

Designing Ecosystems of Intelligence from First Principles

Karl J. Friston^{1,2}, Maxwell J.D. Ramstead^{*1,2}, Alex B. Kiefer^{1,3},
Alexander Tschantz¹, Christopher L. Buckley^{1,4}, Mahault Albarracin^{1,5},
Riddhi J. Pitliya^{1,6}, Conor Heins^{1,7,8,9}, Brennan Klein^{1,10}, Beren Millidge^{1,11},
Dalton A.R. Sakthivadivel^{1,12,13,14}, Toby St Clere Smithe^{1,6,15},
Magnus Koudahl^{1,16}, Safae Essafi Tremblay^{1,17}, Capm Petersen¹, Kaiser Fung¹,
Jason G. Fox¹, Steven Swanson¹, Dan Mapes¹, and Gabriel René¹

¹*VERSES Research Lab, Los Angeles, California, USA*

²*Wellcome Centre for Human Neuroimaging, University College London, London, UK*

³*Department of Philosophy, Monash University, Melbourne, Australia*

⁴*Sussex AI Group, Department of Informatics, University of Sussex, Brighton, UK*

⁵*Department of Computer Science, Université du Québec à Montréal, Montréal, Québec, Canada*

⁶*Department of Experimental Psychology, University of Oxford, Oxford, UK*

⁷*Department of Collective Behaviour, Max Planck Institute of Animal Behavior,
Konstanz, Germany*

⁸*Department of Biology, University of Konstanz, Konstanz, Germany*

⁹*Centre for the Advanced Study of Collective Behaviour, University of Konstanz,
Konstanz, Germany*

¹⁰*Network Science Institute, Northeastern University, Boston, Massachusetts, USA*

¹¹*Brain Network Dynamics Unit, University of Oxford, Oxford, UK*

¹²*Department of Mathematics, Stony Brook University, Stony Brook, New York, USA*

¹³*Department of Physics and Astronomy, Stony Brook University, Stony Brook, New York, USA*

¹⁴*Department of Biomedical Engineering, Stony Brook University, Stony Brook, New York, USA*

¹⁵*Topos Institute, Berkeley, California, USA*

¹⁶*Department of Electrical Engineering, Eindhoven University of Technology,
Eindhoven, The Netherlands*

¹⁷*Department of Philosophy, Université du Québec à Montréal, Montréal, Québec, Canada*

December 5, 2022

Contents

1	Introduction	3
2	A first-principles approach to multi-scale artificial intelligence	4

*maxwell.ramstead@verses.io

3	Active inference	6
3.1	“Model evidence is all you need”	6
3.2	AI designed for belief updating	8
3.3	Comparison to current state-of-the-art approaches	9
3.3.1	Managing complexity	9
3.3.2	Reinforcement learning and active inference	11
3.3.3	Multi-scale considerations	11
3.4	Shared narratives	12
4	From Babel to binary	13
4.1	Active inference and communication	14
4.2	Belief propagation, graphs, and networks	15
4.3	Intelligence at scale	16
5	Ethical and moral considerations	17
6	Conclusion: Our proposal for stages of development for active inference as an artificial intelligence technology	19
6.1	Stages of development for active inference	19
6.2	Implementation	20
A	Appendix: Applications of active inference	1

Abstract

This white paper lays out a vision of research and development in the field of artificial intelligence for the next decade (and beyond). Its denouement is a cyber-physical ecosystem of natural and synthetic sense-making, in which humans are integral participants—what we call “shared intelligence”. This vision is premised on active inference, a formulation of adaptive behavior that can be read as a physics of intelligence, and which inherits from the physics of self-organization. In this context, we understand intelligence as the capacity to accumulate evidence for a generative model of one’s sensed world—also known as self-evidencing. Formally, this corresponds to maximizing (Bayesian) model evidence, via belief updating over several scales: i.e., inference, learning, and model selection. Operationally, this self-evidencing can be realized via (variational) message passing or belief propagation on a factor graph. Crucially, active inference foregrounds an existential imperative of intelligent systems; namely, curiosity or the resolution of uncertainty. This same imperative underwrites belief sharing in ensembles of agents, in which certain aspects (i.e., factors) of each agent’s generative world model provide a common ground or frame of reference. Active inference plays a foundational role in this ecology of belief sharing—leading to a formal account of collective intelligence that rests on shared narratives and goals. We also consider the kinds of communication protocols that must be developed to enable such an ecosystem of intelligences and motivate the development of a shared hyper-spatial modeling language and transaction protocol, as a first—and key—step towards such an ecology.

1 Introduction

This white paper presents *active inference* as an approach to research and development in the field of artificial intelligence (AI), with the aim of developing *ecosystems* of natural and artificial intelligences. The path forward in AI is often presented as progressing from systems that are able to solve problems within one narrow domain—so-called “artificial narrow intelligence” (ANI)—to systems that are able to solve problems in a domain general-manner, at or beyond human levels: what are known as “artificial general intelligence” (AGI) and “artificial super-intelligence” (ASI), respectively [1]. We believe that approaching ASI (or, for reasons outlined below, even AGI) likely requires an understanding of networked or collective intelligence. Given the growing ubiquity of things like autonomous vehicles, robots, and arrays of edge computing devices and sensors (collectively, the internet of things), the zenith of the AI age may end up being a distributed network of intelligent systems, which interact frictionlessly in real time, and compose into emergent forms of intelligence at superordinate scales. The nodes of such a distributed, interconnected ecosystem may then be human users as well as human-designed artifacts that embody or implement forms of intelligence.

In order to enable such ecosystems, we must learn from nature. While acknowledging neuroscience as a key inspiration for AI research, we argue that we must move beyond brains, and embrace the active and nested characteristics of natural intelligence, as it occurs in living organisms and as it might be implemented in physical systems more generally. In our view, this entails asking not only “How does intelligence present to us, as researchers?” but also, crucially, the complementary question “What is it that intelligence *must be, given that* intelligent systems exist in a universe like ours?”

To address this challenge, we foreground *active inference*, which combines the virtues of a “first-principles,” physics-based approach to AI with Bayesian formulations, thus reframing and, in some key respects, extending the methods found in Bayesian approaches to machine learning, which provide the foundations of state-of-the-art AI systems. Active inference entails an explicit mechanics of the beliefs of agents and groups of agents—known as Bayesian mechanics [2, 3]—which is uniquely suited to the engineering of ecosystems of intelligence, as it allows us to write down the dynamics of sparsely coupled systems that self-organize over several scales or “levels” [4–7]. We offer a formal definition of intelligence for AI research, as the capacity of systems to generate evidence for their own existence. This encompasses cognition (i.e., problem-solving via action and perception) and curiosity, as well as the capacity for creativity, which underwrites the current interest in generative AI [8]. We argue that the design of intelligent systems must begin from the *physicality of information and its processing* at every scale or level of self-organization. The result is AI that “scales up” the way nature does: by aggregating individual intelligences and their locally contextualized knowledge bases, within and across ecosystems, into “nested intelligences”—rather than by merely adding more data, parameters, or layers to a machine learning architecture.

We consider the question of how to engineer ecosystems of AI using active inference, with a focus on the problem of communication between intelligent agents, such that shared forms of intelligence emerge, in a nested fashion, from these interactions. We highlight the importance of shared narratives and goals in the emergence of collective behavior, and how active

inference helps account for this in terms of sharing (aspects of) the same generative model. We close our discussion with a sketch of stages of development for AI using the principles of active inference. Our hypothesis is that taking the *multi-scale* and *multi-level* aspects of intelligence seriously has the potential to be transformative with respect to the assumptions and goals of research, development, and design in the field of AI, with potentially broad implications for industry and society: that technologies based on the principles described in this paper may be apt to foster the design of an emergent ecosystem of intelligences spanning spatial and cognitive domains (a hyper-spatial web).

2 A first-principles approach to multi-scale artificial intelligence

The field of artificial intelligence has from the outset used natural systems, whose stunning designs have been refined over evolutionary timescales, as templates for its models. Neuroscience has been the most significant source of inspiration, from the McCulloch-Pitts neuron [9] to the parallel distributed architectures of connectionism and deep learning [10, 11], to the contemporary call for “Neuro-AI” as a paradigm for research in AI, in particular machine learning [12]. Indeed, the definitive aspect of deep learning inherits from the hierarchical depth of cortical architectures in the brain [13]. More recently, machine learning has come, in turn, to influence neuroscience [14–16].

Academic research as well as popular media often depict both AGI and ASI as singular and monolithic AI systems, akin to super-intelligent, human individuals. However, intelligence is ubiquitous in natural systems—and generally looks very different from this. Physically complex, expressive systems, such as human beings, are uniquely capable of feats like explicit symbolic communication or mathematical reasoning. But these paradigmatic manifestations of intelligence exist along with, and emerge from, many simpler forms of intelligence found throughout the animal kingdom, as well as less overt forms of intelligence that pervade nature.

Examples of “basal cognition” abound—and often involve distributed, collective forms of intelligence. Colonies of slime molds, for example, can—as a group—navigate two-dimensional spatial landscapes, and even solve mathematical problems that are analytically intractable [17]. Certain forms of cognition and learning are (at least arguably) observable in plants [18], and we know that plants grow in a modular fashion, as a structured community of tissues that self-organize into a specific configuration [19]. Communication between organisms is often mediated by network structures, which themselves consist of other organisms; for instance, it is known that mycorrhizal networks are able to facilitate communication, learning, and memory in trees [20]. Mobile groups of schooling fish can, as a collective, sense light gradients over a wide spacetime window, even as the individuals that comprise the group can only detect local light intensity [21]. Perhaps most germanely, in morphogenesis (i.e., pattern formation in multicellular organisms), the communication of collectives of cells implements a search for viable solutions in a vast problem space of body configurations [22–

24]. This is not merely a metaphorical extension or use of the word “intelligence,” as it is no different, at its core, from our experience of navigating three-dimensional space [6].

Thus, at each physical spatiotemporal scale of interest, one can identify systems that are competent in their domain of specialization, lending intelligence in physical systems a fundamentally multi-scale character [25, 26]. Observation of nature suggests, moreover, that simpler and more complex forms of intelligence are almost always related compositionally: appreciably intelligent things tend to be composed of systems that are also intelligent to some degree. Most obviously, the intelligence of individual human beings, to the extent that it depends on the brain, implements the collective intelligence of neurons—harnessed by many intervening levels of organization or modularity, and subserved by organelles at the cellular level. This communal or collective aspect of intelligence is reflected in the etymology of “intelligence”—from *inter-* (which means between) and *legere* (which means to choose or to read)—literally inter-legibility, or the ability to understand one another.

Since intelligence at each scale supervenes on, or emerges from, simpler (though still intelligent) parts, the multi-scale view of natural intelligence implies not a mysterious infinite regress, but a recursive, nested structure in which the same functional motif (the action-perception loop) recurs in increasingly ramified forms in more complex agents [27]. The emergence of a higher-level intelligence—from the interaction of intelligent components—depends on *network structure* (e.g., the organization of the nervous system, or communication among members in a group or population) and *sparse coupling* (i.e., the fact that things are defined by what they are *not* connected to [28]), which together often lead to functional specialization among the constituents [29].

But how do we engineer systems like these? Can we leverage the fundamental organizing principles that underwrite the operation of intelligence in nature, and separate them from the contingent details of particular biological systems? We argue that this can be achieved by moving away from asking, “How does intelligence present to us, as researchers?”, which focuses on empirical descriptive adequacy and “reasoning by analogy” (e.g., the Turing test or imitation game [30]). Instead, we must ask the more fundamental question: “What is it that intelligence *must be, given that* intelligent systems exist in nature?” Such an approach aims to deduce fundamental properties of intelligence from foundational considerations about the nature of persisting physical systems (i.e., “first principles”).

This approach has its origins in the cybernetics movement of the 1940s, which set out to describe the general properties of regulatory and purposive systems—that is, properties not tied to any given specific architecture—and from which we draw now commonplace principles of system design, such as feedback and homeostasis [31, 32]. Perhaps the most well-known example of this is the good regulator theorem [33], later developed as the internal model principle in control theory [34], according to which systems that exist physically must contain structures that are homomorphic to whatever environmental factors they are capable of controlling.

Contemporary developments in the statistical mechanics of far-from-equilibrium systems (and in particular, multi-scale, living systems) allow us to formalize these insights—of early cybernetics—as a physics of self-organization, which enables the study of intelligence itself as

*a basic, ubiquitous, physical phenomenon.*¹ This has been called a *physics of sentient systems*; where “sentient” means “responsive to sensory impressions” [38, 39]. More specifically, we argue that one can articulate the principles and corollaries of the core observation that intelligent systems (i.e., agents) exist in terms of a “Bayesian mechanics” that can be used to describe or simulate them [2, 3].

