

Article

Not peer-reviewed version

As One and Many: Relating Individual and Emergent Group-Level Generative Models in Active Inference

[Peter Thestrup Waade](#)*, Christoffer Lundbak Olesen, Jonathan Ehrenreich Laursen, Samuel William Nehrer, [Conor Heins](#), [Karl Friston](#), [Christoph Mathys](#)

Posted Date: 24 October 2024

doi: 10.20944/preprints202410.1895.v1

Keywords: Active Inference; Free Energy Principle; Markov Blanket; Predictive Processing; Cognitive Modelling; Multi-scale; Collective Intelligence; Emergence



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a Creative Commons CC BY 4.0 license, which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Article

As One and Many: Relating Individual and Emergent Group-Level Generative Models in Active Inference

Peter Thestrup Waade ^{1,*} , Christoffer Lundbak Olesen ¹, Jonathan Ehrenreich Laursen ², Samuel William Nehrer ², Conor Heins ³, Karl Friston ⁴ and Christoph Mathys ¹

¹ Interacting Minds Centre, Aarhus University

² School of Communication and Culture, Aarhus University

³ Department of Collective Behavior, Max Planck institute for Animal Behavior

⁴ Queen Sq, Institute of Neurology, University College London

* Correspondence: ptw@cas.au.dk

Abstract: Active inference under the Free Energy Principle has been proposed as an across-scales compatible framework for understanding and modelling behaviour and self-maintenance. Crucially, a collective of active inference agents can, if they maintain a group-level Markov Blanket, constitute a larger group-level active inference agent with a generative model of its own. This potential for computational, scale-free structures speaks to the application of active inference to self-organizing systems across spatiotemporal scales, from cells to human collectives. Due to the difficulty of reconstructing the generative model — that explains the behaviour — of emergent group-level agents, there has been little research on this kind of multi-scale active inference. Here, we propose a data-driven methodology for characterising the relation between the generative model of a group-level agent and the dynamics of its constituent individual agents. We apply methods from computational cognitive modelling and computational psychiatry, applicable for active inference as well as other types of modelling approaches. Using a simple Multi-Armed Bandit task as an example, we employ the new `ActiveInference.jl` library for Julia to simulate a collective of agents who are equipped with a Markov blanket. We use sampling-based parameter estimation to make inferences about the generative model of the group-level agent and we show that there is a non-trivial relationship between the generative models of individual agents and the group-level agent they constitute, even in this simple setting. Finally, we point to a number of ways in which this methodology might be applied to better understand the relations between nested active inference agents across scales.

Keywords: active inference; free energy principle; Markov Blanket; predictive processing; cognitive modelling; multi-scale; collective intelligence; emergence

1. Introduction

Active inference [1–3], under the Free Energy Principle [4–8] has been proposed as a formal framework generally applicable for modeling living and adaptive systems, across spatial and temporal scales of description [9]. From this perspective, perception, learning, action, and cognition in general can be described as approximate Bayesian inference driven by a single optimization process, contingent on a generative model of the environment. Crucially, a system can be considered to do active inference once it maintains a causal boundary - a Markov blanket. It has been suggested that collectives of agents individually performing active inference can maintain such a boundary, leading to a hierarchy of nested simultaneous processes of active inference [10,11]. This potential for nested applicability underwrites much of active inference's ubiquity, and several simulation studies have investigated the emergent processes which ensue. Despite this, very few studies have been able analytically to relate the generative model of a group-level active inference process to the generative models of the individual agents [12]. Here we propose a method for doing this, based on a kind of cognitive modelling inspired by computational psychiatry. In the following, we briefly introduce active inference and the Free Energy Principle, focusing on how it has been applied to multi-agent and multi-scale contexts. We then demonstrate our approach in a simple case, showing how emergent group-level active inference processes have non-trivial relations to their constituent parts. Finally, we discuss the applicability of this approach more broadly.

1.1. Active Inference and the Free Energy Principle

The Free Energy Principle [4–8] is the claim that any system that maintains a stable boundary - formalized as a Markov blanket - can be described as minimizing the variational free energy of its sensory states. A Markov blanket is a set of states that renders states internal to it conditionally independent of external states. Blanket states are further divided into sensory states, which affect — but are not affected by — the internal states, and active states, which affect — but are not affected by — external states. The ensuing separation of internal and external states by blanket states (i.e., conditional independence) is statistical, and not necessarily also causal [13]. In state space formulations, the Free Energy Principle has additionally rested on the assumption that the system is at a non-equilibrium steady state, which is unnecessary in path-integral formulations [8,14]. Maintaining a Markov blanket can be shown to imply a minimization of a negative log-probability (i.e., potential energy or self-information) - called the surprise $\mathfrak{S}$ - of the sensory states, given a generative model, that is, one of how states in the environment generate those states. Surprise is not computable in most cases of interest but can be approximated by a variational free energy upper bound; minimization of surprise then becomes feasible by minimizing the variational free energy. Minimizing variational free energy is equivalent to performing variational (approximate) Bayesian inference, which enables a description of systems that maintain Markov blankets as engaging in active inference (viz., self evidencing).

Active inference is the Free Energy Principle applied to perception, learning, and action. Here, perception is viewed as variational inference on the state of the environment, given sensory states and a generative model, which can itself be updated by Bayesian parameter learning and Bayesian model selection (viz., structure learning). Here, inference corresponds to optimizing a posterior distribution (with an assumed functional form) that has the lowest associated variational free energy given the sensory inputs, which can be shown to be the variational posterior that best approximates the true Bayesian posterior given the constraints imposed on its functional form (e.g., Gaussianity). Actions (or often policies - sequences of actions, with some specified temporal depth) are then selected that minimize the expected free energy; namely, the expectation of the free energy under predicted sensory outcomes, given those actions. This affords a view of action selection as a type of (planning as) inference. Perception and action thus constitute the two possible ways of minimizing free energy: by changing one's expectations, and by changing the world (and thereby the sensory observations). Active inference models come equipped with a preference prior - a prior expectation for the types of sensory states the agent is likely to encounter. The preference prior is usually immutable, which means that the only way to minimize free energy is to act such that the prior is most likely to be realized, and that it therefore encodes the preferred observations of the agent. Technically, these preferred states in general constitute the attracting states of the agent's dynamics; namely, the kind of states that are characteristic of the agent at hand.

Active inference can be applied as a behavioural model irrespective of the Free Energy Principle. It has been applied in a variety of fields, either to understand or to build adaptive systems, including theoretical neurobiology and neuroscience [15], cognitive science and computational psychiatry [16,17], robotics and machine learning [18], and philosophy of mind [19,20], and it generalizes a variety of related approaches, like reinforcement learning, KL-control or expected utility maximization [1,21]. Whether active inference is used to model observed systems or build artificial ones, in order to apply it in any specific context, the first task is to specify a generative model, which provides the constraints that a resulting (active) inference process obeys (i.e., specifies the attracting set of states that characterize the agent). Many types of generative models are used in active inference, including continuous as well as discrete state-space models, and models made for specific contexts as well as more generally applicable models. Here, we specifically employ a widely used discrete state space model, the Partially Observable Markov Decision Process (POMDP) (see Section 2.1 for a full description).

Active inference under the Free Energy Principle is sometimes presented as taking place in nested Markov blanket structures, where smaller agents compose larger-scale blankets that become agents in their own right [10,22]. These group-agents can then in turn be part of even larger-scale

blanket structures - like cells forming organs, which in turn form human bodies and eventually human collectives. In this view, an obvious question to ask is how the active inference agents at higher levels relate to the activity and interactions of the smaller-scale agents that constitute them. We here demonstrate a method for investigating exactly this, based on cognitive modelling. In the following section, we preface the introduction of this method with an overview of the literature to date of active inference in multi-agents settings.

