The assumption of a shared generative model then becomes the assumption that they share a ‘language’ that allows them to make sense of—i.e., assign sufficiently similar semantics to—each other’s messages. The agent’s perceptual capabilities are described in spatial terms, in terms of “close” versus “near” detectors and angle of gaze. These can, however, also be considered abstract variables—a one-bit variable and a two-bit variable—to which their shared semantics refers. The computation each agent performs remains the same: correlating the values of these variables that they obtain by direct interaction with the world with the values that they obtain from the channel, and hence from the (otherwise unperceived) other agents.

The results of the numerical experiments can now be re-interpreted. What the agents discover is that they share an error-correcting code: there is something redundant about their shared world that confers particular symmetries on the correlations between what they ‘see’ and what they ‘hear’. With sufficient experience, each can infer that there is a fast, almost always cyclic permutation symmetry that is coupled to a slower binary oscillation. They share, in other words, a two-dimensional world with an angular and a radial degree of freedom. As noted by one of our reviewers, these arguments might generalize to any metric space that underwrites sense making and communication: “To me, the question here is more general, about the need for a commonly experienced ground, world, spaces in general, allowing a common system of beliefs, and language [to share messages] about it.”

The idea that communication and space are effectively dual concepts—tied together by the ability of each to function as an error-correcting code for the other—is fundamental to much of quantum gravity research (Bain, 2020; Fields et al., 2023b). It is also attested by the co-development of sensory integration and the sense of space via motor babbling in infants (Baranes and Oudeyer, 2009; Saegusa et al., 2009). That relatively simple experiments, motivated by the free energy principle, illuminate this deep connection is surely of interest.

### 6.3 FEDERATED LEARNING AND DISTRIBUTED COGNITION: A PERSPECTIVE FROM COMPUTER SCIENCE

In computer science, distributed machine learning has become increasingly important (Verbraeken et al., 2020), especially in the context of deep learning: recent breakthroughs in training large language models (Biderman et al., 2023) or agents with deep reinforcement learning (OpenAI et al., 2019) are fueled by running the training process on hundreds of GPUs in parallel. In these systems, the optimization method used is stochastic gradient descent (SGD), where parameter updates are calculated by averaging the gradients over a batch of data. Parallelism can hence be exploited by duplicating the model over a number of workers—each calculating a gradient update on a part of a [mini]batch, which is then broadcasted to the others: i.e., data parallelism (Goyal et al., 2017). A different way for parallelizing the training workload is by exploiting the layered structure of deep neural networks, and deploying different layers on different workers: i.e., model parallelism (Shoeybi et al., 2019).

Federated machine learning is a special case of distributed machine learning (Yang et al., 2019), in which the training process is distributed among multiple parties, each with access to its own dataset. Crucially, each party wants to learn the better model without, however, sharing their raw data, e.g., due to privacy or security concerns. When applied to distributed SGD, each party calculates gradients locally on parts of their dataset and shares only these with other parties, e.g., further secured by using secret sharing (Bonawitz et al., 2017), homomorphic encryption (Phong et al., 2018), or differential privacy (Shokri and Shmatikov, 2015).

The processes illustrated in this paper speak to a more liberal notion of federated learning. First, instead of limiting belief-sharing to a noisy gradient estimate—from a random minibatch of data in deep learning—under active inference, agents communicate full belief distributions. This renders communication more effective, as information is only shared when required, i.e., as evaluated by expected free energy. Regarding privacy, the agents themselves agree on what they want to share, by jointly learning the likelihood parameters of the communication modality. This can be thought of as developing their own private code, which can only be deciphered when having access to the same shared belief space. Similarly, differential privacy can be accomplished by only communicating aggregate Dirichlet counts, rather than individual observations.

This perspective on federated learning derives from research on the emergence and dynamics of human communication in computational neuroscience and evolutionary cognitive anthropology. The ensuing approach may enable the design of generic agents that draw their data from different, complementary sources, by composing a series of communicating local models into a network of belief-sharing. This helps to finesse the highly non-trivial problem of automatically designing low-cost models with the

right number and configuration of parameters to enable the emergence of such a network (Friston et al., 2022b; Kauffman and Johnsen, 1991; Odling-Smee et al., 2013).

Decentralised learning—the kind of belief-sharing illustrated in our simulations—requires a high degree of similarity between agents' beliefs, which are learned by sharing data (outcomes). This kind of individual optimization, even when based on a shared language, may not be sufficient in cases where the communication channels between agents are perturbed (e.g., communicating under water distorts the outcomes produced by agents adapted to live on land). How can one design a generic shared model, through communicative interaction, if the signals sent along agents are distorted? Thus, beyond mere similarity or even isomorphism, another condition for belief sharing is the presence of a communication channel that itself is able to maintain the integrity of observable outcomes in the environment; namely, a shared, likelihood model responsible for the structure of the sensory observations to be shared between agents.

This reflects the fact that communication and distributed cognition require that the *communication milieu* should be conducive to shared understanding, and speaks to a tenet of our approach. In the active inference literature, issues regarding the communication milieu are discussed under the banner of *econiche construction* (Creanza and Feldman, 2014; Krakauer et al., 2009; Laland et al., 1999; Lehmann, 2008). This highlights the symmetry of the active inference formulation, which describes not only the manner in which agents attune to their environment, but also how, on average and over time, by accumulating the traces left by actions of agents, the environment also comes to reflect the statistical structure of its denizens (Constant et al., 2018; Friston et al., 2022b; Veissiere et al., 2019). We pursue this from the point of view of ethology.

#### **6.4 CULTURAL NICHE CONSTRUCTION: A PERSPECTIVE FROM NEUROETHOLOGY**

In evolutionary biology, *econiche construction* can be broadly understood as the implicit and explicit modification by organisms of their own environment (Odling-Smee et al., 2013), serving functions at multiple spatial and temporal scales. This modification comes in two flavours. At the phylogenetic scale (that of the entire species), *selective* niche construction generates new feedback loops that can steer selection pressures and change the fitness landscape in a way that benefits the agents that constructed the niche (Han and Hui, 2014). At the ontogenetic scale (of individuals and their development), *developmental* niche construction allows for the reproduction of the life cycle by securing the availability of expected developmental inputs (Stotz, 2017). At the scale of behaviour, cognitive niche construction guides the execution of cognitive functions such as perception, action and learning, by



supporting the performance of those functions (Bertolotti and Magnani, 2017) and canalizing them down specific paths (Constant et al., 2018; Veissiere et al., 2019).