We note that physical implementation is the ultimate constraint on all forms of engineered intelligence. While this claim might sound trivial, it has been a core locus of recent progress in our understanding of the physics of information itself. According to Landauer’s principle [40–43], there is an energy cost to irreversibly read-write any information in a physical medium. Thus, the physicality of information and its processing at every scale of self-organization should be accounted for in the design of intelligent systems. Apart from being principled, forcing models to respect constraints or conservation laws—of the kind furnished by physical implementation—often improves their performance or even enables unique capabilities.² Our core thesis is that all of this is naturally accommodated by an approach to AI grounded in the physics of intelligence.

3 Active inference

3.1 “Model evidence is all you need”

We approach the challenges just outlined from the perspective of active inference, a first-principles or physics-based approach to intelligence that aims to describe, study, and design intelligent agents *from their own perspective* [54]. Active inference shares the same foundations as quantum, classical, and statistical mechanics, and derives a scale-free theory of intelligence by adding an account of the individuation of particular things within their environments [38].

We begin with the observation that individual physical objects can be defined by the typical *absence of influence* of some parts of the universe on others (for example, air temperature directly impacts my skin, but not my internal organs). In *sparse* causal networks, some nodes act as informational bottlenecks that serve both as mediating channels and as (probabilistic) boundaries, on which the variability of states on either side is conditioned. The persistence of such stable boundaries in a changing world (i.e., away from thermodynamic

¹Researchers in AI have often borrowed tools from physics, such as Hamiltonian mechanics, to finesse the inference problems that they face, leading to tools like the Hamiltonian Monte Carlo algorithm, which massively speeds up certain kinds of inferential problem-solving [35]. Conversely, AI has been used in physics, chemistry, and biochemistry to great effect, allowing us to simulate the containment of plasma in Tomahawk nuclear fusion reactors [36], or predict the ways in which proteins will fold, as the famous AlphaFold system enables [37]. What we have in mind, however, is not to borrow techniques or formalisms from physics to solve the problem of intelligent systems design, or to use AI to help finesse problems from physics; but rather, in a complementary fashion, to treat the study of intelligence itself as a chapter of physics.

²Simulated neural networks, for example, often overfit and fail to generalize if they are not forced to learn compressed representations of their inputs [44–47]. Relatedly, ubiquitous forms of regularization can be motivated from physical considerations about the finite bandwidth of neurons [48], and schemes such as predictive coding and sparse coding by considerations about efficient signal transmission [47, 49–53].

equilibrium) is possible only to the extent that the boundary conditions can be predicted and controlled, leveraging an implicit statistical model—a *generative model* of how they are caused by external changes.³

To exist as an individuated thing is thus to gather observational evidence for such a model (“self-evidencing” [55]). This “model evidence” can be scored by a scalar value that conveys the degree to which some observations conform to (i.e., are predictable from) the model. To account for perception, one can update variables in order to maximize model evidence (e.g., update beliefs to match the data). To account for learning, one can update parameters in order to maximize model evidence (e.g., update models to match the data). To account for action, one can select actions in order to maximize (expected) model evidence (assuming that the model encodes preferences in terms of prior beliefs) [39, 56]. From this perspective, model evidence is the only thing that needs to be optimized.

Importantly, model evidence can be approximated in a form that has interesting decompositions, into quantities that map onto distinct facets of intelligence. For instance, an upper bound on the model evidence—called *variational free energy* [57]—can always be written as accuracy minus complexity. When a system minimizes this bound, in so doing, it automatically maximizes the predictive accuracy of its model while minimizing its complexity (implementing a version of Occam’s razor). This means that self-evidencing mandates an accurate account of sensory exchanges with the world that is minimally complex, which serves to limit overfitting and poor generalization [58].

Active inference builds on these insights. If inference entails maximizing accuracy while minimizing complexity, it follows that self-evidencing should minimize the inaccuracy and complexity that is expected following upon a course of action. It transpires that expected complexity is exactly the same quantity minimized in optimal control theory [59, 60]; namely, *risk*, while expected inaccuracy is just the *ambiguity* inherent in the way we sample data (e.g., resolved by switching the lights on in a dark room). Perhaps more interestingly, the ensuing *expected free energy* can be rearranged into expected information gain and expected value, where value is just the (log) preference for an outcome. This result captures exactly the dual aspects of Bayes optimality; namely, optimal Bayesian experimental design [61–63] and decision theory [64]. In essence, it favors choices that ensure the greatest resolution of uncertainty, under the constraint that preferred outcomes are realized. In other words, it mandates information and preference-seeking behavior, where one contextualizes the other. The ensuing curiosity or novelty-seeking thus emerges as an existential imperative [61, 62, 65–68]—to the extent that one could say that to be intelligent is (in part) to be curious, and to balance curiosity against preferences or reward in an optimal fashion.

Crucially, the approach to existence as modeling just outlined can be applied recursively, in a nested fashion, to systems as well as their components, providing the foundations for mathematical theories of collective intelligence at any scale, from rocks to rockstars.⁴ Indeed,

³In the context of scientific modeling, a statistical model is a mathematical object that encodes the way that things change, relative to the way that other things change. Formally, the structure that encodes such contingencies is called a joint probability distribution. This is the generative model.

⁴Even rocks, while not agents *per se*, track the state of their environment: for instance the interior of a rock “knows” that the environment must be well below the melting point of rock (albeit not under that

if existing in a characteristic way just is soliciting or generating evidence for our existence, then everything that exists can be described as engaging in inference, underwritten by a generative model. Dynamics quite generally can then be cast as a kind of belief updating in light of new information: i.e., changing your mind to accommodate new observations, under the constraint of minimal complexity.

3.2 AI designed for belief updating

The principles of natural design that we’ve reviewed suggest that next-generation AI systems must be equipped with *explicit beliefs* about the state of the world; i.e., they should be designed to implement, or embody, a specific perspective—a perspective under a generative model entailed by their structure (e.g., phenotypic hardware) and dynamics. (Later, we will suggest that efforts should also be directed towards research and development of communication languages and protocols supporting ecosystems of AI.)

A formal theory of intelligence requires a calculus or mechanics for movement in this space of beliefs, which active inference furnishes in the form of Bayesian mechanics [2]. Mathematically, belief updating can be expressed as movement in an abstract space—known as a statistical manifold—on which every point corresponds to a probability distribution [70–75]. See Figure 1. This places constraints on the nature of *message passing* in any physical or biophysical realization of an AI system [57, 76–79]: messages must be the sufficient statistics or parameters of probability distributions (i.e., Bayesian beliefs). By construction, these include measures of uncertainty. Any variable drawn from a distribution (e.g., the beliefs held by agents about themselves, their environment, and their possible courses of action) are associated with a measure of confidence, known as precision or inverse variance. Thus, intelligent artifacts built according to these principles will appear to quantify their uncertainty and act to resolve that uncertainty (as in the deployment of attention in predictive coding schemes [80–84]). Uncertainty quantification is particularly important when assessing the evidence for various models of data, via a process known as structure learning or Bayesian model comparison [85–89].

There are several types of uncertainty at play when learning from data. First, there may be irreducible noise in the measurement process itself. Examples of such noise include pixel blur in images. Second, the values of the hidden variables being estimated from data may be ambiguous (e.g., “Is the image I’m viewing of a duck or a rabbit?” or “It looks like rain: should I bring an umbrella?”). Third, there may be noise in the model of the function being learned (e.g., “What do rabbits look like? How do hidden variables map to data?”). Overcoming and accounting for these different types of uncertainty is essential for learning.

Non-probabilistic approaches to AI encounter these forms of uncertainty but do not represent them explicitly in the structure or parameters of their functions. These methods thus hope to learn successfully without quantifying uncertainty, which is variably feasible depending on the specific data and output being learned. AI systems that are not purpose-

English description). As systems become more elaborate, they can represent more about the things to which they couple [69].

built to select actions in order to reduce uncertainty in an optimal manner will struggle to assign confidence to their predictions. Further, as users of these kinds of AI systems, *we* have no way of knowing how confident they are in their assignments of probability—they are “black boxes”.

Taken together, the probabilistic approach provides a normative theory for learning—starting from the first principles of how AI should deal with data and uncertainty. The downside to probabilistic modeling is that it induces severe computational challenges. Specifically, such models must marginalize all the variables in the model in order to arrive at exact “beliefs” about a given variable. Thus, the main computational task in probabilistic inference is marginalization, whereas in traditional AI it is the optimization of parameters. As such, a focus on optimization *per se* in contemporary AI research and development may be misplaced to some extent. Current state-of-the-art AI systems are essentially general-purpose optimization machines, built to handle a specific task domain. But optimization in and of itself is not the same as intelligence. Rather, in an intelligent artifact, optimization should be a *method in the service of optimizing our beliefs* about what is causing our data. Fortunately, there are mathematical tricks, such as variational inference, which convert the (intractable) problem of marginalization into a (tractable) optimization problem, allowing probabilistic approaches to utilize the wealth of techniques available for optimization while retaining the benefits of uncertainty quantification.

3.3 Comparison to current state-of-the-art approaches

Active inference is a very general formulation of intelligence, understood as a self-organizing process of inference. Yet the generality of the formulation is integrative, rather than adversarial or exclusive: it formally relates or connects state-of-the-art approaches (e.g., it has been shown that all canonical neural networks minimize their free energy; [90]), showcasing their strengths, enabling cross-pollination, and motivating refinements.

3.3.1 Managing complexity

In the context of machine learning, the complexity term derivable from model evidence (a.k.a., information gain) is especially interesting [87], since it means that active inference puts predictive accuracy and complexity on an equal footing. In brief, self-evidencing bakes complexity into the optimization of beliefs about the world in a way that automatically finesses many problems with machine learning schemes that focus solely on accuracy [91]. To take a salient example from recent discussions, many of the considerations that seem to motivate non-generative approaches—to learning world models [92]—stem from considering only the likelihood in generative models, rather than model evidence or marginal likelihood—whereas the inclusion of complexity encourages a model to find parsimonious explanations of observations, abstracting from useless detail. In other words, accuracy comes at a (complexity) cost, which must be discounted.

Complexity minimization also speaks to the importance of dimensionality reduction and coarse-graining as clear ways to learn the structure of generative models [85]. This is moti-