1.2. Multi-Agent and Collective Active Inference

Active Inference has since its inception been applied to “social” contexts; that is, contexts where multiple agents interact (e.g., [23,24]). Work on multi-agent active inference traditionally unfolds on one or more of three scales of description: a *within-agent scale*, where the focus is to investigate what kinds of generative models are appropriate for interacting with environments containing other agents; a *between-agent scale*, where the focus is to understand how interactions between multiple agents — ranging in number from dyads to whole populations — mutually shape their behavioural and belief dynamics over time; and a *group-as-agent scale*, where a collective of agents forms an emergent group that possesses a Markov blanket of its own, and which therefore instantiates an active inference agent in its own right. Work has been done on all three scales, but we are not aware of any work that manages to reconstruct the generative model of an emergent group-level agent, or to compare it with the dynamics of its constituent agents at the two lower levels. In the following, we give a brief overview of the literature on multi-agent active inference across the three levels, and proceed to describe how the work presented here complements previous work with a method for accessing the generative models of group-level agents.

A number of theoretical points have been made regarding the types of generative models that must be held by social agents to function. It has been pointed out, for example, that the main statistical regularities in the environment - the parts that are most important to represent properly in a generative model - are social, specifically regarding other agents’ expectations for one’s own behaviour [25,26]. These expectations, which determine the most appropriate way to act in a given situation, can be inferred by attending to socially constructed cues (or “cultural affordances”) in the environment [27] - cues that provide (and are created to provide) what has been called “deontic value” in being informative about obligatory social rules [28].

It is possible to make inferences about the actions of others by selecting among explicit models of the mental processes underlying their behaviour [29], and let it affect one’s behaviour in collaborative tasks [30] - or even to use recursive Theory of Mind models where the level of recursion has to be explicitly limited [31,32]. This allows for explicit perspective-taking, and for interacting with agents different from oneself, but is a complex and computationally costly process that may often not be necessary for coordination. It is argued that humans have an evolved prior preferences for interacting with others whom they are mentally aligned with, that is, with whom they have similar expectations about the shared environment [33] - “shared protentions” in Husserlian terminology [34] - which facilitates communication and coordination in its own right. This is built on canonical simulation work showing that agents can, if they have similar generative models (also called a “shared narrative”), make inferences about each others’ mental states without needing to explicitly model the (infinitely recursive) mutual perspective-taking of the interaction [24], and that these linked active inference agents reach a free energy minimum over time by aligning their generative models [35]. The resulting generalized synchrony - if it is successfully instantiated, which is not always the case [36] - allows agents to communicate their beliefs about the environment to each other [37], and acquire a shared language that allows them to combine knowledge from complementary perspectives (a type of distributed intelligence called “federated inference”) to, for example, track a moving target (in order to hide from it) [38]. This type of generalized synchrony of strategies and belief states can also underpin the coordination of goal-oriented joint action for dyads [39] and teams of agents [40], or be a mechanism determining behaviour in competitive games like the prisoner’s dilemma [41].

On a multi-agent population level, the emergence of generalized synchrony - and the preference for being aligned with interlocutors - is observed in *in vivo* neurons [42–44], and can be a mechanism for implementing cumulative culture [45,46] which can also lead to separate echo-chamber-like epistemic communities that maintain highly precise and difficult-to-change beliefs [47].

The work mentioned above provides important clues for how successful social interaction can be underpinned by active inference - for biological as well as artificial and mixed intelligence systems [48]. It does not, however, engage directly with the proposed multi-scale nature of active inference, because it remains at a within- and between-agent level of description. There is also a strand of work regarding how a collective of (potentially active inference) agents can form an emergent whole with a Markov blanket of its own - which then, of course, can be considered an active inference agent in its own right. The canonical work here is a 'primordial soup' simulation, where it is shown that a system maintaining a Markov blanket leads to the system's internal states carrying information about external states as they come to minimize a variational free energy functional of the blanket states, and therefore model the environment [23]. One way to establish this type of boundary (as in the case of organic morphogenesis) is for the members of the collective to be equipped with a prior expectation of being part of such a structure [49], with cells creating and maintaining the larger structure (as well as differentiating into different roles, or cell types) in the process of reducing free energy [50]. These types of maintained Markov blanket structures can be nested within each other [22], in ways that are found in the brain [51], and which can be related to mental phenomena like the psychopathology of the emergent human mind [52]. There is also work on emergent Markov blankets not related to the brain, and which does not rely on having explicit prior expectations for being in a specific larger structure. Joint free energy minimization can be a mechanism for ant-like agents collectively solving a T-maze [53], or for the self-organization of collective motion, whereby single particles can collectively synchronize their movements by only maintaining a goal prior over simple metrics like relative distance to neighbours [54]. Dyads of agents moving to target locations are designed to form collectives that could be considered to perform Bayesian inference, relative to certain sensory states [55], and certain parameter regimes of spin-glass systems can be analytically related to a collective of active inference agents, whose joint actions implement Bayesian inference at a higher level (namely, via a form of sampling-based inference) [12].

The above work engages with how collectives of agents can come to form larger-scale Markov blanket structures, how they come to solve tasks as a group, and how to relate the dynamics of the agents — forming the internal states of the collective — to Bayesian inference happening at the collective level, all implying that there is an active inference process instantiated at the collective or group level. Despite this, previous studies never explicitly model the inversion of the emergent group's generative model and corresponding action selection, and therefore never consider generative models at multiple levels simultaneously. That is, they largely address inference at within- and between-agent levels, only hinting at group-level active inference proper. The main exception here is the spin glass simulation work by Heins and colleagues [12], where a group-level generative model is analytically related to the activity of constituent agents, although in a way limited to a relatively specific context, and without direct relation to an environment. Fully engaging with a multi-scale active inference account — and the particularly interesting question how the dynamics of the constituent agents relate to the group-level active inference process (for example in the case of psychopathology [52]) — should entail relating the generative models at the different levels to each other (since they define the active inference process), and to the environment.

There are various reasons why this has not been addressed. Firstly, there is not always a well-defined environment to act on and be influenced by *per se* (as with the vacuous environment of self-organizing cells [49] and cooperating dyads [55]), or there might not be easily differentiated internal and blanket states (as with the collaborative ants [53], emergent fish school [54] or jointly vigilant girls hiding from their mother [38]) to implement the inference. A more important reason, however, is probably the fact that the generative model of the group-level agent is difficult to access.

As opposed to the generative models of the constituent agents, which are specified by the modellers and therefore known, the generative model of the group is emergent and *a priori* unknown. This means it has to be reconstructed in order to be compared to the constituent agents, which is not a trivial task. There do, however, exist some standard methods for attempting to reconstruct unknown generative models of systems. One such approach comes from cognitive modelling, a field where the primary task is to infer on the unknown generative models underlying observable behaviour, often of human subjects. From the perspective of active inference, observable behaviour is simply the blanket states of the system. This fits well with cognitive and behavioural modelling, where a Markov blanket-like structure is usually assumed. Here, we use this approach to make inferences about the generative model of a group-level agent. This group-level agent comprises a set of agents set in a Markov blanket structure (we here do not engage with *how* the agents come to be structured in a Markov blanket. Instead, for simplicity, we assume that this process has taken place) and with a simple internal structure and external environment. This lets us use the same generative model structure for the individual agents and the group-level agent. After confirming that the proposed method works for this specific simulation, we go on to make preliminary investigations of the relations between group- and individual-level generative models. The approach we present here should be applicable in many contexts, including when boundaries are dynamically established, or when internal structures and external environments are more complex.