From the perspective of belief-sharing under active inference, at the phylogenetic level, niche construction allows for the selection of environments (i.e., a communication milieu) that best fits the communicative abilities of the agents, thereby preparing agents to synchronize to one another, albeit vicariously, through their environment over developmental time (Bruineberg et al., 2018). Over developmental and intergenerational time, explicit and implicit modifications of the econiche allow the niche to embody knowledge that can be transmitted (horizontally) to the next generation, thereby supporting synchrony and communication (Constant et al., 2018; Manrique and Walker, 2023). The general picture is that of econiche construction as a process that stabilises and maintains the communication milieu by contributing to the selection of organisms endowed with the necessary biological apparatus to sense observations characteristic of the communication milieu (i.e., the selection of the isomorphism), and also, crucially, by encoding or storing those observations in a reliable way. Cultural patterns such as written language would be cases in point of such co-evolved econiches, allowing for synchronisation and sophisticated cognitive functions such as mind reading (Heyes, 2018; Veissiere et al., 2019)

Situations—like the synthetic sentinel simulation considered above—are found in nature. Mammals and birds can exhibit referential communication, in which specific alert calls are used to communicate predator types (Gill and Bierema, 2013; Townsend and Manser, 2013). For example, Japanese great tits (*Parus minor*) use combinations of more than a dozen different notes to share the type of predator identified—including whether they are flying (e.g., crows) or approaching the nest from the ground (e.g., martens)—to broadcast the likely location of the predator (Suzuki, 2014). In the same way that humans combine a finite set of words to create compositional syntax with infinite meanings, tits also use compositional syntax, such as combining ‘scan for danger’ and ‘approach the caller’ notes (Suzuki, 2014; Suzuki et al., 2016). These observations suggest the existence of a hierarchical and factorial (neuronal) architecture that integrates shared messages to form a posterior belief. In avian vocal learning, young birds empirically learn the meaning of calls of adult birds, in which certain neural populations in the higher-level auditory cortex, called the caudomedial nidopallium, respond selectively to the learned song (Yanagihara and Yazaki-Sugiyama, 2016). This further suggests the formation of a sparse representation, consistent with a sparse likelihood mapping acquired via active learning. In short, these empirical observations imply possible targets, when testing process theories based on active inference.

## 6.5 CO-EVOLUTION: A PERSPECTIVE FROM EVOLUTIONARY BIOLOGY

In the simulations above, communication among agents was auditory and synchronous. Federated learning can also occur through other sensory modalities, and importantly through the asynchronous engagement of multiple agents within a jointly-modified niche. For example, in the case of an ant colony, the distribution of pheromones around the nest can be considered as an extended colony-level memory system, akin to an inscribed symbolic language. In this stigmergic setting, the learning process occurs among nestmates indirectly, through their ongoing contact with the niche. From the perspective of an ant nestmate's generative model, the observed pheromone distribution provides valuable information that guides their next movements (Friedman et al., 2021), reflecting learned or inherited associations between that pheromone and its (semantic) meaning. From the perspective of the environment, the modified niche can be seen as a trace of the statistical structure of the ant colony, in line with the agent-niche symmetry discussed above: c.f., (Bruineberg et al., 2018). This type of extended and federated cognition underlies the distributed physiology found in eusocial insect colonies, enabling collective intelligences with scale and scope far beyond any nestmate's body (Friedman et al., 2020)

When considering all the resources of active inference, one can see how federated systems—implementing belief-sharing and econiche construction—may allow efficient training: individual agents learning their respective local model based on shared observations, which would be encoded themselves through a learning process in a generative process that would function as the econiche. Intergenerational transmission of developmental tendencies considered as Bayesian model selection, neatly connects the computational challenges and affordances of federated learning with modern perspectives in evolutionary biology (Friston et al., 2023b).

## 6.6 COLLECTIVE INTELLIGENCE AND GROUP COGNITION: A PERSPECTIVE FROM COMPLEX DYNAMICAL SYSTEMS

This work also sits in close relation to the study of collective intelligence and collective computation in natural and artificial systems. Many popular models of collective inference and learning cast agents as driven by simple rules for transforming environmental and social information into individual decisions, which translate to collective outcomes (Beckers et al., 1990; Couzin et al., 2005; Pratt and Sumpter, 2006; Strandburg-Peshkin et al., 2015). Under the formalism adopted here, these decision rules might be re-cast as approximations to a form of belief-sharing among agents about some shared latent states of the environment (Albarracín et al., 2022; Heins et al., 2023; Krafft et al., 2021).

In contrast to the explicit belief-sharing approach adopted for the current simulations, in many natural collective scenarios (e.g., in animals without explicit communication modalities), these common signals may only indirectly relate to shared contextual variables, and often manifest as noisy, ambiguous sensory input to conspecifics (Couzin et al., 2005; Pérez-Escudero and de Polavieja, 2011; Torney et al., 2015). This sort of indirect information sharing has consequences for group intelligence and performance; for example, by introducing the risk of amplifying noisy or irrelevant information (Albarracín et al., 2022; Couzin et al., 2011; Poel et al., 2022; Sosna et al., 2019). In the context of multi-agent active inference, such maladaptive group outcomes may correspond to local minima in a joint free energy landscape—while individuals may sit at fixed points in their (private) free energy functionals, these fixed points may not always coincide with fixed points in the group’s free energy landscape (Heins et al., 2023).

Further research might help identify the conditions on what types of sensory information channels (and corresponding likelihood models) are sufficient to enable the group to perform optimally, e.g., in the context of consensus decision making or collective perception (Berdahl et al., 2013; Hein et al., 2015; Ward et al., 2008). Furthermore, how exactly conspecifics should update their generative models (e.g., through parameter learning or structure learning) to minimize collective free energy, as opposed to individual-level free energy, is an open question.

## **6.7 HOW THE DESIRE TO SHARE CREATES LANGUAGE: A PERSPECTIVE FROM PSYCHOLOGY**

Most of our beliefs about the world come, not from direct experience, but from other people. All animals, from fruit flies to humans, learn about the world by observing others (Frith, 2010; Kilner et al., 2007; Manrique and Walker, 2023; Rieucou and Giraldeau, 2011). For example, they learn where to go and what to eat. The effect of this learning is that the children will come to share the priors of their parents and of their culture more generally. But humans are unique, since we can learn about the world from others through verbal instruction. Such instruction depends upon communication, mostly via language.

Humans are incredibly social animals (Manrique and Walker, 2023; Torney et al., 2015; Vasil et al., 2020). We not only want to be *liked* by our friends and neighbours (our in-group). We want to be *like* them. And the best way to be like others is to align with them. We achieve such alignment at many levels. At the physical level, we dance, negotiate the city streets, and move furniture together. At a more abstract level, we share our goals. And at the highest, mental level we share our ideas and, in particular, our models of the world. Many advantages accrue from sharing models of the world. Because of our

complementary perspectives, our shared models will be richer and more accurate than any individual model. Furthermore, having shared models, and common ground more generally, communication becomes easier (Bahrami et al., 2010; Clark and Brennan, 1991; Frith and Wentzer, 2013; Heyes and Frith, 2014).