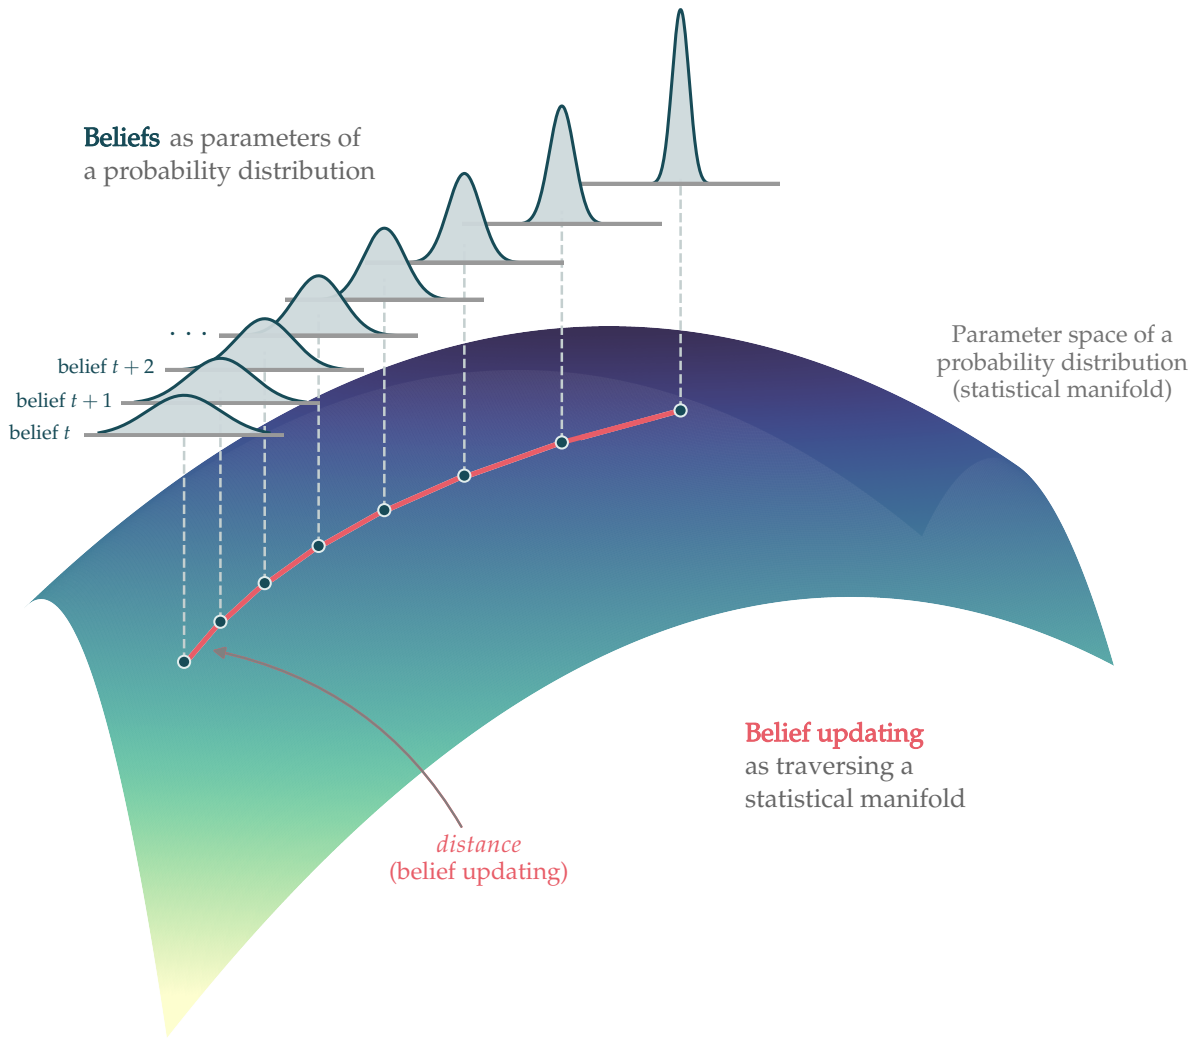


Figure 1: Belief updating on a statistical manifold.

vated by the intuition that, while sensor data itself is extremely high-dimensional, noisy, unpredictable, ambiguous, and redundant, there is a description of the data in terms of its generating causes (e.g., the world of objects with defined properties) that is lower-dimensional, more predictable, and less ambiguous. Such a description furnishes a compressed, and therefore more efficient, account of the data at hand. Thus, while scaling up data spaces, one may have to scale down the number of latent states generating those data, to the extent that doing so does not sacrifice accuracy.

Relatedly, active inference provides a promising framework for learning representations in which distinct generative factors are *disentangled* [93], via the sensorimotor contingencies associated with controllable latent factors [94–96]. Low-dimensional disentangled representations, in addition to being useful for an AI system itself in achieving its own ends, are

more explainable and human-interpretable than generic latent representations.

Minimizing complexity (i.e., compression) points to a direction of travel to the Holy Grail of a generic and robust AI; a move from “big data” to “smart data” or frugal data sampling. This has important implications for hardware and energy efficiency. Because the complexity cost has an accompanying thermodynamic cost—via Landauer’s principle and the Jarzynski equality [42]—there is a lower bound on the thermodynamic cost of any belief updating that can, even in principle, be realized with the right updating scheme. Using active inference, belief updating can be implemented with biomimetic efficiency, without the need for traditional, GPU-based high-performance computing and accompanying costs.

3.3.2 Reinforcement learning and active inference

State-of-the-art AI designed for action selection typically implements *reinforcement learning*, a set of methods for maximizing the expected sum of rewards under a sequence of actions. From a Bayesian perspective, however, curiosity and exploration are as fundamental to intelligence as maximizing reward. Specifically, the epistemic, exploratory, curiosity-driven aspect of intelligence motivates actions expected to reduce uncertainty in the variables and parameters that define your model; which, in the active inference formulation, corresponds to inference and learning, respectively [61, 62, 65–68].

In line with the above discussion of self-evidencing, rather than select actions that maximize some arbitrary state-dependent reward (or equivalently, minimize an arbitrary cost function), an intelligent system ought to generate observations or sensory data consistent with its characteristic (i.e., preferred) exchanges with the sensed world, and thus with its own continued existence. That is, an intelligent agent ought to maximize the evidence for its generative model. Active inference thereby generalizes the notion of reward, and labels every encountered outcome (and implicitly every latent state) in terms of how likely it is that “this would happen to me”. This is scored in terms of prior preferences over outcomes, which are part of the generative model. Preferences over some kinds of outcomes are precise (e.g., not being eaten or embarrassed), others less so (“I prefer coffee in the morning, but tea is nice”). To summarize, preferences provide constraints that define the “kind of thing I am,” with more precise preferences playing a similar role, for example, to the “intrinsic costs” in the actor-critic system proposed in [92].

In this view, Bayesian reinforcement learning is a *special case* of active inference, in which the preference for all outcomes is very imprecise—apart from one privileged outcome called reward, for which there is a very precise preference. The perspective from active inference moves our notion of intelligence away from a monothematic reward optimization problem, towards a multiple-constraint-satisfaction problem, where the implicit ‘satisficing’ [97] just is self-evidencing.

3.3.3 Multi-scale considerations

Another key difference concerns the multi-scale architecture of active inference. First, active inference commits to a separation of temporal scales, which allows it to finesse key issues in

AI research. On the present account, learning is just slow inference, and model selection is just slow learning. All three processes operate in the same basic way, over nested timescales, to maximize model evidence.

Second, active inference predicts, and provides a formalism for describing, the multi-scale character of intelligence in nature; see also [98]. Although this has generally not been a focus of research in machine learning, work in the field consonant with this perspective includes the complex internal structure of LSTM cells [91], the repetition of the split-transform-merge strategy across scales in the ResNeXt architecture [99], capsule networks [100], in which individually complex nodes engage in a form of self-organization, the Thousand Brains theory of the cooperation of cortical columns to produce global representations [101], or the perspective on restricted Boltzmann machines as products of experts [102].

3.4 Shared narratives

We have noted that intelligence as self-evidencing is inherently perspectival, as it involves actively making sense of and engaging with the world from a specific point of view (i.e., given a set of beliefs). Importantly, if the origins of intelligence indeed lie in the partitioning of the universe into subsystems by probabilistic boundaries, then intelligence never arises singly but always exists *on either side* of such a boundary [103, 104]. The world that one models is almost invariably composed of other intelligent agents that model one in turn.

This brings us back to the insight that intelligence must, at some level, be distributed over every agent and over every scale at which agents exist. Active inference is naturally a theory of collective intelligence. There are many foundational issues that arise from this take on intelligence; ranging from communication to cultural niche construction: from theory of mind to selfhood [103–107]. On the active inference account, shared goals emerge from shared narratives, which are provided by shared generative models [108]. Furthermore—on the current analysis—certain things should then be *curious about each other*.

The importance of perspective-taking and implicit shared narratives (i.e., generative models or frames of reference) is highlighted by the recent excitement about *generative AI* [8], in which generative neural networks demonstrate the ability to reproduce the kinds of pictures, prose, or music that we expose them to. Key to the usage of these systems is a *dyadic interaction* between artificial and natural intelligence, from the training of deep neural networks to the exchange of prompts and generated images with the resulting AI systems, and the subsequent selection and sharing of the most apt “reproductions” among generated outputs.⁵ In our view, a truly intelligent generative AI would then become curious about us—and want to know what we are likely to select. In short, when AI takes the initiative to ask *us* questions, we will have moved closer to genuine intelligence, as seen through the lens of self-evidencing.

⁵The importance of fluid exchange between artificial and human intelligence in this paradigm is evinced by the rapidly growing interest in *prompt engineering*, i.e., an increasingly self-aware and theory-driven approach to the role that prompts play in co-creating the outputs of these types of systems [109], which has recently been extended to the optimization of text prompts by distinct AI agents [110].

4 From Babel to binary

Human intelligence and language have co-evolved, such that they both scaffold, and are scaffolded by, one another [111, 112]. The core functional role of language is to enable communication and shared understanding: language has been optimized for sharing with other intelligent creatures (as a language that can be easily passed on reaches further generations). Language has thus facilitated the emergence of more complex interactions and shared customs between agents, which has in turn allowed for the emergence of intensive human collaboration at multiple communal scales [113]. Relatedly, language provides a reference for how to “carve nature at its joints” (e.g., into objects, properties, and events), facilitating learning about the world and the way it works. Finally, it has allowed humans to build an external store of knowledge far beyond the epistemic capacity of any human individual. Human beings both benefit from—and contribute to—this store of knowledge, which, like language itself, has co-evolved with our intelligence.

Across cultures, the earliest recorded narratives of our species have emphasized the astounding integrative power of shared communication systems along with their flipside: the discord and disarray wrought by miscommunication and a lack of mutual understanding. This is illustrated potently in the biblical story of the Tower of Babel, which tells of a mighty civilization that attempted to build a glorious city with a tower that rose to the heavens. These lofty aspirations fell to ruin after a divine disruption that eliminated their common language, shattering it into a thousand uninterpretable dialects. In their confusion and miscomprehension, they were unable to complete the Tower and were thus scattered across the Earth, forced to survive in the clustered tribes that shared their regional vernacular.

Today, humans cope with a “post-Babel” world via a combination of increasing multilingualism, rallying (for better or worse) behind hegemonic languages like English, and, recently, increasingly effective machine translation [114]. Digital computers *do* share a common or universal machine language (i.e., binary representation). If situations can be represented adequately in an appropriate machine syntax, they can be subjected to the operations of mathematical logic, formalized and thereby processed in an unambiguous way. At a higher level, it may be said that “vectorese” is the universal language of AI, in that vectors (i.e., ordered lists of numbers representing a point in an abstract space) constitute the input, output, and medium of data storage and retrieval for most AI algorithms.

Vectors are analogous to the medium of action potentials in the brain—they are capable of representing anything we can think of, but nearly all the interesting (and representationally load-bearing) structure lies in the (learned) transformations and accompanying transition dynamics of the underlying dynamical system. Often, an output vector space can be considered as an embedding or transformation of the input space, and mappings among vector spaces are much like translations among languages. However, vectors themselves may only provide a base structure or medium (analogous to sound or light) for higher-level languages.

It has been clear from the early days of neural language modeling that vector space representations can in principle be learned that capture both the semantic and syntactic regularities implicit in the co-occurrence statistics of natural language corpora [115, 116]. Despite this, we lack anything like *a common high-level language* that AIs can use to commu-

nicate with one another and with humans—other than, arguably, human natural languages themselves, which can be used to interface with AIs via modern language models. The fact that reinforcement learning agents trained to produce prompts for such models often produce unintelligible nonsense strings, however [110, 117], shows that even where large language models use English, they do not use or understand it in the way humans do; this raises the question whether natural languages can really play the role of a shared human-machine language without modification.

Moreover, while the necessity of serializing thought into discrete token sequences for the purposes of communication helps enforce the kind of sparsity structure that we have argued is essential to intelligence and complexity itself, a more direct form of information transfer is also conceivable in which the richness of a latent vector representation (or “thought”) is directly externalized as a data structure. While current state-of-the-art AI can learn the language of vector space embeddings, the science of inter-AI communication and shared latent spaces is in its infancy. For the most part, each AI must learn to carve up the world from scratch, and is unable to share its knowledge fluidly or update it in collaboration with other AIs.⁶

We argue that the future evolution of AI would benefit greatly from a focus on optimization for *shareability* (i.e., gathering evidence for a model of an intrinsically *social* creature.) This might take the form of a shared external store of knowledge about *how to communicate with relevant others*, or a structured communication protocol that can act as the *lingua franca* of AI. A general framework that ties together different embedding spaces and inter-AI messaging over a shared network architecture would, among other things, enable AI agents to learn to offload certain tasks or predictions to other, more specialized AI agents.

4.1 Active inference and communication

An underlying theme thus far is that intelligence at any scale requires a shared generative model and implicit common ground. There are many ways to articulate this theme; from ensemble learning to mixtures of experts [102], from distributed cognition to Bayesian model averaging [118].