2. Materials and Methods

We introduce a methodology for relating generative models and ensuing active inference across levels of nested Markov blankets, based on fitting active inference models to the behaviour of the group-level agent, in the fashion of computational cognitive modelling. The method is applicable as long as: (i) there are clearly definable blanket states (observations and actions) for the group-level agent at each time point; (ii) one or more appropriate generative model structures for the environment of the group-level agent can be identified; and (iii) the resulting behaviour (i.e. blanket states) of the group-level agent is informative, that is, different types of generative models (or parameters of generative models) can be distinguished based on the behaviour. In our illustrative case, we ensure the first condition by composing agents in a fixed Markov blanket structure, and the second and third by choosing a group-level environment for which there are existing well-defined generative models: namely, the Multi-Armed Bandit task (MAB). In addition, this allows us to use the same generative model for the internal agents of the group as for the group level, which means that we can validate the method by confirming *a priori* expected relations. That being said, the method is also in principle applicable with emergent Markov blankets, and with more complex environments.

In the following, we describe the multi-armed bandit task active inference model we use (Section 2.1), as well as the construction of the group-level agent (Section 2.2). Finally, we report the numerical experiments we used as preliminary investigations into the relations between the two levels of description (Section 2.3). All simulations were performed in Julia (version 1.10) [56], relying primarily on a new library for simulating and fitting active inference models, `ActiveInference.jl` (version 0.2), as well as its sister library for fitting behavioural models to data, `ActionModels.jl` (version 0.5), itself an extension to the powerful `Turing.jl` library (version 0.10) [57] for Bayesian parameter estimation. All code and synthesized data is publicly available at osf.io/6z3bd.

2.1. Active Inference and Multi-Armed Bandits

The Multi-Armed Bandit (MAB) task is a ubiquitous task in cognitive science and neuroscience as well as in machine learning [58]. It is simple to design and model, but can stand in for a wide range of environments that depend probabilistically on an agent's actions. As one of the simplest possible instantiations of a complete action-perception loop, it is suitable for establishing the method used here. A MAB task presents an agent (human or artificial) with an environment comprising a "bandit" with multiple arms (in our case three) which the agent can, in a trial-by-trial fashion, select among.

Choosing a bandit arm probabilistically results in an outcome: here, we use binary outcomes, but they can in principle also be continuous or categorical. Traditionally, outcomes are referred to as “wins” and “losses”, and the agent’s task is to maximize the wins and minimize the losses it receives. In active inference, the value of a specific outcome is defined by a subjective preference which is specific to the agent: we therefore simply call them outcomes 1 and 2.

Active inference has already been applied to MAB tasks [59], albeit with a specialized generative model. We here instead use a classic model type that is widely used in the active inference literature, the Partially Observable Markov Decision Process (POMDP) [1] (Figure 1).




Figure 1. The generative model used for active inference based on a Partially Observable Markov Decision Process. Observations o depend on hidden states s according to the observation model A . s changes over time according to the transition model B and the chosen policy π . C is the preference prior which determines the expected free energy G . D and E provide priors for the hidden states and the policies, respectively.

The POMDP is a flexible type of generative model for discrete state spaces in discrete time. POMDP-based generative models for active inference assume that at a given time t , a discrete observation o_t depends on concurrent discrete hidden states s_t , which evolve over time, conditioned on a policy π - a sequence of discrete actions a_t - as parametrized by five matrices A - E . The dependence of o_t on s_t is governed by the A matrix, sometimes called the observation model, which specifies the likelihood of making observations o_t given each possible state of the environment s_t . The dynamics in s_t is governed by the B matrix, sometimes called the transition model, which is a transition probability matrix over the possible states of s_t , given previous hidden states s_{t-1} . The preference prior is encoded by the C matrix, also called the *prior preferences* or the *goal prior*, which encodes a fixed *a priori* expectation over sensory outcomes and thereby define the agent’s preferences. The D matrix defines the agent’s prior over environmental states s_t , and the E matrix defines the agent’s prior over policies π , in effect defining habits that the agent returns to in the absence of other action incentives. The actions a_t of the policy π select between different possible B matrices, thereby controlling the evolution of environmental states s_t . The variational posterior $q(s)$ is constructed by gradient descent on a variational free energy $\mathcal{F}$, which can be rewritten as the sum of the surprise $\mathfrak{S}$ and the divergence between the approximate and the true posterior (Equation (1)). Since changing it can only affect the divergence, the posterior $q(s)$ with the minimal free energy $\mathcal{F}$ is the best approximation to the true posterior.

$$\mathcal{F}[q(s), o] = \underbrace{D_{\text{KL}}[q(s) \| p(s|o)]}_{\text{Divergence}} - \underbrace{\ln p(o)}_{\text{Surprise}} \quad (1)$$

Policies π are then selected to minimize the expected free energy G , which is a combination of the expected information gain and the *pragmatic value* which quantifies the how much observations are in accordance with the preference prior (Equation (2)).

$$G_{\pi} = - \underbrace{\mathbb{E}_{q(\tilde{o}, \tilde{s}|\pi)} [\ln q(\tilde{s}|\tilde{o}, \pi) - \ln q(\tilde{s}|\pi)]}_{\text{Information gain}} - \underbrace{\mathbb{E}_{q(\tilde{o}|\pi)} [\ln p(\tilde{o}|C)]}_{\text{Pragmatic value}} \quad (2)$$

The relative expected free energies forms the posterior over policies G , which is combined with the prior over policies E to form the final probabilities of selecting each policy. The relative weighting of the two is governed by γ , the precision or inverse temperature of a softmax over G . Finally, we here specifically highlight the action precision parameter α , which is implemented as the inverse temperature of a softmax transformation of the (log-)posterior over policies, after it has been marginalized only to specify probabilities for the next action, with higher values leading to more deterministic action selection. Higher α values thereby lead to more deterministic choices; i.e., the agent always selects the most likely action, as opposed to probability matching (at $\alpha = 1$), or to random choices independent of beliefs (as α gets lower). Many additional types of parametrizations can be used to extend the model (e.g., to construct hierarchically, counterfactually, and temporally deep models [60,61] or to enable parameter learning [1]), which we do not employ here. See [1] for a full account of the various possibilities.

A specific POMDP model suitable for MAB tasks was used to implement the individual active inference agents, and also assumed as generative model for the group-level agent. We here describe the details of this MAB-POMDP generative model (see Figure 2). Unless otherwise indicated, all parameter values and settings are the same for all internal agents, as well as for the group agent. The model has three environmental states s , one for each bandit arm. There are two types of observations o , 1 and 2, and three possible actions u , one for selecting each of the three bandit arms. Observation 1 is very likely for one of the arms, and unlikely for the other two (as specified by the A matrix). Agents can choose which bandit arms to play, and their choice is deterministic i.e., not subject to ‘motor errors’ (specified by a precise mapping in the B matrix). Agents have a preference for making observation 1 except when otherwise stated (specified by the C matrix), and have uniform priors over states and policies (D and E matrices). The policy posterior precision γ was 16, and the policy length is 2. With simplicity in mind, we here do not include parameter learning, but instead equip the agents with accurate beliefs about the outcome probabilities, thereby modelling behaviour at the point when learning has successfully converged. Since this means that there is nothing to learn about the environment, agents’ action selection is driven only by the pragmatic value in Equation (2).