In this context a virtuous circle is created. The desire to share encourages the emergence of language. Language enlarges our common ground and makes our models of the world more accurate. This creates a mental niche in which communication is more efficient and our ability to share is enhanced. And all this can be achieved—or, strictly speaking, described—by free energy minimisation.

## 6.8 CONSTRAINTS AND COMMUNICATION: A PERSPECTIVE FROM LINGUISTICS

The level of analysis—in this treatment of communication—falls short of addressing linguistics *per se*. However, there are some crosscutting themes in classical linguistic theories such as Universal Grammar (Chomsky, 2017) and Optimality Theory (Prince and Smolensky, 2007) that deserve mention. Universal Grammar rests on the notion of innate *constraints* on the grammar of possible languages. Similarly, Optimality Theory (Prince and Smolensky, 2007) suggests that the observed forms of language arise from the optimal satisfaction of *constraints*. Both approaches foreground the role of constraints that— from the perspective of communication under active inference—arise naturally from free energy minimisation. For example, the free energy principle is dual to the *constrained* maximum entropy principle (Sakthivadivel, 2022a), where the constraints inherit from the generative model. Intuitively, if language is the broadcasting of beliefs under a generative model, then the structure of language should inherit the factorial (and deep) structure of the underlying generative model. In turn, the structure is constrained by the causal structure of the lived world within which we act—a world that includes our bodies and conspecifics.

On this view, it is literally self-evident (in the sense of self-evidencing) that language is subject to innate constraints; constraints that inherit from learning the structure of generative models apt to explain the co-constructed world. These models necessarily have a grammar; in the sense of dynamics and non-Markovian temporal structure, as evinced in hierarchical structure learning treatments. See for example: (Davis and Johnsrude, 2003; Friston et al., 2017c; George and Hawkins, 2009; MacKay and Peto, 2008; Stoianov et al., 2022; Yildiz et al., 2013; Young et al., 2018; Zorzi et al., 2013). In short, grammar and meaning should be isomorphic with world models.

This view also speaks to the Symbol Grounding Problem (Harnad, 1990); namely, how words acquire meaning. The problem of meaning is dissolved if one commits to the notion that words are (overt or covert) declarations of beliefs, and beliefs are over discrete states. In other words, the meaning of a

word just is [isomorphic to] the beliefs about the current state of the world and the narratives inherent in state transitions. The symbol grounding problem is dissolved in the sense that words are both cause and consequence of world models that underwrite an active engagement with the world. The emergence of shared meaning (c.f., common ground) then reduces to an alignment of the (likelihood) mapping from belief states to words or symbols, among correspondents. It is interesting to speculate about the impact of generative artificial intelligence (AI) in general, and large language models in particular, in light of this view. For example, are the implicit generative models in generative AI sufficiently expressive to constitute a world model that includes the consequences of their action? E.g., (Chalmers, 2023).

## 7 CONCLUSION

---

The account given above offers several key advances on state-of-the-art active inference and structure learning. The technical contributions of this paper are twofold. The first is belief-sharing—the idea that biological or artificial agents with different vantage points can communicate their inferences and inform one another about their shared environment. We saw this in the context of agents, with different viewpoints, speaking to one another to share their percepts to better characterise their world. The second axis is an intelligent procedure for updating beliefs, at the level of learning and model selection. This procedure—based upon the prior belief that our world should be precise, predictable, and sparse—furnishes a (common) sense of purpose in both learning and selection. When combined, these two advances offer a powerful form of federated belief-optimisation that may shed light on biological development and the design of intelligent (eco) systems in general, and belief sharing in particular.

### Software note

Although the generative model – specified by the **(A,B,C,D)** matrices – changes from application to application, the belief updates in Figure 2 are generic and can be implemented using standard routines (here `spm_MDP_VB_XXX.m`). These routines are available as Matlab code in the SPM academic software: <http://www.fil.ion.ucl.ac.uk/spm/>, or available from <https://github.com/spm/>.

### Acknowledgements

We would like to thank our two anonymous reviewers for very helpful guidance in presenting this work. KF is supported by funding for the Wellcome Centre for Human Neuroimaging (Ref: 205103/Z/16/Z), a Canada-UK Artificial Intelligence Initiative (Ref: ES/T01279X/1) and the European Union’s Horizon 2020 Framework Programme for Research and Innovation under the Specific Grant Agreement No. 945539 (Human Brain Project SGA3). AC is supported by a European Research Council Grant (XSCAPE) ERC-2020-SyG.

## Disclosure statement

The authors have no disclosures or conflict of interest.

## References

- Adank, P., Hagoort, P., Bekkering, H., 2010. Imitation improves language comprehension. *Psychol Sci* 21, 1903-1909.
- Albarracín, M., Demekas, D., Ramstead, M.J.D., Heins, C., 2022. Epistemic Communities under Active Inference. *Entropy (Basel)* 24.
- Allan, K., 2013. What is Common Ground?, in: Capone, A., LoPiparo, F., Carapezza, M. (Eds.), *Perspectives on Linguistic Pragmatics*, vol. 2, pp. 285-310.
- Arnold, L., 2003. *Random Dynamical Systems (Springer Monographs in Mathematics)*. Springer-Verlag, Berlin.
- Attias, H., 2003. Planning by Probabilistic Inference, *Proc. of the 9th Int. Workshop on Artificial Intelligence and Statistics*.
- Ay, N., Bertschinger, N., Der, R., Guttler, F., Olbrich, E., 2008. Predictive information and explorative behavior of autonomous robots. *European Physical Journal B* 63, 329-339.
- Bahrami, B., Olsen, K., Latham, P.E., Roepstorff, A., Rees, G., Frith, C.D., 2010. Optimally interacting minds. *Science* 329, 1081-1085.
- Bain, J., 2020. Spacetime as a quantum error-correcting code? *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 71, 26-36.
- Bak, P., Tang, C., Wiesenfeld, K., 1988. Self-organized criticality. *Physical review. A, General physics* 38, 364-374.
- Baranes, A., Oudeyer, P.Y., 2009. R-IAC: Robust Intrinsically Motivated Exploration and Active Learning. *IEEE Transactions on Autonomous Mental Development* 1, 155-169.
- Barlow, H., 1961. Possible principles underlying the transformations of sensory messages, in: Rosenblith, W. (Ed.), *Sensory Communication*. MIT Press, Cambridge, MA, pp. 217-234.
- Barsalou, L.W., 2008. Grounded cognition. *Annual review of psychology* 59, 617-645.
- Beal, M.J., 2003. *Variational Algorithms for Approximate Bayesian Inference*. PhD. Thesis, University College London.
- Beckers, R., Deneubourg, J.L., Goss, S., Pasteels, J.M., 1990. Collective decision making through food recruitment. *Insectes Sociaux* 37, 258-267.
- Berdahl, A., Torney, C.J., Ioannou, C.C., Faria, J.J., Couzin, I.D., 2013. Emergent sensing of complex environments by mobile animal groups. *Science* 339, 574-576.
- Berger, J.O., 2011. *Statistical decision theory and Bayesian analysis*. Springer, New York; London.
- Berlyne, D.E., 1954. A THEORY OF HUMAN CURIOSITY. *British Journal of Psychology* 45, 180-191.
- Bertolotti, T., Magnani, L., 2017. Theoretical considerations on cognitive niche construction. *Synthese* 194, 4757-4779.
- Biderman, S., Schoelkopf, H., Anthony, Q., Bradley, H., O'Brien, K., Hallahan, E., Aflah Khan, M., Purohit, S., Sai Prashanth, U., Raff, E., Skowron, A., Sutawika, L., van der Wal, O., 2023. Pythia: A Suite for Analyzing Large Language Models Across Training and Scaling, p. arXiv:2304.01373.