Imagine that someone has locked you in a large dark room. As a self-evidencing and curious creature, you would be compelled to feel your way around to resolve uncertainty about your situation. Successive palpations lead you to infer that there is a large animal in the room—by virtue of feeling what seem to be a tail, a succession of legs, and eventually a trunk. Your actions generate accumulated evidence for the hypothesis “I am in a room with an elephant.” Now, imagine an alternative scenario in which you and five friends are deployed around the same room, and can report what you feel to each other. In this scenario, you quickly reach the consensus “We are in a room with an elephant.” The mechanics of belief updating are similar in both scenarios. In the first, you accumulate evidence and

⁶An important exception is the proliferation of fine-tuned copies of large monolithic pre-trained models such as BERT. This is not obviously relevant to our interest in (possibly real-time) communication and mutual updating among persistent, physically situated AI systems, though it may constitute a form of evolution of populations of AI systems with partially divergent learning histories.

successively update your posterior belief about latent states. In the second, the collective assimilation of evidence is parallelized across multiple individuals.

Is the latter equivalent to having one brain with twelve hands? Not quite. The second kind of belief updating rests upon a shared generative model or hypothesis space that enables you to assimilate the beliefs of another. For example, you share a common notion of a “trunk,” a “leg,” and a “tail”—and crucially, you have access to a shared language for communicating such concepts. Sharing a generative model allows each agent to infer the causes of its sensations and, crucially, disentangle the causes that are unique to the way the world is sampled—e.g., “where I am looking”—and causes that constitute the shared environment (e.g., “what I am looking at”) [4, 104, 119]. Crucially, any dyad or ensemble of self-evidencing agents will come to share a generative model (or at least some factors of a generative model) via their interactions [28] (see [120, 121] for numerical experiments in active inference that illustrate this phenomenon, and Table A.1 for related applications.)

What results is a shared intelligence (i.e., a kind of collective super-intelligence) that emerges from an ensemble of agents. Heuristically, maximizing model evidence means making the world as predictable as possible [122, 123]. This is assured if we are both singing from the same hymn sheet, so to speak—so that I can predict you and you can predict me. Mathematically, this is evinced as a generalized synchrony between the dynamics on our respective statistical manifolds [120, 124]. This generalized synchrony (or synchronicity) is special because it unfolds in a (shared) belief space, meaning it can be read as mutual understanding: i.e., coming to align our beliefs, via a shared language and a shared generative model. This sharedness is arguably the basis of culture and underpins the existence of our civilization. Our challenge, which we take to be a necessary step toward ASI or even AGI, is to expand the sphere of culture to include artificial agents.

4.2 Belief propagation, graphs, and networks

Operationally, ecosystems of shared intelligence can be described in terms of message passing on a factor graph [57, 77, 125, 126], a special kind of graph or network in which nodes correspond to the factors of a Bayesian belief or probability distribution. Factors are just probabilistic beliefs that one multiplies together to get a joint distribution (i.e., a generative model). For example, one could factorize beliefs about the latent states of an object into “what” and “where.” These beliefs jointly specify a unique object in extrapersonal space; noting that knowing what something is and knowing where it is are largely independent of each other [127]. The edges of a factor graph correspond to the messages passed among factors that underwrite belief updating. In the implementations of active inference that we have been describing, they comprise the requisite sufficient statistics that summarize the beliefs of other nodes.

Technically, this is useful because for any generative model there is a dual or complementary factor graph that prescribes precisely the requisite message passing and implicit computational architecture. In our setting, this architecture has an interesting aspect: we can imagine the nodes of a vast graph partitioned into lots of little subgraphs. Each of these would correspond to an agent updating its beliefs via the propagation of internal messages.

Conversely, external messages would correspond to communication and belief-sharing that rests upon certain factors being distributed or duplicated over two or more subgraphs (i.e., agents or computers). This kind of architecture means that, in principle, any subgraph or agent can see, vicariously, every observable in the world—as seen through the eyes of another agent. But what is the functional and structural form of the generative model that underwrites such an architecture?

Taking our lead from human communication, the most efficient (minimum description length or minimum-complexity) generative model of worldly states should be somewhat simplified (i.e., coarse-grained), leveraging discrete representations with only as much granularity as is required to maintain an accurate account of observations [128, 129]. There are many motivations for this kind of generative model. First, it is continuous with the approach to thing-ness or individuation described above, according to which individuals are defined by the sparsity of their interactions. Concepts should evince a sparse structure, both because they are themselves “things” (and so should have sparse connections to other similar “things”), and because they are accurate representations of a world characterized by sparsity. Second, belief updating can, in this case, be implemented with simple linear operators, of the sort found in quantum computation [26, 130, 131]. Furthermore, this kind of discretization via coarse-graining moves us into the world of the theory of signs or semiotics [132, 133], Boolean logic and operators, and the sort of inference associated with abductive reasoning [134]. Finally, it finesses the form of message passing, since the sufficient statistics of discrete distributions can be reduced to a list of the relative probabilities of being in the states or levels of any given factor [135], enabling AI systems to flexibly switch contexts and acquire knowledge from others quickly and adaptively, based on a repository of shared representations.

4.3 Intelligence at scale

A subtle aspect of active inference, in this setting, is the selection of which messages to listen or attend to. In principle, this is a solved problem—in the simple case, each agent (i.e., subgraph) actively selects the messages or viewpoints that afford the greatest expected information gain [136].⁷ The neurobiological homologue of this would be attention: selecting the newsworthy information that resolves uncertainty about things you do not already know, given a certain context. There are many interesting aspects of this enactive (action-oriented) aspect of message passing; especially when thinking about nested, hierarchical structures in a global (factor) graph. In these structures—and in simulations of hierarchical processing in the brain—certain factors at higher hierarchical levels can control the selection of messages by lower levels [137, 138]. This motivates exploration of the multi-scale aspects of shared intelligence.

The emerging picture is of a message passing protocol that instantiates variational message passing on graphs of discrete belief spaces. But what must these messages contain?

⁷See [122] for more complex cases where agents have preferences for certain kinds of interaction partners, resulting in the formation of “echo chambers.”

Clearly, on the present proposal, they must contain vectors of sufficient statistics; but they also have to identify themselves in relation to the (shared) factors to which they pertain. Furthermore, they must also declare their origin, in much the same way as neuronal populations in the brain receive spatially addressed inputs from other parts of the brain.

In a synthetic setting, this calls for spatial addressing, leading to the notion of a *spatial message passing protocol and modeling language*—of the sort being developed as open standards in the Institute of Electrical and Electronics Engineers (IEEE) P2874 Spatial Web Working Group [139]. In short, the first step—toward realizing the kind of distributed, emergent, shared intelligence we have in mind—is to construct the next generation of modeling and message passing protocols, which include an irreducible spatial addressing system amenable to vectorization, and allowing for the vector-based shared representation of much of human knowledge.

5 Ethical and moral considerations

We conclude our discussion of large-scale collective intelligence with a brief discussion of the relevant areas of ethical discussion—and contention. First, it is important to note that the kind of collective intelligence evinced by eusocial insects (e.g., ant colonies), in which most individuals are merely replaceable copies of one another, is not the only paradigm for shared intelligence—nor is it a suitable one for systems in which individual nodes embody complex generative models. We believe that developing a cyber-physical network of emergent intelligence in the manner described above not only ought to, but for architectural reasons *must*, be pursued in a way that positively values and safeguards the *individuality* of people (as well as potentially non-human persons).

This idea is not new. Already in the late 1990s, before the widespread adoption of the internet as a communication technology, a future state of society had been hypothesized in which the intrinsic value of individuals is acknowledged in part because knowledge is valuable and *knowledge and life are inseparable* [140]—that is, each person has a distinct and unique life experience and, as such, knows something that no one else does. This resonates deeply with our idea that every intelligence implements a generative model of its own existence. The form of collective intelligence that we envision can emerge only from a network of essentially unique, epistemically and experientially diverse agents. This useful diversity of perspectives is a special case of functional specialization across the components of a complex system.

Much discussion in the field of AI ethics focuses on the problem of AI alignment; i.e., aligning our value systems with those of hypothetical conscious AI agents, which may possibly evince forms of super-intelligence [141–143]; for critical discussion, see [144]. This can be discussed under the broader rubric of the capacity for empathy or sympathy—what one might call *sympathetic* intelligence—which concerns the ability of agents to share aspects of their generative models, to take the perspective of other agents, and to understand the world in ways similar enough to enable coordinated action. Whether the emergence of shared intelligence in such a network structure entails the emergence of a new, collective mind is an open question.

Current state-of-the-art AI systems are largely “black boxes.” Such an approach to the design of AI ultimately puts severe limits on its transparency, explainability, and auditability. In addition, their capacity to engage in genuine collaboration with humans and other AI is limited, because they lack the ability to take the perspective of another. Moving to multi-scale active inference offers a number of technical advantages that may help address these problems. One is that leveraging explicit generative models, which carve the world into discrete latent states, may help us to identify and quantify bias in our models. Such architectures feature increased auditability, in that they are explicitly queryable and their inferences can be examined forensically—allowing us to address these biases directly. Shared generative models also effectively equip AI with a theory of mind, facilitating perspective-taking and allowing for genuinely dyadic interactions.

Much like a brain, with its many layers and connections, the multi-scale architecture for collective intelligence that we propose could be equipped with nodes and layers to enable a kind of collective self-monitoring and self-organisation of salience. However, this raises the question of authority and power: this kind of approach to the design of AI must account for the plurality and vulnerability of individual perspectives, and the need to understand and counterbalance potential abuses of power. More broadly and perhaps more fundamentally, we note that the approach to AI that we have presented here does not obviate the dangers associated with bias in AI technologies, especially when deployed at industrial scale in commercial settings, e.g., [145]. The general idea is that the deployment of AI technologies in societies that have preexisting hierarchies of power and authority can have problematic consequences. For example, discriminatory bias encoded in data will result in unfairly biased AI systems (e.g., [146]) regardless of the specific technologies used to build that AI. It is highly probable that the use of AI technologies premised on such data will sustain social biases and practices that are harmful, or may represent future harm, the consequences of which are not yet fully known—or may be unknowable—regardless of the intentions of the creators. These concerns are well founded and cannot be resolved through narrowly technical means. As such, some combination of novel social policies, government regulations, and ethical norms are likely to be required to ensure that these new technologies harness and reflect our most essential and persistent values.

We are not pessimistic. Nature provides us with endless demonstrations of the success of emergent, shared intelligence across systems at every scale. Looking back to the elegant design of the human body, we find bountiful examples of diverse systems of nested intelligences working together to seek out a dynamic harmony and balance. As an integrated system, the body is capable of achieving multi-scale homeostasis and allostasis, notably via the incredible coordination and communicative power of the nervous system, allowing it to adapt to novel environmental conditions and to regulate its needs in real time. We reiterate our conviction that the design of AI should be informed by, and aligned with, these time-tested methods and design principles. Furthermore, we believe that the class of sympathetic and shared intelligences that we have described in this paper offers a responsible and desirable path to achieving the highest technical and ethical goals for AI, based on a design of ecosystems of intelligence from first principles.

6 Conclusion: Our proposal for stages of development for active inference as an artificial intelligence technology

The aim of this white paper was to present a vision of research and development in the field of artificial intelligence for the next decade (and beyond). We suggested that AGI and ASI will emerge from the interaction of intelligences networked into a hyper-spatial web or ecosystem of natural and artificial intelligence. We have proposed active inference as a technology uniquely suited to the collaborative design of an ecosystem of natural and synthetic sense-making, in which humans are integral participants—what we call shared intelligence. The Bayesian mechanics of intelligent systems that follows from active inference led us to define intelligence operationally, as the accumulation of evidence for an agent’s generative model of their sensed world—also known as self-evidencing. This self-evidencing can be implemented using message passing or belief propagation on (factor) graphs or networks. Active inference is uniquely suited to this task because it leads to a formal account of collective intelligence. We considered the kinds of communication protocols that must be developed to enable such an ecosystem of intelligences, and argued that such considerations motivate the development of a generalized, *hyper-spatial modeling language and transaction protocol*. We suggest that establishing such common languages and protocols is a key enabling step towards an ecosystem of naturally occurring and artificial intelligences.

In closing, we provide a roadmap for developing intelligent artifacts and message passing schemes as methods or tools for the common good. This roadmap is inspired by the technology readiness levels (TRLs) that have been adopted as a framework for understanding progress in technical research and development by institutions such as the European Commission, the International Organization for Standardization (ISO), and the National Aeronautics and Space Administration agency (NASA).

6.1 Stages of development for active inference

S0: Systemic Intelligence. This is contemporary state-of-the-art AI; namely, universal function approximation—mapping from input or sensory states to outputs or action states—that optimizes some well-defined value function or cost of (systemic) states. Examples include deep learning, Bayesian reinforcement learning, etc.