Figure 2. Specification of the POMDP model

2.2. Cognitive Modelling for Emergent Agents

We placed the active inference agents in a Markov blanket (Figure 3). This involves: (i) a set of *sensory agents* (in this case one) who observe the environment and acts accordingly; (ii) a set of *internal agents*, who observe the actions of the sensory agent and act accordingly; and (iii) a set of *active agents* (in this case one), who observes the actions of (some of) the internal and/or sensory agents (in this case all the internal agents) and whose actions affect the environment. The sensory agent is a simple information transfer agent, i.e., it simply copies the input it gets from the environment and passes it on to the internal agents. The active agent is likewise kept as a simple aggregator, who observes the actions of all the internal agents and chooses an action with a probability proportional to the amount of times it has been made by an internal agent (thereby functioning as a probabilistic voting mechanism). The observations of the sensory agents and the actions of the active agents correspond to the observations and actions of the group-level agent, which can now be used for the model-fitting. We use standard methodology from cognitive modelling to infer the generative model implied by the behaviour of the group-level agent [62]. We use Markov-Chain Monte Carlo (MCMC)-based approximate Bayesian inference to estimate the parameter values that best explain behaviour [63,64].

The sensory and active agents could easily be implemented as active inference agents equipped with appropriate generative models. However, for simplicity we use a simple rule-based approximation to construct these agents. Owing to the simplicity of the setup, we can now use the same generative model for the individual agents and as we do for the group agent. This is because the internal agents receive and produce exactly the kind of observations and actions that a MAB task also would produce. This is not the generally the case, however; see Section 4.1 for a discussion of ways in which to extend this setup.




Figure 3. The agent group structure

2.3. Simulation Experiments

We conducted a series of four simulation experiments that all share the same basic setup and simulation environment, each introducing a variation from the basic setup. For simplicity, we here focus on a single parameter at the group level which has a clear effect on behaviour: the action precision α . α is implemented as the temperature of a softmax over the agent's action probabilities. Thus, agents with higher α will be more likely to select the action which is optimal relative to their preferences and current beliefs about the environment - lower α will lead to more stochastic, or "noisy",

behaviour behaviour. As a preliminary, we performed a parameter recovery study to confirm that the α parameter can accurately be inferred. We simulated behaviour using α values between 0 and 2. In the recovery study and in the simulation experiments, we used a Multi-Armed Bandit task (MAB) with three bandits. The probabilities of (the generally preferred) outcome 1 for the three bandits were 0.8, 0.2, and 0.2. Because α is non-negative, we used a wide half-normal distribution as prior (mean=0, SD=4, truncated to be non-negative). When comparing generative values with estimates, we used the median of the parameter posteriors as a point estimate.

In the first simulation experiment, we made all internal agents identical, and compared the α values of the internal agents to that of the group-agent. In the second experiment, we let α values vary between agents, and compared the mean of the internal agents' α values with the group α . We constructed the α values so as to control the mean of the distribution of internal agents' parameter values. This was done by drawing a set of weights from a Dirichlet distribution with sufficient statistics $\alpha_i = 1.5, i = 1, 2, \dots, n$. These weights were multiplied by the number of internal agents and the desired mean, in order to create the action precision α alpha values for the active inference agents. For the third experiment, we changed the active agent, so that it no longer probabilistically aggregated votes, but instead always chose the option with the most votes. For the fourth experiment, we let the prior preferences of the internal agents vary. We implemented this by drawing the relative preference for observation 1, for each internal agent, from a beta distribution with sufficient statistics $\alpha = 0.8, \beta = 0.8$. For each experiment, we ran the simulation separately with 4, 8, 16 and 100 internal agents. We use α values for internal agents between 0 and 1. The goal of the third and fourth experiments was to investigate relations between the group α and other characteristics of the internal agents (other than their α value).

3. Results

3.1. Parameter Recovery

The parameter recovery results shows that α can well be recovered in the range between 0 and 1 (Figure 4). The variance of inferred alphas increase at true α values above 1, with higher values resulting in inferred α values at around 2 or 4, irrespective of the true value. This indicates that α cannot be recovered well at values higher than 1. The efficiency of estimating parameters such as α depends on there being discernible consequences for behaviour when varying the parameter. When α is large, behaviour becomes largely deterministic, so that there is no effect of further increasing α . In the following, we therefore do not use values above 1.

3.2. Simulation Experiments

The results of the four simulation experiments are shown in Figure 5. We see that the group α in the first experiment is equal to the shared α of the internal agents. In the second experiment, we find that the inferred group α scales with the mean of the internal agent α values, although sub-linearly, and that having larger numbers of internal agents results in less variance in the group α values. In the third experiment, we see that the deterministic voting aggregation leads to much steeper increases in group α as internal agent α increases, with the steepness further increasing as the number of internal agents increases. This quickly reaches α values high enough to display the distinct pattern of clustering at around 4, the part of parameter space where α no longer can be inferred well. For the fourth experiment, we find that internal agents with varying preferences lead to much lower group α values, and that higher numbers of internal agents again lead to less variance in the group α values.




Figure 4. Parameter recovery results for the α (action precision) parameter. Each point is a recovery attempt, with the x-axis representing the generative α value of the agent, and the y-axis representing the inferred value. The gray line depicts for the identity function, where recovery is perfect, and is not a regression line.




Figure 5. Results of the four simulation experiments. In A), all individual agents have the same α values, which becomes identical to the group-level α . In B), agents had varying α values constructed to have a specified mean, leading to a sub-linear relation between the group-level and average individual α . In C), the active agent deterministically, as opposed to stochastically, selects the action with the most votes, leading to group-level agents with high action precision α . In D), internal agents have conflicting preferences, leading to noisy low- α group-level agents. The inferred α value of the group agent is displayed against the mean of the internal agents' α value. Each point is a single simulation, with colours distinguishing between different numbers of internal agents (red = 4 agents, blue = 8 agents, yellow = 16 agents, green = 100 agents). Gray lines depict the identity function, where the group α is identical to the mean of the internal agents' α value, and is not a regression line.

4. Discussion

We have presented simulation methods and numerical studies of collective active inference, with the generative models specified at the individual as well as ensemble levels and at corresponding spatiotemporal scales. This enabled us to investigate how the active inference process at the group level is influenced by the (inter-)active inference of individual agents. To make this possible, we employed methods from cognitive modelling, fitting an active inference model to the behaviour of the emergent group-level agent (i.e., blanket states).

As a proof-of-principle, we created a simple group of agents collaborating on a Multi-Armed Bandit (MAB) task, and focused on the action precision parameter α of the group-level agent. We found that, with the POMDP generative model apt for Multi-Armed Bandit tasks, α recovered well in a range between 0 and 1. At higher α values, actions became fully deterministic to the point where different values had no influence on behaviour, making inference converge on values around 4 (the specific level is directly determined by the standard deviation of the prior). We therefore limited further analysis to α levels between 0 and 1 for internal agents. We then proceeded to simulate group behaviour and estimated α levels of group agents. We found that the group-level agent had the same α as the constituent agents when they all had the same value (Figure 5A). This was expected in this specific case, where agents simply voted on the action that the group should take, aggregating identical agents results in a larger-scale version of that same agent.

When the α levels of individual agents vary, the group level agent enacts an α value similar to the mean of the individual agents (Figure 5B); essentially, the generative model of the larger level agent became an unweighted Bayesian model average of the constituent agents. The group-level α scaled sub-linearly with the average of the individual agents' α values - this sub-linear relation is possibly related to the way the internal agent α values are constructed. Firstly, the internal agent α values must vary around the specified mean, but are simultaneously bounded to be non-negative. In order to avoid extreme values, they are therefore sampled from a non-uniform Dirichlet distribution, which leads them to be non-symmetrically distributed, with more α values below the mean than above it (Figure 5B). If the group-level α reflects the majority, rather than the mean, of the internal agents' α values, this could lead to the observed sub-linear relation. In addition to this, differences in higher α values are less distinguishable behaviourally. This means that it is possible to increase the α values of some internal agents (and therefore the mean of the distribution) without increasing the inferred group-level α , also possibly contributing to a sub-linear scaling.