- Blakemore, S.J., Frith, C.D., Wolpert, D.M., 1999. Spatio-temporal prediction modulates the perception of self-produced stimuli. *J Cogn Neurosci* 11, 551-559.
- Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H.B., Patel, S., Ramage, D., Segal, A., Seth, K., 2017. Practical Secure Aggregation for Privacy-Preserving Machine Learning, Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security. Association for Computing Machinery, Dallas, Texas, USA, pp. 1175–1191.
- Botvinick, M., Toussaint, M., 2012. Planning as inference. *Trends in Cognitive Sciences* 16, 485-488.
- Braun, D.A., Ortega, P.A., Theodorou, E., Schaal, S., 2011. Path integral control and bounded rationality, Adaptive Dynamic Programming And Reinforcement Learning (ADPRL), 2011 IEEE Symposium on. IEEE, Paris, pp. 202 - 209.
- Bruineberg, J., Rietveld, E., Parr, T., van Maanen, L., Friston, K.J., 2018. Free-energy minimization in joint agent-environment systems: A niche construction perspective. *Journal of theoretical biology* 455, 161-178.
- Chalmers, D.J., 2023. Could a Large Language Model be Conscious?, p. arXiv:2303.07103.
- Chomsky, N., 2006. *Language and Mind*. Cambridge University Press.
- Chomsky, N., 2017. Language architecture and its import for evolution. *Neuroscience and biobehavioral reviews* 81, 295-300.
- Clark, H.H., Brennan, S.E., 1991. Grounding in communication, in: Resnick, L., B, L., John, M., Teasley, S., D (Eds.), *Perspectives on Socially Shared Cognition*. American Psychological Association, pp. 13--1991.
- Constant, A., Ramstead, M.J.D., Veissiere, S.P.L., Campbell, J.O., Friston, K.J., 2018. A variational approach to niche construction. *J R Soc Interface* 15.
- Constant, A., Ramstead, M.J.D., Veissiere, S.P.L., Friston, K., 2019. Regimes of Expectations: An Active Inference Model of Social Conformity and Human Decision Making. *Front Psychol* 10, 679.
- Couzin, I.D., Ioannou, C.C., Demirel, G., Gross, T., Torney, C.J., Hartnett, A., Conradt, L., Levin, S.A., Leonard, N.E., 2011. Uninformed individuals promote democratic consensus in animal groups. *Science* 334, 1578-1580.
- Couzin, I.D., Krause, J., Franks, N.R., Levin, S.A., 2005. Effective leadership and decision-making in animal groups on the move. *Nature* 433, 513-516.
- Crauel, H., Flandoli, F., 1994. Attractors for Random Dynamical-Systems. *Probab Theory Rel* 100, 365-393.
- Creanza, N., Feldman, M.W., 2014. Complexity in models of cultural niche construction with selection and homophily. *Proc Natl Acad Sci U S A* 111 Suppl 3, 10830-10837.
- Dauwels, J., 2007. On Variational Message Passing on Factor Graphs, 2007 IEEE International Symposium on Information Theory, pp. 2546-2550.
- Davis, M.H., Johnsrude, I.S., 2003. Hierarchical processing in spoken language comprehension. *J Neurosci* 23, 3423-3431.
- Ellis, G.F.R., Noble, D., O'Connor, T., 2011. Top-down causation: an integrating theme within and across the sciences? *Interface focus* 2, 1-3.
- Ellison, C.J., Mahoney, J.R., James, R.G., Crutchfield, J.P., Reichardt, J., 2011. Information symmetries in irreversible processes. *Chaos* 21, 037107.
- England, J.L., 2013. Statistical physics of self-replication. *The Journal of chemical physics* 139, 121923.
- England, J.L., 2015. Dissipative adaptation in driven self-assembly. *Nature nanotechnology* 10, 919-923.
- Fields, C., Fabrocini, F., Friston, K., Glazebrook, J.F., Hazan, H., Levin, M., Marciandò, A., 2023a. Control Flow in Active Inference Systems—Part I: Classical and Quantum