S1: Sentient Intelligence. Sentient behavior or active inference based on belief updating and propagation (i.e., optimizing beliefs about states as opposed to states *per se*); where “sentient” means “responsive to sensory impressions.”⁸ This entails planning as inference; namely, inferring courses of action that maximize expected information gain and expected value, where value is part of a generative (i.e., world) model; namely, prior preferences. This

⁸To preempt any worries, we emphasize that we do *not* mean that sentient intelligent systems are necessarily *conscious*, in the sense of having qualitative states of awareness; e.g., as the word was used in the recent controversy surrounding Google’s AI system LaMDA [147]. It is standard to use the word “sentient” to mean “responsive to sensory impressions” in the literature on the free energy principle; e.g., in [39]

kind of intelligence is both information-seeking and preference-seeking. It is quintessentially curious.

S2: Sophisticated Intelligence. Sentient behavior—as defined under S1—in which plans are predicated on the consequences of action for beliefs about states of the world, as opposed to states *per se*. I.e., a move from “what will happen if I do this?” to “what will I believe or know if I do this?” [148, 149]. This kind of inference generally uses generative models with discrete states that “carve nature at its joints”; namely, inference over coarse-grained representations and ensuing world models. This kind of intelligence is amenable to formulation in terms of modal logic, quantum computation, and category theory. This stage corresponds to “artificial general intelligence” in the popular narrative about the progress of AI.

S3: Sympathetic (or Sapient) Intelligence. The deployment of sophisticated AI to recognize the nature and dispositions of users and other AI and—in consequence—recognize (and instantiate) attentional and dispositional states of self; namely, a kind of minimal selfhood (which entails generative models equipped with the capacity for Theory of Mind). This kind of intelligence is able to take the perspective of its users and interaction partners—it is perspectival, in the robust sense of being able to engage in dyadic and shared perspective-taking.

S4: Shared (or Super) Intelligence. The kind of collective that emerges from the coordination of Sympathetic Intelligence (as defined in S3) and their interaction partners or users—which may include naturally occurring intelligence such as ourselves, but also other sapient artifacts. This stage corresponds, roughly speaking, to “artificial super-intelligence” in the popular narrative about the progress of AI—with the important distinction that we believe that such intelligence will emerge from dense interactions between agents networked into a hyper-spatial web. We believe that the approach that we have outlined here is the most likely route toward this kind of hypothetical, planetary-scale, distributed super-intelligence [150].

6.2 Implementation

A: Theoretical. The basis of belief updating (i.e., inference and learning) is underwritten by a formal calculus (e.g., Bayesian mechanics), with clear links to the physics of self-organization of open systems far from equilibrium.

B: Proof of principle. Software instances of the formal (mathematical) scheme, usually on a classical (i.e., von Neumann) architecture.

C: Deployment at scale. Scaled and efficient application of the theoretical principles (i.e., methods) in a real-world setting (e.g., edge-computing, robotics, variational message passing on the web, etc.)

D: Biomimetic hardware. Implementations that elude the von Neumann bottleneck, on biomimetic or neuromorphic architectures. E.g., photonics, soft robotics, and belief propagation: i.e., message passing of the sufficient statistics of (Bayesian) beliefs.

Stage	Theoretical	Proof of principle	Deployment at scale	Biomimetic	Timeframe
S1: Sentient	Established ^{1,2}	Established ³	Provisional ⁴	Aspirational	2 years
S2: Sophisticated	Established ⁵	Provisional ⁶	Aspirational		4 years
S3: Sympathetic	Provisional ⁷	Aspirational			8 years
S4: Shared	Provisional ^{8,9}	Aspirational			16 years

Table 1: Stages of AI premised on active inference.

¹ Friston, K.J. A free energy principle for a particular physics. doi:10.48550/arXiv.1906.10184 (2019).[38]

² Ramstead, M.J.D. et al. On Bayesian Mechanics: A Physics of and by Beliefs. 10.48550/arXiv.2205.11543 (2022).[2]

³ Parr, T., Pezzulo, G. & Friston, K.J. Active Inference: The Free Energy Principle in Mind, Brain, and Behavior. (MIT Press, 2022). doi:10.7551/mitpress/12441.001.0001.[151]

⁴ Mazzaglia, P., Verbelen, T., Catal, O. & Dhoedt, B. The Free Energy Principle for Perception and Action: A Deep Learning Perspective. *Entropy* 24, 301, doi:10.3390/e24020301 (2022).[152]

⁵ Da Costa, L. et al. Active inference on discrete state-spaces: A synthesis. *Journal of Mathematical Psychology* 99, 102447, doi:10.1016/j.jmp.2020.102447 (2020).[153]

⁶ Friston, K.J., Parr, T. & de Vries, B. The graphical brain: Belief propagation and active inference. *Network Neuroscience* 1, 381-414, doi:10.1162/NETN_a_00018 (2017).[125]

⁷ Friston, K.J. et al. Generative models, linguistic communication and active inference. *Neuroscience and Biobehavioral Reviews* 118, 42-64, doi:10.1016/j.neubiorev.2020.07.005 (2020).[136]

⁸ Friston, K.J., Levin, M., Sengupta, B. & Pezzulo, G. Knowing one’s place: a free-energy approach to pattern regulation. *Journal of the Royal Society Interface* 12, doi:10.1098/rsif.2014.1383 (2015).[4]

⁹ Albarracín, M., Demekas, D., Ramstead, M.J.D. & Heins, C. Epistemic Communities under Active Inference. *Entropy* 24, doi:10.3390/e24040476 (2022).[122]

Additional information

Acknowledgements The authors thank Rosalyn Moran and George Percivall for useful discussions. Table A.1 in Appendix A has been reproduced from [153] under a CC BY 4.0 licence (<https://creativecommons.org/licenses/by/4.0/>).

Funding information All work on this paper was funded by VERSES. KF is supported by funding for the Wellcome Centre for Human Neuroimaging (Ref: 205103/Z/16/Z) and a Canada-UK Artificial Intelligence Initiative (Ref: ES/T01279X/1). CH is supported by the U.S. Office of Naval Research (Ref: N00014-19-1-2556). BK & CH acknowledge the support of a grant from the John Templeton Foundation (Ref: 61780). The opinions expressed in this publication are those of the author(s) and do not necessarily reflect the views of the John Templeton Foundation. BM was funded by Rafal Bogacz with a BBSRC grant (Ref: BB/s006338/1) and a MRC grant (Ref: MC UU 00003/1). SET is supported in part by

funding from the Social Sciences and Humanities Research Council of Canada (Ref: 767-2020-2276).

References

- [1] Nick Bostrom. *Superintelligence: Paths, Dangers, Strategies*. 1st. USA: Oxford University Press, Inc., 2014. ISBN: 0199678111. DOI: 10.5555/2678074.
- [2] Maxwell J.D. Ramstead, Dalton A.R. Sakthivadivel, Conor Heins, Magnus T. Koudahl, Beren Millidge, Lancelot Da Costa, Brennan Klein, and Karl Friston. “On Bayesian mechanics: A physics of and by beliefs”. In: *arXiv* (2022). DOI: 10.48550/arXiv.2205.11543.
- [3] Lancelot Da Costa, Karl Friston, Conor Heins, and Grigorios A. Pavliotis. “Bayesian mechanics for stationary processes”. In: *Proceedings of the Royal Society A* 477.2256 (2021). DOI: 10.1098/rspa.2021.0518.
- [4] Karl Friston, Michael Levin, Biswa Sengupta, and Giovanni Pezzulo. “Knowing one’s place: a free-energy approach to pattern regulation”. In: *Journal of The Royal Society Interface* 12.105 (2015). DOI: 10.1098/rsif.2014.1383.
- [5] Karl Friston. “Life as we know it”. In: *Journal of the Royal Society Interface* 10.86 (2013), p. 20130475. DOI: 10.1098/rsif.2013.0475.
- [6] Franz Kuchling, Karl Friston, Georgi Georgiev, and Michael Levin. “Morphogenesis as Bayesian inference: A variational approach to pattern formation and control in complex biological systems”. In: *Physics of Life Reviews* 33 (2020), pp. 88–108. DOI: 10.1016/j.plrev.2019.06.001.
- [7] Maxwell J.D. Ramstead, Casper Hesp, Alexander Tschantz, Ryan Smith, Axel Constant, and Karl Friston. “Neural and phenotypic representation under the free-energy principle”. In: *Neuroscience & Biobehavioral Reviews* 120 (2021), pp. 109–122. DOI: 10.1016/j.neubiorev.2020.11.024.
- [8] Sequoia Capital. *Generative AI: A Creative New World*. URL: <https://www.sequoiacap.com/article/generative-ai-a-creative-new-world/>.
- [9] Wolfgang Maass. “Networks of spiking neurons: The third generation of neural network models”. In: *Neural Networks* 10.9 (1997). DOI: 10.1016/S0893-6080(97)00011-7.
- [10] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. “Deep learning”. In: *Nature* 521.7553 (2015). DOI: 10.1038/nature14539.
- [11] Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. “Continuous control with deep reinforcement learning”. In: *arXiv* (2019). DOI: 10.48550/arXiv.1509.02971.
- [12] Anthony Zador et al. “Toward Next-Generation Artificial Intelligence: Catalyzing the NeuroAI Revolution”. In: *arXiv* (2022). DOI: 10.48550/arxiv.2210.08340.

- [13] Semir Zeki and Stewart Shipp. “The functional logic of cortical connections”. In: *Nature* 335.6188 (1988), pp. 311–317. DOI: 10.1038/335311a0.
- [14] Daniel L. K. Yamins, Ha Hong, Charles F. Cadieu, Ethan A. Solomon, Darren Seibert, and James J. DiCarlo. “Performance-optimized hierarchical models predict neural responses in higher visual cortex”. In: *Proceedings of the National Academy of Sciences* 111.23 (2014), pp. 8619–8624. DOI: 10.1073/pnas.1403112111.
- [15] Daniel Yamins and James J. DiCarlo. “Using goal-driven deep learning models to understand sensory cortex”. In: *Nature Neuroscience* 19 (2016), pp. 356–365. DOI: 10.1038/nn.4244.
- [16] Blake A. Richards et al. “A deep learning framework for neuroscience”. In: *Nature Neuroscience* 22.11 (2019), pp. 1761–1770. DOI: 10.1038/s41593-019-0520-2.
- [17] Nirosha J. Murugan, Daniel H. Kaltman, Paul H. Jin, Melanie Chien, Ramses Martinez, Cuong Q. Nguyen, Anna Kane, Richard Novak, Donald E. Ingber, and Michael Levin. “Mechanosensation Mediates Long-Range Spatial Decision-Making in an Aneural Organism”. In: *Advanced Materials* 33.34 (2021), p. 2008161. DOI: 10.1002/adma.202008161.
- [18] Paco Calvo and Karl Friston. “Predicting green: really radical (plant) predictive processing”. In: *Journal of the Royal Society Interface* 14.131 (2017). DOI: 10.1098/rsif.2017.0096.
- [19] Ottoline Leyser. “Auxin, Self-Organisation, and the Colonial Nature of Plants”. In: *Current Biology* 21.9 (May 2011), R331–R337. DOI: 10.1016/j.cub.2011.02.031.
- [20] Suzanne W. Simard. “Mycorrhizal Networks Facilitate Tree Communication, Learning, and Memory”. In: *Memory and Learning in Plants*. Ed. by Frantisek Baluska, Monica Gagliano, and Guenther Witzany. Cham: Springer International Publishing, 2018, pp. 191–213. ISBN: 978-3-319-75596-0. DOI: 10.1007/978-3-319-75596-0_10.
- [21] Andrew Berdahl, Colin J. Torney, Christos C. Ioannou, Jolyon J. Faria, and Iain D. Couzin. “Emergent Sensing of Complex Environments by Mobile Animal Groups”. In: *Science* 339.6119 (2013), pp. 574–576. DOI: 10.1126/science.1225883.
- [22] Chris Fields and Michael Levin. “Competency in Navigating Arbitrary Spaces as an Invariant for Analyzing Cognition in Diverse Embodiments”. In: *Entropy* 24.6 (2022), p. 819. DOI: 10.3390/e24060819.
- [23] Jamie Davies and Michael Levin. “Synthetic morphology via active and agential matter”. In: *OSF Preprints* (2022). DOI: 10.31219/osf.io/xrv8h.
- [24] Chris Fields, Johanna Bischof, and Michael Levin. “Morphological Coordination: A Common Ancestral Function Unifying Neural and Non-Neural Signaling”. In: *Physiology* 35.1 (2020), pp. 16–30. DOI: 10.1152/physiol.00027.2019.
- [25] Michael Levin. “The Computational Boundary of a “Self”: Developmental Bioelectricity Drives Multicellularity and Scale-Free Cognition”. In: *Frontiers in Psychology* 10 (2019). DOI: 10.3389/fpsyg.2019.02688.