We also showcase two examples of how varying other aspects than α of the individual agents can lead to changes in the group-level α value. First, we see that a deterministic vote aggregation (where the action with the most votes is always chosen, as opposed to being chosen with a probability equal to the proportion of votes it received) led to markedly higher α values in the group-level agent (Figure 5C). The group level α now scaled super-linearly with the mean of the individual agents' α values, with the slope increasing with the amount of internal agents, now quickly leading to group α values reaching the upper bound of identifiability, where actions are fully deterministic. This was likely due to the diminishing effect of the internal agents' stochasticity on the deterministic voting aggregation as the number of agents increases. This can be interpreted as a simple consequence of the law of large numbers, where the variance of a statistical estimator decreases as the number of samples increases, such that the 'optimal' action (highest-value) is chosen more reliably as the number of agents increases. Finally, we see that when individual agents had varying preferences, the group α was strongly reduced (Figure 5D). This was due to the agents' votes now contradicting each other, resulting in a much more stochastic active agent, and therefore a lower α for the group-level agent. We also see that larger amounts of internal agents reduced the variances in group α estimates. This was due to the way we had constructed the set of preferences, making it more likely to have a balanced set of contradicting preferences for higher numbers of internal agents.

In sum, this represents an example of how to investigate relations between individual and group-level emergent generative models of active inference agents. This makes it possible formally to relate

all three levels — *within-agent*, *between-agent* and *group-as-agent* levels — simultaneously. In the last section, we discuss the limitations of our method, and ways in which it could be extended in the future.

4.1. Applications and Extensions

We deliberately used a simple group agent which was constructed with a predefined Markov blanket structure, and an environment chosen for its well-defined generative model, which can also be used for internal agents. However, the method we used does, in principle, apply to a much broader range of scenarios, namely whenever there is a identifiable Markov blanket the group level.

The method inherits a set of limitations from relying on methods from computational cognitive modelling. Most importantly, it infers the generative model of the group-agent from its observable behaviour (i.e., blanket states). This is only possible insofar as the behaviour is informative about the generative model; when several generative models (or parameter values) result in the same behaviour, it is not possible to distinguish between them. In our example, this is the case when α values are high enough enough that action becomes fully deterministic, so that variations in α cannot be distinguished. There are other types of degeneracy: a uniform preference prior would lead to random behaviour indistinguishable from that associated with a very low α value, for example. Additionally, as in most computational modelling, there is an in principle infinite space of possible generative models that the emergent agent might have. Since searching this space is not practically possible, the approach relies on the modeller specifying a model structure, or a set of model structures to compare with Bayesian model selection. This, in turn, requires an environment for which a suitable generative model is identifiable. Notably, it is not actually *a priori* necessary that the emergent group agent *has* a suitable generative model of the environment, especially in non-evolved settings, or when a group agent fails to act appropriately; however, it is often in practice assumed that agents have environment-appropriate generative models. Finally, even if a generative model is identifiable, all the challenges of doing approximate Bayesian inference (with Markov-Chain Monte Carlo methods, or with other approaches) apply here, such as the amount of computational resources required to fit models, or the risk that samplers might not be able successfully to converge when fitting the model. Methods for addressing these types of problems apply here [63], except that simulation means that large quantities of data can be generated to make model fitting more precise, and that environments can be designed specifically to allow for distinguishing between different generative models. A challenge specific to this method is that a group-level Markov blanket has to be defined clearly, specifically requiring that sensory and active states can be defined for which it is possible to define a suitable generative model.

Despite these limitations, a cognitive modelling-based approach offers a tractable way to access the generative model of emergent agents with unknown generative models of the environment, making it possible to compare them to the generative models of individual constituent agents. This approach enables describing more than just the *spatially* self-similar nature of nested Markov blankets by elucidating the *cognitively* or *computationally* fractal nature of nested agents making inference over their environment. Notably, this approach is applicable to computational models of cognition in general, although it is particularly theoretically compatible with active inference modelling. Cognitive computational modelling in general assumes a Markov Blanket structure (with the observable behaviour of an agent consisting of its observations and actions).

There are multiple ways to extend the simulations reported here, where internal agents within a static Markov Blanket vote for a group action to be taken. One might add parameter learning, so as to allow inferring over time the consequences of one's actions, such as is usually the goal in Multi-Armed Bandit tasks. One might also infer on other parameters than α for the group-agent, such as γ , the contents of the various matrices, learning rates for parameter learning, *etc.* One could also implement the sensory and active agents as active inference agents proper. This would facilitate going beyond simple aggregator and information transmission agents, for example including sensory agents that can distort information about the environment, or multiple (perhaps competing) agents in either role. It would also be possible to let agents take actions that reflect their level of certainty (which could

lead to the group-agent becoming certainty-weighted Bayesian model average of the internal agents, and would more generally make it possible further to explore the effect of individual-agent certainty dynamics on the group-agent). It is also of interest to change the communication structure of the internal agents; so that, instead of simple voting, agents interact in a more or less complex fashion, like a network structure where only a subset of internal agents pass their actions directly on to the active agents. Finally, one can also go beyond the simple POMDP-model used here, including more complex POMDP's, continuous state space active inference models, or combinations thereof. The Multi-Armed Bandit task POMDP is general enough that it can be used for the single agents as well as for the constituent agent. However, this is not in any way required. Indeed, internal agents will often find themselves in very different environments from that of the group, mandating different generative models. This could also include models with explicit Theory of Mind, allowing agents explicitly to infer whether other agents are collaborating or competing, and whether active and sensory agents might for example have reasons to be untruthful. This would also allow agents to have explicit 'altruistic' preference priors that include other agents' reaching their goals. From here, it is not difficult to replace the Multi-Armed Bandit task environment with another group-agent, for example letting the two groups compete in game-theoretic tasks as with single agents. This lends itself easily to greater-than-two-scales simulations of group-agents constituted by sets of other group-agents, in turn made up of individual agents, *etc.*

There would also be value in applying this method to other types of simulations, including those with dynamically emerging Markov Blankets. Many of the simulations reviewed in Section 1.2 do not have explicit environments for the group-agent, which makes it difficult to specify group-level generative models. It is not, however, difficult to imagine fitting a generative model to the behaviour of the active inferants [53] as if they were a single agents solving the T-maze task; the active inference fish school [54] as if it was a single fish-like agent; or the multiple Markov-Blanketed cell groups that come together to create groups of groups of cells [22], as if each group was a single cell. One might in principle even apply the same methods to deep neural network structures like variational auto-encoders, in order to infer implicit beliefs of the network.

There are also other approaches that could be of value to pursue in this context, in addition to simply comparing parameter values (or generative model types) of individual and group-level agents. One obvious direction would be to look at the relations between dynamically-changing beliefs (or prediction errors, *etc.*) at the group-level on one hand, and the equally dynamic beliefs or actions of the internal agents on the other - in the same way one might compare belief states to neural activity in model-based computational neuroscience [65]. Insofar as this were successful, it would help elucidate the relationship between internal (agent) state dynamics and emergent active inference processes, perhaps allowing for later reversing the inference to use internal state dynamics to help make inference about group-level beliefs even in the absence of observable environments. One might also combine the method with evolutionary algorithms, potentially at different scales; for example, one could test the hypothesis that selection pressure at the group-level fosters altruistic preferences in individual members, or attention to the needs of the group, while selection on the individual level fosters a stronger orientation (or stronger preference prior) toward accomplishing one's own goals. One would also be able to distinguish the average success (or negative free energy) of the individual agents *as distinct from* the negative free energy of the group agent: this would be an interesting numerical study because the variational free energy is an extensive quantity. This means that the joint free energy of a group or ensemble is simply the sum of the individual free energies (i.e., if every member of the ensemble fulfils its prior preferences, then the group should also act as if it were fulfilling its prior preference). Whether this means the free energy under a group level generative model corresponds to the sum of free energies under individual generative models is an outstanding question.