- Formulations of Active Inference. *IEEE Transactions on Molecular, Biological and Multi-Scale Communications* 9, 235-245.
- Fields, C., Friston, K., Glazebrook, J.F., Levin, M., 2021a. A free energy principle for generic quantum systems, p. arXiv:2112.15242.
- Fields, C., Glazebrook, J.F., Levin, M., 2021b. Minimal physicalism as a scale-free substrate for cognition and consciousness. *Neuroscience of consciousness* 2021, niab013.
- Fields, C., Glazebrook, J.F., Levin, M., 2022. Neurons as hierarchies of quantum reference frames, p. arXiv:2201.00921.
- Fields, C., Glazebrook, J.F., Marciano, A., 2023b. Communication protocols and quantum error-correcting codes from the perspective of topological quantum field theory, p. arXiv:2303.16461.
- Friedman, D.A., Johnson, B.R., Linksvayer, T.A., 2020. Distributed physiology and the molecular basis of social life in eusocial insects. *Horm Behav* 122, 104757.
- Friedman, D.A., Tschantz, A., Ramstead, M.J.D., Friston, K., Constant, A., 2021. Active Inferants: An Active Inference Framework for Ant Colony Behavior. *Frontiers in Behavioral Neuroscience* 15.
- Friston, K., Da Costa, L., Sakthivadivel, D.A.R., Heins, C., Pavliotis, G.A., Ramstead, M., Parr, T., 2022a. Path integrals, particular kinds, and strange things, p. arXiv:2210.12761.
- Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., O'Doherty, J., Pezzulo, G., 2016. Active inference and learning. *Neuroscience and biobehavioral reviews* 68, 862-879.
- Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., Pezzulo, G., 2017a. Active Inference: A Process Theory. *Neural Comput* 29, 1-49.
- Friston, K., Friedman, D.A., Constant, A., Bleu Knight, V., Parr, T., Campbell, J.O., 2023a. A variational synthesis of evolutionary and developmental dynamics, p. arXiv:2303.04898.
- Friston, K., Friedman, D.A., Constant, A., Knight, V.B., Fields, C., Parr, T., Campbell, J.O., 2023b. A Variational Synthesis of Evolutionary and Developmental Dynamics. *Entropy* 25, 964.
- Friston, K., Mattout, J., Kilner, J., 2011. Action understanding and active inference. *Biological cybernetics* 104, 137-160.
- Friston, K., Parr, T., Zeidman, P., 2018. Bayesian model reduction. arXiv preprint arXiv:1805.07092.
- Friston, K., Penny, W., 2011. Post hoc Bayesian model selection. *Neuroimage* 56, 2089-2099.
- Friston, K., Rigoli, F., Ognibene, D., Mathys, C., Fitzgerald, T., Pezzulo, G., 2015. Active inference and epistemic value. *Cogn Neurosci*, 1-28.
- Friston, K., Schwartenbeck, P., FitzGerald, T., Moutoussis, M., Behrens, T., Dolan, R.J., 2014. The anatomy of choice: dopamine and decision-making. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences* 369.
- Friston, K.J., Frith, C.D., 2015. Active inference, communication and hermeneutics. *Cortex; a journal devoted to the study of the nervous system and behavior* 68, 129-143.
- Friston, K.J., Parr, T., de Vries, B., 2017b. The graphical brain: Belief propagation and active inference. *Network neuroscience (Cambridge, Mass.)* 1, 381-414.
- Friston, K.J., Parr, T., Yufik, Y., Sajid, N., Price, C.J., Holmes, E., 2020. Generative models, linguistic communication and active inference. *Neuroscience and biobehavioral reviews* 118, 42-64.
- Friston, K.J., Ramstead, M.J.D., Kiefer, A.B., Tschantz, A., Buckley, C.L., Albarracin, M., Pitliya, R.J., Heins, C., Klein, B., Millidge, B., Sakthivadivel, D.A.R., St Clere Smithe, T., Koudahl, M., Essafi Tremblay, S., Petersen, C., Fung, K., Fox, J.G., Swanson, S., Mapes, D., René, G., 2022b. Designing Ecosystems of Intelligence from First Principles, p. arXiv:2212.01354.



- Friston, K.J., Rosch, R., Parr, T., Price, C., Bowman, H., 2017c. Deep temporal models and active inference. *Neuroscience and biobehavioral reviews* 77, 388-402.
- Friston, K.J., Sajid, N., Quiroga-Martinez, D.R., Parr, T., Price, C.J., Holmes, E., 2021. Active listening. *Hearing research* 399, 107998.
- Frith, C., 2010. What is consciousness for? *Pragmatics & Cognition* 18, 497-551.
- Frith, C., Wentzer, T., 2013. Neural Hermeneutics, in: Kaldis, B. (Ed.), *Encyclopedia of Philosophy and the Social Sciences*. SAGE publications, Inc, Thousand Oaks.
- Garrod, S., Pickering, M.J., 2009. Joint action, interactive alignment, and dialog. *Top Cogn Sci* 1, 292-304.
- George, D., Hawkins, J., 2009. Towards a mathematical theory of cortical micro-circuits. *PLoS Comput Biol* 5, e1000532.
- Gershenson, C., 2012. Guiding the self-organization of random Boolean networks. *Theory in biosciences = Theorie in den Biowissenschaften* 131, 181-191.
- Gershman, S.J., Niv, Y., 2010. Learning latent structure: carving nature at its joints. *Curr Opin Neurobiol* 20, 251-256.
- Ghavamzadeh, M., Mannor, S., Pineau, J., Tamar, A., 2016. Bayesian Reinforcement Learning: A Survey. *arXiv arXiv:1609.04436*.
- Ghazanfar, A.A., Takahashi, D.Y., 2014. The evolution of speech: vision, rhythm, cooperation. *Trends Cogn Sci* 18, 543-553.
- Gill, S., Bierema, A.M.K., 2013. On the Meaning of Alarm Calls: A Review of Functional Reference in Avian Alarm Calling. *Ethology* 119, 449-461.
- Goyal, P., Dollár, P., Girshick, R., Noordhuis, P., Wesolowski, L., Kyrola, A., Tulloch, A., Jia, Y., He, K., 2017. Accurate, Large Minibatch SGD: Training ImageNet in 1 Hour, p. *arXiv:1706.02677*.
- Gros, C., 2009. Cognitive Computation with Autonomously Active Neural Networks: An Emerging Field. *Cognitive Computation* 1, 77-90.
- Han, X., Hui, C., 2014. Niche construction on environmental gradients: the formation of fitness valley and stratified genotypic distributions. *PLoS One* 9, e99775.
- Harnad, S., 1990. The Symbol Grounding Problem. *Physica D* 42, 335-346.
- 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.
- Hein, A.M., Rosenthal, S.B., Hagstrom, G.I., Berdahl, A., Torney, C.J., Couzin, I.D., 2015. The evolution of distributed sensing and collective computation in animal populations. *eLife* 4, e10955.
- Heins, C., Klein, B., Demekas, D., Aguilera, M., Buckley, C.L., 2023. Spin Glass Systems as Collective Active Inference, in: Buckley, C.L., Cialfi, D., Lanillos, P., Ramstead, M., Sajid, N., Shimazaki, H., Verbelen, T. (Eds.), *Active Inference*. Springer Nature Switzerland, Cham, pp. 75-98.
- Heyes, C.M., 2018. Cognitive gadgets : the cultural evolution of thinking.
- Heyes, C.M., Frith, C.D., 2014. The cultural evolution of mind reading. *Science* 344, 1243091.
- Higgins, I., Chang, L., Langston, V., Hassabis, D., Summerfield, C., Tsao, D., Botvinick, M., 2021. Unsupervised deep learning identifies semantic disentanglement in single inferotemporal face patch neurons. *Nat Commun* 12, 6456.
- Hohwy, J., 2016. The Self-Evidencing Brain. *Nous* 50, 259-285.
- Howard, R., 1966. Information Value Theory. *IEEE Transactions on Systems, Science and Cybernetics SSC-2*, 22-26.
- Huber, R., Ghilardi, M.F., Massimini, M., Tononi, G., 2004. Local sleep and learning. *Nature* 430, 78-81.
- Hunt, B., Ott, E., Yorke, J., 1997. Differentiable synchronisation of chaos. *Phys Rev E* 55, 4029-4034.