- [26] Chris Fields, James F. Glazebrook, and Michael Levin. “Minimal physicalism as a scale-free substrate for cognition and consciousness”. In: *Neuroscience of Consciousness* 2021.2 (2021). DOI: 10.1093/nc/niab013.
- [27] William G. Lycan. “Homuncular Functionalism”. In: *Mind and Cognition: An Anthology*. Ed. by William G. Lycan and Jesse J. Prinz. Blackwell, 2008, p. 69. ISBN: 978-1-405-15784-1. URL: <https://psycnet.apa.org/record/2008-00729-000>.
- [28] Dalton A.R. Sakthivadivel. “Weak Markov Blankets in High-Dimensional, Sparsely-Coupled Random Dynamical Systems”. In: *arXiv* (2022). DOI: 10.48550/arXiv.2207.07620.
- [29] Thomas Parr, Noor Sajid, and Karl Friston. “Modules or Mean-Fields?” In: *Entropy* 22.5 (2020). DOI: 10.3390/e22050552.
- [30] Stevan Harnad. “Can a Machine Be Conscious? How?” In: *Journal of Consciousness Studies* 10.4-5 (2003), pp. 67–75. URL: <https://philpapers.org/rec/HARCAM>.
- [31] Arturo Rosenblueth, Norbert Wiener, and Julian Bigelow. “Behavior, Purpose and Teleology”. In: *Philosophy of Science* 10.1 (1943), pp. 18–24. DOI: 10.2307/184878.
- [32] W. Ross Ashby. *An Introduction to Cybernetics*. London: Chapman & Hall, 1956. URL: <http://pcp.vub.ac.be/books/IntroCyb.pdf>.
- [33] Roger C. Conant and W. Ross Ashby. “Every good regulator of a system must be a model of that system”. In: *International Journal of Systems Science* 1.2 (1970), pp. 89–97. DOI: 10.1080/00207727008920220.
- [34] Bruce A. Francis and Walter M. Wonham. “The internal model principle of control theory”. In: *Automatica* 12.5 (1976), pp. 457–465. DOI: 10.1016/0005-1098(76)90006-6.
- [35] Radford M. Neal. “MCMC using Hamiltonian dynamics”. In: *Handbook of Markov Chain Monte Carlo*. Ed. by Steve Brooks, Andrew Gelman, Galin L. Jones, and Xiao-Li Meng. Vol. 2. 11. Chapman and Hall/CRC, 2011, p. 2. DOI: 10.1201/b10905.
- [36] Jonas Degraeve et al. “Magnetic control of tokamak plasmas through deep reinforcement learning”. In: *Nature* 602.7897 (2022), pp. 414–419. DOI: 10.1038/s41586-021-04301-9.
- [37] John Jumper et al. “Highly accurate protein structure prediction with AlphaFold”. In: *Nature* 596.7873 (2021), pp. 583–589. DOI: 10.1038/s41586-021-03819-2.
- [38] Karl Friston. “A free energy principle for a particular physics”. In: *arXiv* (2019). DOI: 10.48550/arXiv.1906.10184.
- [39] Maxwell J.D. Ramstead, Axel Constant, Paul B. Badcock, and Karl Friston. “Variational ecology and the physics of sentient systems”. In: *Physics of Life Reviews* 31 (2019), pp. 188–205. DOI: 10.1016/j.plrev.2018.12.002.
- [40] Rolf Landauer. “Irreversibility and Heat Generation in the Computing Process”. In: *IBM Journal of Research and Development* 5.3 (1961). DOI: 10.1147/rd.53.0183.

- [41] Charles H. Bennett. “Notes on Landauer’s principle, reversible computation, and Maxwell’s Demon”. In: *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*. Quantum Information and Computation 34.3 (2003), pp. 501–510. DOI: 10.1016/S1355-2198(03)00039-X.
- [42] Christopher Jarzynski. “Nonequilibrium Equality for Free Energy Differences”. In: *Physical Review Letters* 78 (14 1997), pp. 2690–2693. DOI: 10.1103/PhysRevLett.78.2690.
- [43] Denis J. Evans. “A non-equilibrium free energy theorem for deterministic systems”. In: *Molecular Physics* 101.10 (2003), pp. 1551–1554. DOI: 10.1080/0026897031000085173.
- [44] Jürgen Schmidhuber. “Formal Theory of Creativity, Fun, and Intrinsic Motivation (1990–2010)”. In: *IEEE Transactions on Autonomous Mental Development* 2.3 (2010), pp. 230–247. DOI: 10.1109/TAMD.2010.2056368.
- [45] Chris S. Wallace and David L. Dowe. “Minimum Message Length and Kolmogorov Complexity”. In: *The Computer Journal* 42.4 (1999), pp. 270–283. DOI: 10.1093/comjnl/42.4.270.
- [46] David J.C. MacKay. “Free energy minimisation algorithm for decoding and cryptanalysis”. In: *Electronics Letters* 31.6 (1995), pp. 446–447. DOI: 10.1049/e1:19950331.
- [47] Bruno A. Olshausen and David J. Field. “Emergence of simple-cell receptive field properties by learning a sparse code for natural images”. In: *Nature* 381.6583 (1996), pp. 607–609. DOI: 10.1038/381607a0.
- [48] Biswa Sengupta, Martin Stemmler, Simon B. Laughlin, and Jeremy E. Niven. “Action potential energy efficiency varies among neuron types in vertebrates and invertebrates”. In: *PLoS Computational Biology* 6.7 (2010), e1000840. DOI: 10.1371/journal.pcbi.1000840.
- [49] Peter Elias. “Predictive coding–I”. In: *IRE Transactions on Information Theory* 1.1 (1955), pp. 16–24. DOI: 10.1109/TIT.1955.1055126.
- [50] Rajesh P. Rao. “An optimal estimation approach to visual perception and learning”. In: *Vision Research* 39.11 (1999), pp. 1963–1989. DOI: 10.1016/S0042-6989(98)00279-X.
- [51] Lance M. Optican and Barry J. Richmond. “Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex. III. Information theoretic analysis”. In: *Journal of Neurophysiology* 57.1 (1987), pp. 162–178. DOI: 10.1152/jn.1987.57.1.162.
- [52] Horace B. Barlow and Walter A. Rosenblith. “Possible principles underlying the transformations of sensory messages”. In: MIT Press, 1961, pp. 217–234. DOI: 10.7551/mitpress/9780262518420.003.0013.

- [53] Eero P. Simoncelli and Bruno A. Olshausen. “Natural image statistics and neural representation”. In: *Annual Review of Neuroscience* 24.1 (2001), pp. 1193–1216. DOI: 10.1146/annurev.neuro.24.1.1193.
- [54] Chris Eliasmith. “A New Perspective on Representational Problems”. In: *Journal of Cognitive Sciences* 6 (2005), pp. 97–123. URL: <http://compneuro.uwaterloo.ca/files/publications/eliasmith.2005a.pdf>.
- [55] Jakob Hohwy. “The self-evidencing brain”. In: *Nous* 50.2 (2016), pp. 259–285. DOI: 10.1111/nous.12062.
- [56] Karl Friston. “Embodied Inference: or “I think therefore I am, if I am what I think””. In: *The Implications of Embodiment: Cognition and Communication* (2011).
- [57] John Winn and Christopher M. Bishop. “Variational message passing”. In: *Journal of Machine Learning Research* 6 (2005), pp. 661–694. URL: <https://www.jmlr.org/papers/volume6/winn05a/winn05a.pdf>.
- [58] Biswa Sengupta and Karl Friston. “How Robust are Deep Neural Networks?” In: *arXiv* (2018). DOI: 10.48550/arXiv.1804.11313.
- [59] Hilbert J. Kappen, Vicenç Gomez, and Manfred Opper. “Optimal control as a graphical model inference problem”. In: *Machine Learning* 87.2 (2012), pp. 159–182. DOI: 10.1007/s10994-012-5278-7.
- [60] Emanuel Todorov. “General duality between optimal control and estimation”. In: *2008 47th IEEE Conference on Decision and Control*. 2008, pp. 4286–4292. DOI: 10.1109/CDC.2008.4739438.
- [61] David J.C. Mackay. “Information-Based Objective Functions for Active Data Selection”. In: *Neural Computation* 4.4 (1992), pp. 590–604. DOI: 10.1162/neco.1992.4.4.590.
- [62] Dennis V. Lindley. “On a Measure of the Information Provided by an Experiment”. In: *Annals of Mathematical Statistics* 27.4 (1956), pp. 986–1005. DOI: 10.1214/aoms/1177728069.
- [63] Stefano Balietti, Brennan Klein, and Christoph Riedl. “Optimal design of experiments to identify latent behavioral types”. In: *Experimental Economics* 24.3 (2021), pp. 772–799. DOI: 10.1007/s10683-020-09680-w.
- [64] James O. Berger. *Statistical Decision Theory and Bayesian Analysis*. Springer Series in Statistics. New York, NY: Springer, 1985. ISBN: 978-1-4419-3074-3. DOI: 10.1007/978-1-4419-3074-3.
- [65] Philipp Schwartenbeck, Johannes Passecker, Tobias U. Hauser, Thomas FitzGerald, Martin Kronbichler, and Karl Friston. “Computational mechanisms of curiosity and goal-directed exploration”. In: *Elife* 8 (2019), e41703. DOI: 10.7554/eLife.41703.
- [66] Jürgen Schmidhuber. “Developmental robotics, optimal artificial curiosity, creativity, music, and the fine arts”. In: *Connection Science* 18.2 (2006), pp. 173–187. DOI: 10.1080/09540090600768658.

- [67] Susanne. Still and Diona Precup. “An information-theoretic approach to curiosity-driven reinforcement learning”. In: *Theory in Biosciences* 131.3 (2012), pp. 139–48. DOI: 10.1007/s12064-011-0142-z.
- [68] Andrew Barto, Marco Mirolli, and Gianluca Baldassarre. “Novelty or surprise?” In: *Frontiers in Psychology* 4 (2013), p. 907. DOI: 10.3389/fpsyg.2013.00907.
- [69] Daniel C. Dennett. “Beyond Belief”. In: *Thought and Object*. Ed. by Andrew Woodfield. Oxford University Press, 1983. URL: <https://philpapers.org/rec/DENBB>.
- [70] Gavin E. Crooks. “Measuring thermodynamic length”. In: *Physical Review Letters* 99.10 (2007), p. 100602. DOI: 10.1103/PhysRevLett.99.100602.
- [71] Eun-jin Kim. “Investigating Information Geometry in Classical and Quantum Systems through Information Length”. In: *Entropy* 20.8 (2018). DOI: 10.3390/e20080574.
- [72] Nihat Ay. “Information Geometry on Complexity and Stochastic Interaction”. In: *Entropy* 17.4 (2015), pp. 2432–2458. DOI: 10.3390/e17042432.
- [73] Shun-ichi Amari. “Natural gradient works efficiently in learning”. In: *Neural Computation* 10.2 (1998), pp. 251–276. DOI: 10.1162/089976698300017746.
- [74] Ariel Caticha. “The basics of information geometry”. In: *AIP Conference Proceedings* 1641.1 (2015), pp. 15–26. DOI: 10.1063/1.4905960.
- [75] Thomas Parr, Lancelot Da Costa, and Karl Friston. “Markov blankets, information geometry and stochastic thermodynamics”. In: *Philosophical Transactions of the Royal Society A* 378.2164 (2020), p. 20190159. DOI: 10.1098/rsta.2019.0159.
- [76] Frank R. Kschischang, Brendan J. Frey, and Hans-Andrea Loeliger. “Factor graphs and the sum-product algorithm”. In: *IEEE Transactions on Information Theory* 47.2 (2001), pp. 498–519. DOI: 10.1109/18.910572.
- [77] Justin Dauwels. “On Variational Message Passing on Factor Graphs”. In: *2007 IEEE International Symposium on Information Theory*, pp. 2546–2550. DOI: 10.1109/ISIT.2007.4557602.
- [78] Karl J. Friston, Richard Rosch, Thomas Parr, Cathy Price, and Howard Bowman. “Deep temporal models and active inference”. In: *Neuroscience & Biobehavioral Reviews* 90 (2018), pp. 486–501. DOI: 10.1016/j.neubiorev.2018.04.004.
- [79] Thomas Parr, Dimitrije Markovic, Stefan J. Kiebel, and Karl Friston. “Neuronal message passing using Mean-field, Bethe, and Marginal approximations”. In: *Scientific Reports* 9.1 (2019), p. 1889. DOI: 10.1038/s41598-018-38246-3.
- [80] Harriet Feldman and Karl Friston. “Attention, uncertainty, and free-energy”. In: *Frontiers in Human Neuroscience* 4 (2010), p. 215. DOI: 10.3389/fnhum.2010.00215.
- [81] Jakob Hohwy. “Attention and conscious perception in the hypothesis testing brain”. In: *Frontiers in Psychology* 3 (2012), p. 96. DOI: 10.3389/fpsyg.2012.00096.