One might also extend the detection of group-level Markov Blankets; they could, for example, exist at slower timescales (as detectable for example by the renormalization group method [66]). There might also be ways to integrate multiple states as active or sensory states at the group level (perhaps in

the fashion of integrated information, as in [67] or [68]). One might for example take the central mass of the active inference fish school [54] as the group-level active state; although it is not immediately obvious that this could be considered an active blanket state in the usual sense. Finally, one could also use this method for other types of systems with unknown generative models, such as animals or artificial systems, as well as laboratory-grown organoids and “dishbrains” [69]. In general, this method could be used to further the field of active inference and applications of the Free Energy Principle, as well as multi-scale adaptive processes, by providing access to the scale-free cognitive and computational processes they imply.

Author Contributions: Conceptualization, P.T.W.; methodology, P.T.W. and C.M.; software, P.T.W. and C.L.O.; validation, C.L.O. and P.T.W.; formal analysis, P.T.W. and C.L.O.; investigation, P.T.W. and C.L.O.; resources, P.T.W. and C.M.; data curation, P.T.W. and C.L.O.; writing—original draft preparation, P.T.W.; writing—review and editing, P.T.W., C.M., C.H., K.F., C.L.O., S.W.N. and J.E.L.; visualization, S.W.N., J.E.L. and C.L.O.; supervision, C.M. and K.F.; project administration, P.T.W.; funding acquisition, N/A. All authors have read and agreed to the published version of the manuscript.

Funding: C.M. acknowledges funding from Aarhus University Research Foundation (grant no. AUFF-E-2019-7-10) and from the Carlsberg Foundation (grant no. CF21-0439).

Institutional Review Board Statement: Not applicable

Informed Consent Statement: Not applicable

Data Availability Statement: The original data presented in the study are openly available in OSF at osf.io/6z3bd.

Conflicts of Interest: The authors declare no conflicts of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

Abbreviations

The following abbreviations are used in this manuscript:

FEP	Free Energy Principle
POMDP	Partially-Observed Markov Decision Process
MAB	multi-armed bandit task

Appendix A




Figure A1. Distributions of α values across internal agents when using varying α values.

References

1. Parr, T.; Pezzulo, G.; Friston, K.J. *Active Inference: The Free Energy Principle in Mind, Brain, and Behavior*; The MIT Press, 2022. doi:10.7551/mitpress/12441.001.0001.
2. Friston, K.; FitzGerald, T.; Rigoli, F.; Schwartenbeck, P.; O'Doherty, J.; Pezzulo, G. Active inference and learning. *Neuroscience & Biobehavioral Reviews* **2016**, *68*, 862–879. doi:10.1016/j.neubiorev.2016.06.022.
3. Friston, K.; FitzGerald, T.; Rigoli, F.; Schwartenbeck, P.; Pezzulo, G. Active inference: A process theory. *Neural Computation* **2017**. doi:10.1162/NECO_a_00912.
4. Friston, K.J.; Stephan, K.E. Free-energy and the brain. *Synthese* **2007**, *159*, 417–458. doi:10.1007/s11229-007-9237-y.
5. Friston, K. The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience* **2010**, *11*, 127–138. Number: 2 Publisher: Nature Publishing Group, doi:10.1038/nrn2787.
6. Friston, K. The free-energy principle: a rough guide to the brain? *Trends in Cognitive Sciences* **2009**, *13*, 293–301. doi:10.1016/j.tics.2009.04.005.
7. Friston, K. A free energy principle for a particular physics, 2019. arXiv:1906.10184 [q-bio], doi:10.48550/arXiv.1906.10184.
8. Friston, K.; Da Costa, L.; Sajid, N.; Heins, C.; Ueltzhöffer, K.; Pavliotis, G.A.; Parr, T. The free energy principle made simpler but not too simple. *Physics Reports* **2023**, *1024*, 1–29. doi:10.1016/j.physrep.2023.07.001.
9. Ramstead, M.J.D.; Badcock, P.B.; Friston, K.J. Answering Schrödinger's question: A free-energy formulation. *Physics of life reviews* **2018**, *24*. doi:10.1016/j.plrev.2017.09.001.
10. Kirchhoff, M.; Parr, T.; Palacios, E.; Friston, K.; Kiverstein, J. The markov blankets of life: Autonomy, active inference and the free energy principle. *Journal of the Royal Society Interface* **2018**, *15*. doi:10.1098/rsif.2017.0792.
11. Hesp, C.; Ramstead, M.; Constant, A.; Badcock, P.; Kirchhoff, M.; Friston, K. A Multi-scale View of the Emergent Complexity of Life: A Free-Energy Proposal; Springer, Cham, 2019; pp. 195–227. doi:10.1007/978-3-030-00075-2_7.
12. Heins, C.; Klein, B.; Demekas, D.; Aguilera, M.; Buckley, C.L. Spin Glass Systems as Collective Active Inference. *Active Inference*; Buckley, C.L.; Cialfi, D.; Lanillos, P.; Ramstead, M.; Sajid, N.; Shimazaki, H.; Verbelen, T., Eds.; Springer Nature Switzerland: Cham, 2023; pp. 75–98. doi:10.1007/978-3-031-28719-0_6.
13. Biehl, M.; Pollock, F.A.; Kanai, R. A Technical Critique of Some Parts of the Free Energy Principle. *Entropy* **2021**, *23*, 293. Number: 3 Publisher: Multidisciplinary Digital Publishing Institute, doi:10.3390/e23030293.
14. Friston, K.; Da Costa, L.; Sakthivadivel, D.A.R.; Heins, C.; Pavliotis, G.A.; Ramstead, M.; Parr, T. Path integrals, particular kinds, and strange things. *Physics of Life Reviews* **2023**, *47*, 35–62. doi:10.1016/j.plrev.2023.08.016.
15. Friston, K.; Kiebel, S. Predictive coding under the free-energy principle. *Philosophical Transactions of the Royal Society B: Biological Sciences* **2009**, *364*, 1211–1221. doi:10.1098/rstb.2008.0300.
16. Friston, K. Computational psychiatry: from synapses to sentience. *Molecular Psychiatry* **2023**, *28*, 256–268. Publisher: Nature Publishing Group, doi:10.1038/s41380-022-01743-z.
17. Schwartenbeck, P.; Friston, K. Computational Phenotyping in Psychiatry: A Worked Example. *eNeuro* **2016**, *3*, ENEURO.0049–16.2016. doi:10.1523/ENEURO.0049-16.2016.
18. Lanillos, P.; Meo, C.; Pezzato, C.; Meera, A.A.; Baioumy, M.; Ohata, W.; Tschantz, A.; Millidge, B.; Wisse, M.; Buckley, C.L.; Tani, J. Active Inference in Robotics and Artificial Agents: Survey and Challenges, 2021. arXiv:2112.01871 [cs], doi:10.48550/arXiv.2112.01871.
19. Hohwy, J. *The Predictive Mind*; Oxford University Press, 2013. doi:10.1093/acprof:oso/9780199682737.001.0001.
20. Hohwy, J. Conscious Self-Evidencing. *Review of Philosophy and Psychology* **2022**, *13*, 809–828. doi:10.1007/s13164-021-00578-x.
21. Friston, K.J.; Daunizeau, J.; Kiebel, S.J. Reinforcement Learning or Active Inference? *PLOS ONE* **2009**, *4*, e6421. Publisher: Public Library of Science, doi:10.1371/journal.pone.0006421.
22. Palacios, E.R.; Razi, A.; Parr, T.; Kirchhoff, M.; Friston, K. On Markov blankets and hierarchical self-organisation. *Journal of Theoretical Biology* **2020**, *486*, 110089. doi:10.1016/j.jtbi.2019.110089.
23. Friston, K. Life as we know it. *Journal of The Royal Society Interface* **2013**, *10*, 20130475. Publisher: Royal Society, doi:10.1098/rsif.2013.0475.
24. Friston, K.; Frith, C. A Duet for one. *Consciousness and Cognition* **2015**, *36*, 390–405. doi:10.1016/j.concog.2014.12.003.