- Ishii, S., Yoshida, W., Yoshimoto, J., 2002. Control of exploitation-exploration meta-parameter in reinforcement learning. *Neural Netw* 15, 665-687.
- Isomura, T., Parr, T., Friston, K., 2019. Bayesian Filtering with Multiple Internal Models: Toward a Theory of Social Intelligence. *Neural Comput* 31, 2390-2431.
- Itti, L., Baldi, P., 2009. Bayesian Surprise Attracts Human Attention. *Vision Res.* 49, 1295-1306.
- Jafri, H.H., Singh, R.K., Ramaswamy, R., 2016. Generalized synchrony of coupled stochastic processes with multiplicative noise. *Phys Rev E* 94, 052216.
- Jardri, R., Deneve, S., 2013. Circular inferences in schizophrenia. *Brain* 136, 3227-3241.
- Jeffery, K., Pollack, R., Rovelli, C., 2019. On the Statistical Mechanics of Life: Schrodinger Revisited. *Entropy* 21, 1-21.
- Kairouz, P., McMahan, H.B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A.N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., D'Oliveira, R.G.L., Eichner, H., El Rouayheb, S., Evans, D., Gardner, J., Garrett, Z., Gascon, A., Ghazi, B., Gibbons, P.B., Gruteser, M., Harchaoui, Z., He, C.Y., He, L., Huo, Z.Y., Hutchinson, B., Hsu, J., Jaggi, M., Javidi, T., Joshi, G., Khodak, M., Konecny, J., Korolova, A., Koushanfar, F., Koyejo, S., Lepoint, T., Liu, Y., Mittal, P., Mohri, M., Nock, R., Olgun, A., Pagh, R., Qi, H., Ramage, D., Raskar, R., Raykova, M., Song, D., Song, W.K., Stich, S.U., Sun, Z.T., Suresh, A.T., Tramer, F., Vepakomma, P., Wang, J.Y., Xiong, L., Xu, Z., Yang, Q., Yu, F.X., Yu, H., Zhao, S., 2021. Advances and Open Problems in Federated Learning. *Foundations and Trends in Machine Learning* 14, 1-210.
- Kass, R.E., Raftery, A.E., 1995. Bayes Factors. *Journal of the American Statistical Association* 90, 773-795.
- Kastel, N., Hesp, C., 2021. Ideas Worth Spreading: A Free Energy Proposal for Cumulative Cultural Dynamics, in: Kamp, M., Koprinska, I., Bibal, A., Bouadi, T., Frénay, B., Galárraga, L., Oramas, J., Adilova, L., Krishnamurthy, Y., Kang, B., LARGERON, C., Lijffijt, J., Viard, T., Welke, P., Ruocco, M., Aune, E., Gallicchio, C., Schiele, G., Pernkopf, F., Blott, M., Fröning, H., Schindler, G., Guidotti, R., Monreale, A., Rinzivillo, S., Biecek, P., Ntoutsis, E., Pechenizkiy, M., Rosenhahn, B., Buckley, C., Cialfi, D., Lanillos, P., Ramstead, M., Verbelen, T., Ferreira, P.M., Andresini, G., Malerba, D., Medeiros, I., Fournier-Viger, P., Nawaz, M.S., Ventura, S., Sun, M., Zhou, M., Bitetta, V., Bordino, I., Ferretti, A., Gullo, F., Ponti, G., Severini, L., Ribeiro, R., Gama, J., Gavaldà, R., Cooper, L., Ghazaleh, N., Richiardi, J., Roqueiro, D., Saldana Miranda, D., Sechidis, K., Graça, G. (Eds.), *Machine Learning and Principles and Practice of Knowledge Discovery in Databases*. Springer International Publishing, Cham, pp. 784-798.
- Kauffman, S.A., Johnsen, S., 1991. Coevolution to the edge of chaos: coupled fitness landscapes, poised states, and coevolutionary avalanches. *Journal of theoretical biology* 149, 467-505.
- Kaufmann, R., Gupta, P., Taylor, J., 2021. An Active Inference Model of Collective Intelligence. *Entropy (Basel)* 23, 830.
- Kilner, J.M., Friston, K.J., Frith, C.D., 2007. Predictive coding: an account of the mirror neuron system. *Cognitive processing* 8, 159-166.
- Klyubin, A.S., Polani, D., Nehaniv, C.L., 2005. Empowerment: A Universal Agent-Centric Measure of Control. In *Proc. CEC 2005*. IEEE 1, 128-135.
- Knill, E., Laflamme, R., 1997. Theory of quantum error-correcting codes. *Physical Review A* 55, 900-911.
- Krafft, P.M., Shmueli, E., Griffiths, T.L., Tenenbaum, J.B., Pentland, A.S., 2021. Bayesian collective learning emerges from heuristic social learning. *Cognition* 212, 104469.
- Krakauer, D.C., Page, K.M., Erwin, D.H., 2009. Diversity, dilemmas, and monopolies of niche construction. *Am Nat* 173, 26-40.