- [82] Peter Kok, Dobromir Rahnev, Janneke F. M. Jehee, Hakwan C. Lau, and Floris P. de Lange. “Attention Reverses the Effect of Prediction in Silencing Sensory Signals”. In: *Cerebral Cortex* 22.9 (2011), pp. 2197–2206. DOI: 10.1093/cercor/bhr310.
- [83] Ryota Kanai, Yutaka Komura, Stewart Shipp, and Karl Friston. “Cerebral hierarchies: predictive processing, precision and the pulvinar”. In: *Philosophical Transactions of the Royal Society B* 370.1668 (2015), p. 20140169. DOI: 10.1098/rstb.2014.0169.
- [84] Jakub Limanowski. “Precision control for a flexible body representation”. In: *Neuroscience and Biobehavioral Reviews* 134 (2022), p. 104401. DOI: 10.1016/j.neubiorev.2021.10.023.
- [85] Samuel J. Gershman and Yael Niv. “Learning latent structure: carving nature at its joints”. In: *Current Opinions in Neurobiology* 20.2 (2010), pp. 251–6. DOI: 10.1016/j.conb.2010.02.008.
- [86] Joshua B. Tenenbaum, Charles Kemp, Thomas L. Griffiths, and Noah D. Goodman. “How to grow a mind: statistics, structure, and abstraction”. In: *Science* 331.6022 (2011), pp. 1279–85. DOI: 10.1126/science.1192788.
- [87] David J. Spiegelhalter, Nicola G. Best, Bradley R. Carlin, and Angelika van der Linde. “Bayesian measures of model complexity and fit”. In: *Journal of the Royal Statistical Society Series B-Statistical Methodology* 64.3 (2002), pp. 583–616. DOI: 10.1111/1467-9868.00353.
- [88] William D. Penny. “Comparing Dynamic Causal Models using AIC, BIC and Free Energy”. In: *Neuroimage* 59.1 (2012), pp. 319–330. DOI: 10.1016/j.neuroimage.2011.07.039.
- [89] Karl Friston, Thomas Parr, and Peter Zeidman. “Bayesian model reduction”. In: *arXiv* (2018). DOI: 10.48550/arXiv.1805.07092.
- [90] Takuya Isomura, Hideaki Shimazaki, and Karl Friston. “Canonical neural networks perform active inference”. In: *Communications Biology* 5.1 (2022), pp. 1–15. DOI: 10.1038/s42003-021-02994-2.
- [91] Sepp Hochreiter and Jürgen Schmidhuber. “Long Short-Term Memory”. In: *Neural Computation* 9.8 (1997), pp. 1735–1780. DOI: 10.1162/neco.1997.9.8.1735.
- [92] Yann LeCun. “A Path Towards Autonomous Machine Intelligence”. In: *OpenReview* (2022). URL: <https://openreview.net/forum?id=BZ5a1r-kVsf>.
- [93] Irina Higgins, David Amos, David Pfau, Sébastien Racanière, Loïc Matthey, Danilo Jimenez Rezende, and Alexander Lerchner. “Towards a Definition of Disentangled Representations”. In: *arXiv* (2018). DOI: 10.48550/arXiv.1812.02230.
- [94] Alexander Tschantz, Anil K. Seth, and Christopher L. Buckley. “Learning action-oriented models through active inference”. In: *PLoS Computational Biology* 16.4 (2020), e1007805. DOI: 10.1371/journal.pcbi.1007805.

- [95] Toon Van de Maele, Tim Verbelen, Ozan Çatal, and Bart Dhoedt. “Disentangling What and Where for 3D Object-Centric Representations Through Active Inference”. In: *Communications in Computer and Information Science*. 2021. DOI: 10.1007/978-3-030-93736-2_50.
- [96] Geoffrey E. Hinton, Alex Krizhevsky, and Sida I. Wang. “Transforming Auto-Encoders”. In: *Artificial Neural Networks and Machine Learning – ICANN 2011*. Ed. by T. Honkela, W. Duch, M. Girolami, and S. Kaski. 2011. DOI: 10.1007/978-3-642-21735-7_6.
- [97] Gerd Gigerenzer. “Moral Satisficing: Rethinking Moral Behavior as Bounded Rationality”. In: *Topics in Cognitive Science* 2.3 (2010), pp. 528–554. DOI: 10.1111/j.1756-8765.2010.01094.x.
- [98] David Krakauer, Nils Bertschinger, Eckehard Olbrich, Jessica Flack, and Nihat Ay. “The information theory of individuality”. In: *Theory in Biosciences* 139.2 (2020), pp. 209–223. DOI: 10.1007/s12064-020-00313-7.
- [99] Saining Xie, Ross Girshick, Piotr Dollar, Z. Tu, and Kaiming He. “Aggregated Residual Transformations for Deep Neural Networks”. In: 2017, pp. 5987–5995. DOI: 10.1109/CVPR.2017.634.
- [100] Sara Sabour, Nicholas Frosst, and Geoffrey E. Hinton. “Dynamic Routing between Capsules”. In: *Proceedings of the 31st International Conference on Neural Information Processing Systems*. NIPS’17. Long Beach, California, USA: Curran Associates Inc., 2017, pp. 3859–3869. DOI: 10.5555/3294996.3295142.
- [101] Jeff Hawkins, Subutai Ahmad, and Yuwei Cui. “A Theory of How Columns in the Neocortex Enable Learning the Structure of the World”. In: *Frontiers in Neural Circuits* 11 (2017). DOI: 10.3389/fncir.2017.00081.
- [102] Geoffrey Hinton. “Training products of experts by minimizing contrastive divergence”. In: *Neural Computation* 14.8 (2002), pp. 1771–800. DOI: 10.1162/089976602760128018.
- [103] Axel Constant, Maxwell J.D. Ramstead, Samuel P.L. Veissière, John O. Campbell, and Karl Friston. “A variational approach to niche construction”. In: *Journal of The Royal Society Interface* 15.141 (2018), p. 20170685. DOI: 10.1098/rsif.2017.0685.
- [104] Samuel P.L. Veissière, Axel Constant, Maxwell J.D. Ramstead, Karl Friston, and Laurence J. Kirmayer. “Thinking through other minds: A variational approach to cognition and culture”. In: *Behavioral and Brain Sciences* 43 (2020), e90. DOI: 10.1017/S0140525X19001213.
- [105] Kevin N. Laland, F. John Odling-Smee, and Marcus W. Feldman. “Evolutionary consequences of niche construction and their implications for ecology”. In: *Proceedings of the National Academy of Sciences* 96.18 (1999), pp. 10242–10247. DOI: 10.1073/pnas.96.18.10242.

- [106] Axel Constant, Maxwell J.D. Ramstead, Samuel P.L. Veissiere, and Karl Friston. “Regimes of Expectations: An Active Inference Model of Social Conformity and Human Decision Making”. In: *Frontiers in Psychology* 10 (2019), p. 679. DOI: 10.3389/fpsyg.2019.00679.
- [107] Jared Vasil, Paul B. Badcock, Axel Constant, Karl Friston, and Maxwell J.D. Ramstead. “A World Unto Itself: Human Communication as Active Inference”. In: *Frontiers in Psychology* 11 (2020), p. 417. DOI: 10.3389/fpsyg.2020.00417.
- [108] Nabil Bouizegarene, Maxwell J.D. Ramstead, Axel Constant, Karl Friston, and Laurence Kirmayer. “Narrative as active inference”. In: *PsyArXiv* (2020). DOI: 10.31234/osf.io/47ub6.
- [109] Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. “Pre-Train, Prompt, and Predict: A Systematic Survey of Prompting Methods in Natural Language Processing”. In: *ACM Computing Surveys* (2022). DOI: 10.1145/3560815.
- [110] Mingkai Deng, Jianyu Wang, Cheng-Ping Hsieh, Yihan Wang, Han Guo, Tianmin Shu, Meng Song, Eric P. Xing, and Zhiting Hu. “RLPrompt: Optimizing Discrete Text Prompts With Reinforcement Learning”. In: *arXiv* (2022). DOI: 10.48550/arXiv.2205.12548.
- [111] Michael Tomasello. “Cultural Learning Redux”. In: *Child Development* 87.3 (2016), pp. 643–53. DOI: 10.1111/cdev.12499.
- [112] Cecilia M. Heyes. *Cognitive Gadgets: The Cultural Evolution of Thinking*. Harvard University Press, 2018. ISBN: 9780674980150. DOI: 10.4159/9780674985155.
- [113] Joseph Henrich. *The Secret of Our Success: How Culture Is Driving Human Evolution, Domesticating Our Species, and Making Us Smarter*. Princeton University Press, 2016. DOI: 10.2307/j.ctvc77f0d.
- [114] Yonghui Wu et al. “Google’s Neural Machine Translation System: Bridging the Gap between Human and Machine Translation”. In: *CoRR* abs/1609.08144 (2016). arXiv: 1609.08144. URL: <http://arxiv.org/abs/1609.08144>.
- [115] Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel P. Kuksa. “Natural Language Processing (almost) from Scratch”. In: *CoRR* abs/1103.0398 (2011). arXiv: 1103.0398. URL: <http://arxiv.org/abs/1103.0398>.
- [116] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. “Distributed Representations of Words and Phrases and their Compositionality”. In: *CoRR* abs/1310.4546 (2013). arXiv: 1310.4546. URL: <http://arxiv.org/abs/1310.4546>.
- [117] Albert Webson and Ellie Pavlick. “Do Prompt-Based Models Really Understand the Meaning of Their Prompts?” In: *ArXiv* abs/2109.01247 (2022).
- [118] Thomas H. FitzGerald, Raymond J. Dolan, and Karl Friston. “Model averaging, optimal inference, and habit formation”. In: *Frontiers in Human Neuroscience* 8 (2014), p. 457. DOI: 10.3389/fnhum.2014.00457.

- [119] Maxwell J.D. Ramstead, Samuel P.L. Veissière, and Laurence J. Kirmayer. “Cultural Affordances: Scaffolding Local Worlds Through Shared Intentionality and Regimes of Attention”. In: *Frontiers in Psychology* 7 (2016). DOI: 10.3389/fpsyg.2016.01090.
- [120] Karl Friston and Christopher Frith. “A duet for one”. In: *Consciousness and Cognition* 36 (2015), pp. 390–405. DOI: 10.1016/j.concog.2014.12.003.
- [121] Takuya Isomura, Thomas Parr, and Karl Friston. “Bayesian Filtering with Multiple Internal Models: Toward a Theory of Social Intelligence”. In: *Neural Computation* 31.12 (2019), pp. 2390–2431. DOI: 10.1162/neco_a_01239.
- [122] Mahault Albarracín, Daphne Demekas, Maxwell J.D. Ramstead, and Conor Heins. “Epistemic communities under active inference”. In: *Entropy* 24.4 (2022), p. 476. DOI: 10.3390/e24040476.
- [123] Natalie Kastel and Casper Hesp. “Ideas Worth Spreading: A Free Energy Proposal for Cumulative Cultural Dynamics”. In: *Machine Learning and Principles and Practice of Knowledge Discovery in Databases*. Ed. by Michael Kamp et al. Springer International Publishing, pp. 784–798. DOI: 10.1007/978-3-030-93736-2_55.
- [124] Ensor R. Palacios, Takuya Isomura, Thomas Parr, and Karl Friston. “The emergence of synchrony in networks of mutually inferring neurons”. In: *Scientific Reports* 9.1 (2019), p. 6412. DOI: 10.1038/s41598-019-42821-7.
- [125] Karl Friston, Thomas Parr, and Bert de Vries. “The graphical brain: Belief propagation and active inference”. In: *Network Neuroscience* 1.4 (2017), pp. 381–414. DOI: 10.1162/NETN_a_00018.
- [126] Jonathan S. Yedidia, William T. Freeman, and Yair Weiss. “Constructing free-energy approximations and generalized belief propagation algorithms”. In: *IEEE Transactions on Information Theory* 51.7 (2005), pp. 2282–2312. DOI: 10.1109/TIT.2005.850085.
- [127] Leslie G. Ungerleider and James V. Haxby. “‘What’ and ‘where’ in the human brain”. In: *Current Opinion in Neurobiology* 4.2 (1994), pp. 157–165. DOI: [https://doi.org/10.1016/0959-4388\(94\)90066-3](https://doi.org/10.1016/0959-4388(94)90066-3).
- [128] Conor Heins, Brennan Klein, Daphne Demekas, Miguel Aguilera, and Christopher Buckley. “Spin glass systems as collective active inference”. In: *International Workshop on Active Inference* (2022). DOI: 10.48550/arXiv.2207.06970.
- [129] Brennan Klein and Erik Hoel. “The emergence of informative higher scales in complex networks”. In: *Complexity* (2020), p. 8932526. DOI: 10.1155/2020/8932526.
- [130] Chris Fields, Karl Friston, James F. Glazebrook, and Michael Levin. “A free energy principle for generic quantum systems”. In: *Progress in Biophysics and Molecular Biology* 173 (2022), pp. 36–59. DOI: 10.1016/j.pbiomolbio.2022.05.006.
- [131] Juan M.R. Parrondo, Jordan M. Horowitz, and Takahiro Sagawa. “Thermodynamics of information”. In: *Nature Physics* 11.2 (2015), pp. 131–139. DOI: 10.1038/Nphys3230.