25. Veissière, S.P.L.; Constant, A.; Ramstead, M.J.D.; Friston, K.J.; Kirmayer, L.J. Thinking through other minds: A variational approach to cognition and culture. *Behavioral and Brain Sciences* **2020**, *43*, e90. doi:10.1017/S0140525X19001213.
26. Veissière, S.P.L.; Constant, A.; Ramstead, M.J.D.; Friston, K.J.; Kirmayer, L.J. TTOM in action: Refining the variational approach to cognition and culture. *Behavioral and Brain Sciences* **2020**, *43*, e120. doi:10.1017/S0140525X20000011.
27. Ramstead, M.J.D.; Veissière, S.P.L.; Kirmayer, L.J. Cultural Affordances: Scaffolding Local Worlds Through Shared Intentionality and Regimes of Attention. *Frontiers in Psychology* **2016**, *7*. Publisher: Frontiers, doi:10.3389/fpsyg.2016.01090.
28. Constant, A.; Ramstead, M.J.D.; Veissière, S.P.L.; Friston, K. Regimes of Expectations: An Active Inference Model of Social Conformity and Human Decision Making. *Frontiers in Psychology* **2019**, *10*. Publisher: Frontiers, doi:10.3389/fpsyg.2019.00679.
29. Isomura, T.; Parr, T.; Friston, K. Bayesian Filtering with Multiple Internal Models: Toward a Theory of Social Intelligence. *Neural Computation* **2019**, *31*, 2390–2431. doi:10.1162/neco_a_01239.
30. Pöppel, J.; Kahl, S.; Kopp, S. Resonating Minds—Emergent Collaboration Through Hierarchical Active Inference. *Cognitive Computation* **2022**, *14*, 581–601. doi:10.1007/s12559-021-09960-4.
31. Yoshida, W.; Dolan, R.J.; Friston, K.J. Game Theory of Mind. *PLOS Computational Biology* **2008**, *4*, e1000254. Publisher: Public Library of Science, doi:10.1371/journal.pcbi.1000254.
32. Devaine, M.; Hollard, G.; Daunizeau, J. Theory of Mind: Did Evolution Fool Us? *PLOS ONE* **2014**, *9*, e87619. Publisher: Public Library of Science, doi:10.1371/journal.pone.0087619.
33. Vasil, J.; Badcock, P.B.; Constant, A.; Friston, K.; Ramstead, M.J.D. A World Unto Itself: Human Communication as Active Inference. *Frontiers in Psychology* **2020**, *11*, 417. doi:10.3389/fpsyg.2020.00417.
34. Albarracin, M.; Pitliya, R.J.; St. Clere Smithe, T.; Friedman, D.A.; Friston, K.; Ramstead, M.J.D. Shared Protentions in Multi-Agent Active Inference. *Entropy* **2024**, *26*, 303. Number: 4 Publisher: Multidisciplinary Digital Publishing Institute, doi:10.3390/e26040303.
35. Friston, K.J.; Frith, C.D. Active inference, communication and hermeneutics. *Cortex; a Journal Devoted to the Study of the Nervous System and Behavior* **2015**, *68*, 129–143. doi:10.1016/j.cortex.2015.03.025.
36. Medrano, J.; Sajid, N. A Broken Duet: Multistable Dynamics in Dyadic Interactions. *Entropy* **2024**, *26*, 731. Number: 9 Publisher: Multidisciplinary Digital Publishing Institute, doi:10.3390/e26090731.
37. Friston, K.J.; Parr, T.; Yufik, Y.; Sajid, N.; Price, C.J.; Holmes, E. Generative models, linguistic communication and active inference. *Neuroscience & Biobehavioral Reviews* **2020**, *118*, 42–64. doi:10.1016/j.neubiorev.2020.07.005.
38. Friston, K.J.; Parr, T.; Heins, C.; Constant, A.; Friedman, D.; Isomura, T.; Fields, C.; Verbelen, T.; Ramstead, M.; Clippinger, J.; Frith, C.D. Federated inference and belief sharing. *Neuroscience & Biobehavioral Reviews* **2024**, *156*, 105500. doi:10.1016/j.neubiorev.2023.105500.
39. Maisto, D.; Donnarumma, F.; Pezzulo, G. Interactive inference: a multi-agent model of cooperative joint actions. *IEEE Transactions on Systems, Man, and Cybernetics: Systems* **2024**, *54*, 704–715. arXiv:2210.13113 [cs, math, q-bio], doi:10.1109/TSMC.2023.3312585.
40. Levchuk, G.; Pattipati, K.; Serfaty, D.; Fouse, A.; McCormack, R. Chapter 4 - Active Inference in Multiagent Systems: Context-Driven Collaboration and Decentralized Purpose-Driven Team Adaptation. In *Artificial Intelligence for the Internet of Everything*; Lawless, W.; Mittu, R.; Sofge, D.; Moskowitz, I.S.; Russell, S., Eds.; Academic Press, 2019; pp. 67–85. doi:10.1016/B978-0-12-817636-8.00004-1.
41. Demekas, D.; Heins, C.; Klein, B. An Analytical Model of Active Inference in the Iterated Prisoner's Dilemma. Active Inference; Buckley, C.L.; Cialfi, D.; Lanillos, P.; Ramstead, M.; Sajid, N.; Shimazaki, H.; Verbelen, T.; Wisse, M., Eds.; Springer Nature Switzerland: Cham, 2024; pp. 145–172. doi:10.1007/978-3-031-47958-8_10.
42. Isomura, T.; Friston, K. In vitro neural networks minimise variational free energy. *Scientific Reports* **2018**, *8*, 16926. doi:10.1038/s41598-018-35221-w.
43. Palacios, E.R.; Isomura, T.; Parr, T.; Friston, K. The emergence of synchrony in networks of mutually inferring neurons. *Scientific Reports* **2019**, *9*, 6412. doi:10.1038/s41598-019-42821-7.
44. Isomura, T.; Shimazaki, H.; Friston, K.J. Canonical neural networks perform active inference. *Communications Biology* **2022**, *5*, 55. doi:10.1038/s42003-021-02994-2.
45. Kastel, N.; Hesp, C. Ideas Worth Spreading: A Free Energy Proposal for Cumulative Cultural Dynamics. Machine Learning and Principles and Practice of Knowledge Discovery in Databases; Kamp, M.; Koprinska, I.; Bibal, A.; Bouadi, T.; Frénay, B.; Galárraga, L.; Oramas, J.; Adilova, L.; Krishnamurthy, Y.; Kang, B.;