- Kschischang, F.R., Frey, B.J., Loeliger, H.A., 2001. Factor graphs and the sum-product algorithm. *Ieee T Inform Theory* 47, 498-519.
- Laland, K., Matthews, B., Feldman, M.W., 2016. An introduction to niche construction theory. *Evolutionary ecology* 30, 191-202.
- Laland, K.N., Odling-Smee, F.J., Feldman, M.W., 1999. Evolutionary consequences of niche construction and their implications for ecology. *Proceedings of the National Academy of Sciences of the United States of America* 96, 10242-10247.
- Laughlin, S.B., 2001. Efficiency and complexity in neural coding. *Novartis Found Symp.* 239, 177-187.
- Lehmann, L., 2008. The adaptive dynamics of niche constructing traits in spatially subdivided populations: evolving posthumous extended phenotypes. *Evolution* 62, 549-566.
- Levin, M., 2019. The Computational Boundary of a "Self": Developmental Bioelectricity Drives Multicellularity and Scale-Free Cognition. *Front Psychol* 10, 2688.
- Limanowski, J., 2017. (Dis-)Attending to the Body, in: Metzinger, T.K., Wiese, W. (Eds.), *Philosophy and Predictive Processing*. MIND Group, Frankfurt am Main.
- Lindley, D.V., 1956. On a Measure of the Information Provided by an Experiment. *Annals of Mathematical Statistics* 27, 986-1005.
- Linsker, R., 1990a. Perceptual Neural Organization - Some Approaches Based on Network Models and Information-Theory. *Annual Review of Neuroscience* 13, 257-281.
- Linsker, R., 1990b. Perceptual neural organization: some approaches based on network models and information theory. *Annu Rev Neurosci.* 13, 257-281.
- Mackay, D.J.C., 1992. Information-Based Objective Functions for Active Data Selection. *Neural Computation* 4, 590-604.
- MacKay, D.J.C., Peto, L.C.B., 2008. A hierarchical Dirichlet language model. *Natural Language Engineering* 1, 289-308.
- Manrique, H.M., Walker, M.J., 2023. To copy or not to copy? That is the question! From chimpanzees to the foundation of human technological culture. *Phys Life Rev* 45, 6-24.
- Meder, B., Nelson, J.D., 2012. Information search with situation-specific reward functions. *Judgment and Decision Making* 7, 119-148.
- Moens, V., Zenon, A., 2019. Learning and forgetting using reinforced Bayesian change detection. *PLoS Comput Biol* 15, e1006713.
- Namikawa, J., 2005. Chaotic itinerancy and power-law residence time distribution in stochastic dynamical systems. *Physical review. E, Statistical, nonlinear, and soft matter physics* 72, 026204.
- Navarro, D.J., Perfors, A.F., 2011. Hypothesis generation, sparse categories, and the positive test strategy. *Psychol Rev* 118, 120-134.
- Nelson, J.D., McKenzie, C.R., Cottrell, G.W., Sejnowski, T.J., 2010. Experience matters: information acquisition optimizes probability gain. *Psychol Sci* 21, 960-969.
- Odling-Smee, F.J., Laland, K.N., Feldman, M.W., 2013. *Niche Construction*.
- Olshausen, B.A., Field, D.J., 1996a. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* 381, 607-609.
- Olshausen, B.A., Field, D.J., 1996b. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* 381, 607-609.
- OpenAI, Berner, C., Brockman, G., Chan, B., Cheung, V., Debiak, P., Dennison, C., Farhi, D., Fischer, Q., Hashme, S., Hesse, C., Józefowicz, R., Gray, S., Olsson, C., Pachocki, J., Petrov, M., Pinto, H.P.d.O., Raiman, J., Salimans, T., Schlatter, J., Schneider, J., Sidor, S., Sutskever, I., Tang, J., Wolski, F., Zhang, S., 2019. Dota 2 with Large Scale Deep Reinforcement Learning, p. arXiv:1912.06680.
- Ortega, P.A., Braun, D.A., 2013. Thermodynamics as a theory of decision-making with information-processing costs. *Proc. R. Soc. A* 469 2153.

- Palacios, E.R., Razi, A., Parr, T., Kirchhoff, M., Friston, K., 2017. Biological Self-organisation and Markov blankets. *bioRxiv*.
- Parr, T., Da Costa, L., Friston, K., 2020. Markov blankets, information geometry and stochastic thermodynamics. *Philosophical transactions. Series A, Mathematical, physical, and engineering sciences* 378, 20190159.
- Parr, T., Friston, K.J., 2018. The Anatomy of Inference: Generative Models and Brain Structure. *Frontiers in computational neuroscience* 12, 90.
- Parr, T., Friston, K.J., 2019. Generalised free energy and active inference. *Biological cybernetics* 113, 495-513.
- Parr, T., Markovic, D., Kiebel, S.J., Friston, K.J., 2019. Neuronal message passing using Mean-field, Bethe, and Marginal approximations. *Scientific reports* 9, 1889.
- Parr, T., Pezzulo, G., Friston, K.J., 2022. Active Inference: The Free Energy Principle in Mind, Brain, and Behavior. MIT Press, Cambridge.
- Pellet, J.P., Elisseeff, A., 2008. Using Markov blankets for causal structure learning. *Journal of Machine Learning Research* 9, 1295-1342.
- Pérez-Escudero, A., de Polavieja, G.G., 2011. Collective animal behavior from Bayesian estimation and probability matching. *PLoS Comput Biol* 7, e1002282.
- Phong, L.T., Aono, Y., Hayashi, T., Wang, L.H., Moriai, S., 2018. Privacy-Preserving Deep Learning via Additively Homomorphic Encryption. *Ieee Transactions on Information Forensics and Security* 13, 1333-1345.
- Poel, W., Daniels, B.C., Sosna, M.M.G., Twomey, C.R., Leblanc, S.P., Couzin, I.D., Romanczuk, P., 2022. Subcritical escape waves in schooling fish. *Sci Adv* 8, eabm6385.
- Pratt, S.C., Sumpter, D.J., 2006. A tunable algorithm for collective decision-making. *Proc Natl Acad Sci U S A* 103, 15906-15910.
- Prince, A., Smolensky, P., 2007. Optimality Theory : Constraint Interaction in Generative Grammar. John Wiley & Sons Chichester, Chichester.
- Ramstead, M.J.D., Sakthivadivel, D.A.R., Heins, C., Koudahl, M., Millidge, B., Da Costa, L., Klein, B., Friston, K.J., 2022. On Bayesian Mechanics: A Physics of and by Beliefs, p. arXiv:2205.11543.
- Rieucou, G., Giraldeau, L.A., 2011. Exploring the costs and benefits of social information use: an appraisal of current experimental evidence. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences* 366, 949-957.
- Saegusa, R., Metta, G., Sandini, G., Sakka, S., 2009. Active Motor Babbling for Sensorimotor Learning, *IEEE International Conference on Robotics and Biomimetics (ROBIO)*, Bangkok, THAILAND, pp. 794-+.
- Sakthivadivel, D.A.R., 2022a. A Constraint Geometry for Inference and Integration, p. arXiv:2203.08119.
- Sakthivadivel, D.A.R., 2022b. Weak Markov Blankets in High-Dimensional, Sparsely-Coupled Random Dynamical Systems, p. arXiv:2207.07620.
- Sanchez, E.H., Serrurier, M., Ortner, M., 2019. Learning Disentangled Representations via Mutual Information Estimation, p. arXiv:1912.03915.
- Savage, L.J., 1954. *The Foundations of Statistics*. Wiley, New York.
- Schmidhuber, J., 1991. Curious model-building control systems. In *Proc. International Joint Conference on Neural Networks*, Singapore. *IEEE* 2, 1458-1463.
- Schmidhuber, J., 2010. Formal Theory of Creativity, Fun, and Intrinsic Motivation (1990-2010). *Ieee Transactions on Autonomous Mental Development* 2, 230-247.
- Schwartenbeck, P., Passecker, J., Hauser, T.U., FitzGerald, T.H., Kronbichler, M., Friston, K.J., 2019. Computational mechanisms of curiosity and goal-directed exploration. *eLife* 8, e41703.

- Shoeybi, M., Patwary, M., Puri, R., LeGresley, P., Casper, J., Catanzaro, B., 2019. Megatron-LM: Training Multi-Billion Parameter Language Models Using Model Parallelism, p. arXiv:1909.08053.
- Shokri, R., Shmatikov, V., 2015. Privacy-preserving deep learning, 2015 53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton). IEEE Press, Monticello, IL, USA, pp. 909-910.
- Smith, R., Schwartenbeck, P., Parr, T., Friston, K.J., 2020. An Active Inference Approach to Modeling Structure Learning: Concept Learning as an Example Case. *Frontiers in computational neuroscience* 14, 41.
- Sosna, M., Twomey, C., Bak-Coleman, J., Poel, W., Daniels, B., Romanczuk, P., Couzin, I., 2019. Individual and collective encoding of risk in animal groups. *Proceedings of the National Academy of Sciences* 116, 201905585.
- Spielman, D.A., Srivastava, N., 2011. Graph Sparsification by Effective Resistances. *Siam Journal on Computing* 40, 1913-1926.
- Steels, L., 2011. Modeling the cultural evolution of language. *Phys Life Rev* 8, 339-356.
- Stickgold, R., Walker, M.P., 2007. Sleep-dependent memory consolidation and reconsolidation. *Sleep Med* 8, 331-343.
- Still, S., Precup, D., 2012. An information-theoretic approach to curiosity-driven reinforcement learning. *Theory in biosciences = Theorie in den Biowissenschaften* 131, 139-148.
- Stoianov, I., Maisto, D., Pezzulo, G., 2022. The hippocampal formation as a hierarchical generative model supporting generative replay and continual learning. *Prog Neurobiol* 217, 102329.
- Stotz, K., 2017. Why developmental niche construction is not selective niche construction: and why it matters. *Interface focus* 7, 20160157.
- Strandburg-Peshkin, A., Farine, D.R., Couzin, I.D., Crofoot, M.C., 2015. GROUP DECISIONS. Shared decision-making drives collective movement in wild baboons. *Science* 348, 1358-1361.
- Suzuki, T.N., 2014. Communication about predator type by a bird using discrete, graded and combinatorial variation in alarm calls. *Animal Behaviour* 87, 59-65.
- Suzuki, T.N., Wheatcroft, D., Griesser, M., 2016. Experimental evidence for compositional syntax in bird calls. *Nat Commun* 7, 10986.
- 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.R., Tenenbaum, J.B., Gershman, S.J., 2016. Toward the neural implementation of structure learning. *Curr Opin Neurobiol* 37, 99-105.
- Tipping, M.E., 2001. Sparse Bayesian learning and the relevance vector machine. *Journal of Machine Learning Research* 1, 211-244.
- Tomasello, M., 2016. Cultural Learning Redux. *Child development* 87, 643-653.
- Torney, C.J., Lorenzi, T., Couzin, I.D., Levin, S.A., 2015. Social information use and the evolution of unresponsiveness in collective systems. *Journal of The Royal Society Interface* 12, 20140893.
- Toutounji, H., Pipa, G., 2014. Spatiotemporal computations of an excitable and plastic brain: neuronal plasticity leads to noise-robust and noise-constructive computations. *PLoS Comput Biol* 10, e1003512.
- Townsend, S.W., Manser, M.B., 2013. Functionally Referential Communication in Mammals: The Past, Present and the Future. *Ethology* 119, 1-11.
- van den Broek, J.L., Wiegerinck, W.A.J.J., Kappen, H.J., 2010. Risk-sensitive path integral control. *UAI* 6, 1-8.
- Van Dijk, S.G., Polani, D., 2013. Informational Constraints-Driven Organization in Goal-Directed Behavior. *Advances in Complex Systems* 16, 1350016.

- Vasil, J., Badcock, P.B., Constant, A., Friston, K., Ramstead, M.J.D., 2020. A World Unto Itself: Human Communication as Active Inference. *Front Psychol* 11, 417.
- Veissiere, S.P.L., Constant, A., Ramstead, M.J.D., Friston, K.J., Kirmayer, L.J., 2019. Thinking through other minds: A variational approach to cognition and culture. *The Behavioral and brain sciences* 43, e90.
- Verbraeken, J., Wolting, M., Katzy, J., Kloppenburg, J., Verbelen, T., Rellermeyer, J.S., 2020. A Survey on Distributed Machine Learning. *Acm Comput Surv* 53, Article 30.
- Vigorito, C.M., Barto, A.G., 2010. Intrinsically Motivated Hierarchical Skill Learning in Structured Environments. *Ieee Transactions on Autonomous Mental Development* 2, 132-143.
- Ward, A.J., Sumpter, D.J., Couzin, I.D., Hart, P.J., Krause, J., 2008. Quorum decision-making facilitates information transfer in fish shoals. *Proc Natl Acad Sci U S A* 105, 6948-6953.
- Williford, K., Bennequin, D., Friston, K., Rudrauf, D., 2018. The Projective Consciousness Model and Phenomenal Selfhood. *Front Psychol* 9, 2571.
- Wilson, M., Wilson, T.P., 2005. An oscillator model of the timing of turn-taking. *Psychon Bull Rev* 12, 957-968.
- Winn, J., Bishop, C.M., 2005. Variational message passing. *Journal of Machine Learning Research* 6, 661-694.
- Yanagihara, S., Yazaki-Sugiyama, Y., 2016. Auditory experience-dependent cortical circuit shaping for memory formation in bird song learning. *Nat Commun* 7, 11946.
- Yang, Q., Liu, Y., Chen, T.J., Tong, Y.X., 2019. Federated Machine Learning: Concept and Applications. *Acm T Intel Syst Tec* 10, Article 12.
- Yildiz, I.B., von Kriegstein, K., Kiebel, S.J., 2013. From birdsong to human speech recognition: bayesian inference on a hierarchy of nonlinear dynamical systems. *PLoS Comput Biol* 9, e1003219.
- Young, T., Hazarika, D., Poria, S., Cambria, E., 2018. Recent Trends in Deep Learning Based Natural Language Processing. *Ieee Computational Intelligence Magazine* 13, 55-75.
- Zorzi, M., Testolin, A., Stoianov, I.P., 2013. Modeling language and cognition with deep unsupervised learning: a tutorial overview. *Front Psychol* 4, 515.

## Highlights

- This paper offers a proof of concept that communication in general—and language in particular—is an emergent property of conspecifics that seek evidence for generative models of their shared, co-constructed world.
- It uses numerical studies to show that nested free energy minimising—or evidence maximising—processes are sufficient to explain (i) the (horizontal) emergence of language among agents that can broadcast their beliefs and (ii) its (vertical) transmission over generations.
- By grounding these processes in formal (Bayesian) terms, one can join various perspectives on communication; ranging from generalised synchrony in dynamical systems to cultural niche construction in evolutionary psychology.