- Largerone, 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.; Gavalda, R.; Cooper, L.; Ghazaleh, N.; Richiardi, J.; Roqueiro, D.; Saldana Miranda, D.; Sechidis, K.; Graça, G., Eds.; Springer International Publishing: Cham, 2021; pp. 784–798. doi:10.1007/978-3-030-93736-2_55.
46. Kastel, N.; Hesp, C.; Ridderinkhof, K.R.; Friston, K.J. Small steps for mankind: Modeling the emergence of cumulative culture from joint active inference communication. *Frontiers in Neurorobotics* **2023**, *16*. Publisher: Frontiers, doi:10.3389/fnbot.2022.944986.
 47. Albarracin, M.; Demekas, D.; Ramstead, M.J.D.; Heins, C. Epistemic Communities under Active Inference. *Entropy* **2022**, *24*, 476. Number: 4 Publisher: Multidisciplinary Digital Publishing Institute, doi:10.3390/e24040476.
 48. Friston, K.J.; Ramstead, M.J.; Kiefer, A.B.; Tschantz, A.; Buckley, C.L.; Albarracin, M.; Pitliya, R.J.; Heins, C.; Klein, B.; Millidge, B.; Sakthivadivel, D.A.; St Clere Smithe, T.; Koudahl, M.; Tremblay, S.E.; Petersen, C.; Fung, K.; Fox, J.G.; Swanson, S.; Mapes, D.; René, G. Designing ecosystems of intelligence from first principles. *Collective Intelligence* **2024**, *3*, 26339137231222481. Publisher: SAGE Publications, doi:10.1177/26339137231222481.
 49. Friston, K.; Levin, M.; Sengupta, B.; Pezzulo, G. Knowing one's place: a free-energy approach to pattern regulation. *Journal of the Royal Society, Interface* **2015**, *12*, 20141383. doi:10.1098/rsif.2014.1383.
 50. Kuchling, F.; Friston, K.; Georgiev, G.; Levin, M. Morphogenesis as Bayesian inference: A variational approach to pattern formation and control in complex biological systems. *Physics of Life Reviews* **2019**. Publisher: Elsevier B.V., doi:10.1016/j.plrev.2019.06.001.
 51. Friston, K.J.; Fagerholm, E.D.; Zarghami, T.S.; Parr, T.; Hipólito, I.; Magrou, L.; Razi, A. Parcels and particles: Markov blankets in the brain. *Network Neuroscience* **2021**, *5*, 211–251. doi:10.1162/netn_a_00175.
 52. Pio-Lopez, L.; Kuchling, F.; Tung, A.; Pezzulo, G.; Levin, M. Active inference, morphogenesis, and computational psychiatry. *Frontiers in Computational Neuroscience* **2022**, *16*. Publisher: Frontiers, doi:10.3389/fncom.2022.988977.
 53. Friedman, D.A.; Tschantz, A.; Ramstead, M.J.D.; Friston, K.; Constant, A. Active Inferants: An Active Inference Framework for Ant Colony Behavior. *Frontiers in Behavioral Neuroscience* **2021**, *15*. Publisher: Frontiers, doi:10.3389/fnbeh.2021.647732.
 54. Heins, C.; Millidge, B.; Da Costa, L.; Mann, R.P.; Friston, K.J.; Couzin, I.D. Collective behavior from surprise minimization. *Proceedings of the National Academy of Sciences* **2024**, *121*, e2320239121. Publisher: Proceedings of the National Academy of Sciences, doi:10.1073/pnas.2320239121.
 55. Kaufmann, R.; Gupta, P.; Taylor, J. An Active Inference Model of Collective Intelligence. *Entropy* **2021**, *23*, 830. Number: 7 Publisher: Multidisciplinary Digital Publishing Institute, doi:10.3390/e23070830.
 56. Bezanson, J.; Karpinski, S.; Shah, V.B.; Edelman, A. Julia: A Fast Dynamic Language for Technical Computing, 2012. arXiv:1209.5145 [cs], doi:10.48550/arXiv.1209.5145.
 57. Ge, H.; Xu, K.; Ghahramani, Z. Turing: A Language for Flexible Probabilistic Inference. Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics. PMLR, 2018, pp. 1682–1690. ISSN: 2640-3498.
 58. Slivkins, A. Introduction to Multi-Armed Bandits. *Foundations and Trends® in Machine Learning* **2019**, *12*, 1–286. Publisher: Now Publishers, Inc., doi:10.1561/22000000068.
 59. Marković, D.; Stojić, H.; Schwöbel, S.; Kiebel, S.J. An empirical evaluation of active inference in multi-armed bandits. *Neural Networks* **2021**, *144*, 229–246. doi:10.1016/j.neunet.2021.08.018.
 60. Friston, K.J.; Rosch, R.; Parr, T.; Price, C.; Bowman, H. Deep temporal models and active inference. *Neuroscience & Biobehavioral Reviews* **2017**, *77*, 388–402. doi:10.1016/j.neubiorev.2017.04.009.
 61. Friston, K.; Heins, C.; Verbelen, T.; Costa, L.D.; Salvatori, T.; Markovic, D.; Tschantz, A.; Koudahl, M.; Buckley, C.; Parr, T. From pixels to planning: scale-free active inference, 2024. arXiv:2407.20292, doi:10.48550/arXiv.2407.20292.
 62. Daunizeau, J.; Ouden, H.E.M.d.; Pessiglione, M.; Kiebel, S.J.; Stephan, K.E.; Friston, K.J. Observing the Observer (I): Meta-Bayesian Models of Learning and Decision-Making. *PLOS ONE* **2010**, *5*, e15554. Publisher: Public Library of Science, doi:10.1371/journal.pone.0015554.

63. Lee, M.D.; Wagenmakers, E.J. *Bayesian Cognitive Modeling: A Practical Course*, 1 ed.; Cambridge University Press, 2014. doi:10.1017/CBO9781139087759.
64. Blei, D.M.; Kucukelbir, A.; McAuliffe, J.D. Variational Inference: A Review for Statisticians. *Journal of the American Statistical Association* **2017**, *112*, 859–877. arXiv:1601.00670 [cs, stat], doi:10.1080/01621459.2017.1285773.
65. Palmeri, T.J.; Love, B.C.; Turner, B.M. Model-based cognitive neuroscience. *Journal of mathematical psychology* **2017**, *76*, 59–64. doi:10.1016/j.jmp.2016.10.010.
66. Ma, S.k. Introduction to the Renormalization Group. *Reviews of Modern Physics* **1973**, *45*, 589–614. Publisher: American Physical Society, doi:10.1103/RevModPhys.45.589.
67. Rosas, F.E.; Mediano, P.A.M.; Jensen, H.J.; Seth, A.K.; Barrett, A.B.; Carhart-Harris, R.L.; Bor, D. Reconciling emergences: An information-theoretic approach to identify causal emergence in multivariate data. *PLOS Computational Biology* **2020**, *16*, e1008289. Publisher: Public Library of Science, doi:10.1371/journal.pcbi.1008289.
68. Albantakis, L.; Barbosa, L.; Findlay, G.; Grasso, M.; Haun, A.M.; Marshall, W.; Mayner, W.G.P.; Zaeemzadeh, A.; Boly, M.; Juel, B.E.; Sasai, S.; Fujii, K.; David, I.; Hendren, J.; Lang, J.P.; Tononi, G. Integrated information theory (IIT) 4.0: Formulating the properties of phenomenal existence in physical terms. *PLOS Computational Biology* **2023**, *19*, e1011465. Publisher: Public Library of Science, doi:10.1371/journal.pcbi.1011465.
69. Kagan, B.J.; Kitchen, A.C.; Tran, N.T.; Habibollahi, F.; Khajehnejad, M.; Parker, B.J.; Bhat, A.; Rollo, B.; Razi, A.; Friston, K.J. In vitro neurons learn and exhibit sentience when embodied in a simulated game-world. *Neuron* **2022**, *110*, 3952–3969.e8. doi:10.1016/j.neuron.2022.09.001.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.