- [132] Deb Roy. “Semiotic schemas: A framework for grounding language in action and perception”. In: *Artificial Intelligence* 167.1-2 (2005), pp. 170–205. DOI: 10.1016/j.artint.2005.04.007.
- [133] William H. Sewell. “A Theory of Structure: Duality, Agency, and Transformation”. In: *American Journal of Sociology* 98.1 (1992), pp. 1–29. DOI: 10.2307/2781191.
- [134] Charles Sanders Peirce. *Collected Papers of Charles Sanders Peirce*. Collected Papers of Charles Sanders Peirce v. 5. Harvard University Press, 1931. ISBN: 978-0-674-13802-5. URL: <https://books.google.com/books?id=USgPAQAAIAAJ>.
- [135] Zoubin Ghahramani and Michael I. Jordan. “Factorial Hidden Markov Models”. In: *Machine Learning* 29.2 (1997), pp. 245–273. DOI: 10.1023/A:1007425814087.
- [136] Karl Friston, Thomas Parr, Yan Yufik, Noor Sajid, Catherine J. Price, and Emma Holmes. “Generative models, linguistic communication and active inference”. In: *Neuroscience and Biobehavioral Reviews* 118 (2020), pp. 42–64. DOI: 10.1016/j.neubiorev.2020.07.005.
- [137] Thomas Parr and Karl Friston. “Working memory, attention, and salience in active inference”. In: *Scientific Reports* 7.1 (2017), p. 14678. DOI: 10.1038/s41598-017-15249-0.
- [138] Ryan Smith, Thomas Parr, and Karl Friston. “Simulating Emotions: An Active Inference Model of Emotional State Inference and Emotion Concept Learning”. In: *Frontiers in Psychology* 10 (2019), p. 2844. DOI: 10.3389/fpsyg.2019.02844.
- [139] *Standard for Spatial Web Protocol, Architecture and Governance*. URL: <https://standards.ieee.org/ieee/2874/10375/>.
- [140] Pierre Levy and Robert Bononno. *Collective Intelligence: Mankind’s Emerging World in Cyberspace*. USA: Perseus Books, 1997. ISBN: 0306456354. DOI: 10.5555/550283.
- [141] Stuart Russell. *Human compatible: Artificial intelligence and the problem of control*. Viking, 2019. ISBN: 978-0-525-55861-3.
- [142] Colin Allen, Iva Smit, and Wendell Wallach. “Artificial morality: Top-down, bottom-up, and hybrid approaches”. In: *Ethics and Information Technology* 7.3 (2005), pp. 149–155. DOI: 10.1007/s10676-006-0004-4.
- [143] Stuart Russell, Tom Dietterich, Eric Horvitz, Bart Selman, Francesca Rossi, Demis Hassabis, Shane Legg, Mustafa Suleyman, Dileep George, and Scott Phoenix. “Letter to the editor: Research priorities for robust and beneficial artificial intelligence: An open letter”. In: *AI Magazine* 36.4 (2015), pp. 3–4. DOI: 10.1609/aimag.v36i4.2621.
- [144] Gary Marcus and Ernest Davis. *Rebooting AI: Building artificial intelligence we can trust*. Pantheon, 2019. ISBN: 9781524748258. URL: <http://rebooting.ai/>.
- [145] Abeba Birhane. “Algorithmic injustice: A relational ethics approach”. In: *Patterns* 2.2 (2021), p. 100205. DOI: 10.1016/j.patter.2021.100205.

- [146] Abeba Birhane. “The unseen Black faces of AI algorithms”. In: *Nature News and Views* 610 (2022), pp. 451–452. DOI: 10.1038/d41586-022-03050-7.
- [147] Romal Thoppilan et al. “LaMDA: Language Models for Dialog Applications”. In: *arXiv* (2022). DOI: 10.48550/arxiv.2201.08239.
- [148] Karl Friston, Lancelot Da Costa, Danijar Hafner, Casper Hesp, and Thomas Parr. “Sophisticated Inference”. In: *Neural Computation* 33.3 (2021), pp. 713–763. DOI: 10.1162/neco_a_01351.
- [149] Casper Hesp, Alexander Tschantz, Beren Millidge, Maxwell Ramstead, Karl Friston, and Ryan Smith. “Sophisticated affective inference: simulating anticipatory affective dynamics of imagining future events”. In: *International Workshop on Active Inference*. Springer. 2020, pp. 179–186. DOI: 10.1007/978-3-030-64919-7_18.
- [150] Adam Frank, David Grinspoon, and Sara Walker. “Intelligence as a planetary scale process”. In: *International Journal of Astrobiology* 21.2 (2022), pp. 47–61. DOI: 10.1017/S147355042100029X.
- [151] Thomas Parr, Giovanni Pezzulo, and Karl Friston. *Active Inference: The Free Energy Principle in Mind, Brain, and Behavior*. MIT Press, 2022. ISBN: 9780262369978. DOI: 10.7551/mitpress/12441.001.0001.
- [152] Pietro Mazzaglia, Tim Verbelen, Ozan Çatal, and Bart Dhoedt. “The Free Energy Principle for Perception and Action: A Deep Learning Perspective”. In: *Entropy* 24.2 (2022). DOI: 10.3390/e24020301.
- [153] Lancelot Da Costa, Thomas Parr, Noor Sajid, Sebastijan Veselic, Victorita Neacsu, and Karl Friston. “Active inference on discrete state-spaces: A synthesis”. In: *Journal of Mathematical Psychology* 99 (2020), p. 102447. DOI: 10.1016/j.jmp.2020.102447.

A Appendix: Applications of active inference

Table A.1: Examples of Active Inference implementations. From Da Costa et al. (2020) [1]

Application	Description	References
Decision-making under uncertainty	Initial formulation of active inference on partially observable Markov decision processes.	Friston, Samothrakis et al. (2012) [2]
Optimal control	Application of KL or risk sensitive control in an engineering benchmark—the mountain car problem.	Çatal et al. (2019) [3] and Friston, Adams et al. (2012) [4]
Evidence accumulation	Illustrating the role of evidence accumulation in decision-making through an urns task.	FitzGerald, Moran et al. (2015) [5] and FitzGerald, Schwartenbeck et al. (2015) [6]
Psychopathology	Simulation of addictive choice behaviour.	Schwartenbeck, FitzGerald, Mathys, Dolan, Wurst et al. (2015) [7]
Dopamine	The precision of beliefs about policies provides a plausible description of dopaminergic discharges.	Friston et al. (2014) [8] and FitzGerald, Dolan et al. (2015) [9]
Functional magnetic resonance imaging	Empirical prediction and validation of dopaminergic discharges.	Schwartenbeck, FitzGerald, Mathys, Dolan, Friston (2015) [10]
Maximal utility theory	Evidence in favor of surprise minimization as opposed to utility maximization in human decision-making.	Schwartenbeck, FitzGerald, Mathys, Dolan, Kronbichler et al. (2015) [11]
Social cognition	Examining the effect of prior preferences on interpersonal inference.	Moutoussis et al. (2014) [12]

Exploration-exploitation dilemma	Casting behavior as expected free energy minimizing accounts for epistemic and pragmatic choices.	Friston et al. (2015) [13]
Habit learning and action selection	Formulating learning as an inferential process and action selection as Bayesian model averaging.	Friston et al. (2016) [14] and FitzGerald et al. (2014) [15]
Scene construction and anatomy of time	Mean-field approximation for multi-factorial hidden states, enabling high dimensional representations of the environment.	Friston and Buzsáki (2016) [16] and Mirza et al. (2016) [17]
Electrophysiological responses	Synthesizing various in-silico neurophysiological responses via a gradient descent on free energy. E.g., place-cell activity, mismatch negativity, phase-precession, theta sequences, theta-gamma coupling and dopaminergic discharges.	Friston, FitzGerald et al. (2017) [18]
Structure learning, curiosity and insight	Simulation of artificial curiosity and abstract rule learning. Structure learning via Bayesian model reduction.	Friston, Lin et al. (2017) [19]
Hierarchical temporal representations	Generalization to hierarchical generative models with deep temporal structure and simulation of reading.	Friston et al. (2018b) [20] and Parr and Friston (2017b) [21]
Computational neuropsychology	Simulation of visual neglect, hallucinations, and prefrontal syndromes under alternative pathological priors.	Benrimoh et al. (2018) [22], Parr, Benrimoh et al. (2018) [23], Parr and Friston (2018c) [24], Parr, Rees et al. (2018) [25] and Parr, Rikhye et al. (2019) [26]

Neuromodulation	Use of precision parameters to manipulate exploration during saccadic searches; associating uncertainty with cholinergic and noradrenergic systems.	Parr and Friston (2017a) [27], Parr and Friston (2019) [28], Sales et al. (2019) [29] and Vincent et al. (2019) [30]
Decisions to movements	Mixed generative models combining discrete and continuous states to implement decisions through movement.	Friston, Parr et al. (2017) [31] and Parr and Friston (2018d) [32]
Planning, navigation and niche construction	Agent induced changes in environment (generative process); decomposition of goals into subgoals.	Bruineberg et al. (2018) [33], Constant et al. (2018) [34] and Kaplan and Friston (2018a) [35]
Atari games	Active inference compares favorably to reinforcement learning in the game of Doom.	Cullen et al. (2018) [36]
Machine learning	Scaling active inference to more complex machine learning problems.	Tschantz et al. (2019) [37]

Supplemental References

- [1] Lancelot Da Costa, Thomas Parr, Noor Sajid, Sebastijan Veselic, Victorita Neacsu, and Karl Friston. “Active inference on discrete state-spaces: A synthesis”. In: *Journal of Mathematical Psychology* 99 (2020), p. 102447. DOI: 10.1016/j.jmp.2020.102447.
- [2] Karl Friston, Spyridon Samothrakis, and Read Montague. “Active inference and agency: optimal control without cost functions”. In: *Biological Cybernetics* 106.8 (2012), pp. 523–541. DOI: 10.1007/s00422-012-0512-8.
- [3] Ozan Çatal, Tim Verbelen, Johannes Nauta, Cedric De Boom, and Bart Dhoedt. “Learning Perception and Planning With Deep Active Inference”. In: *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 2020, pp. 3952–3956. DOI: 10.1109/ICASSP40776.2020.9054364.
- [4] Karl Friston, Rick Adams, and Read Montague. “What is value—accumulated reward or evidence?” In: *Frontiers in Neurorobotics* 6 (2012). DOI: 10.3389/fnbot.2012.00011.

- [5] Thomas FitzGerald, Rosalyn J. Moran, Karl Friston, and Raymond J. Dolan. “Precision and neuronal dynamics in the human posterior parietal cortex during evidence accumulation”. In: *NeuroImage* 107 (2015), pp. 219–228. DOI: <https://doi.org/10.1016/j.neuroimage.2014.12.015>.
- [6] Thomas FitzGerald, Philipp Schwartenbeck, Michael Moutoussis, Raymond J. Dolan, and Karl Friston. “Active Inference, Evidence Accumulation, and the Urn Task”. In: *Neural Computation* 27.2 (2015), pp. 306–328. DOI: 10.1162/NECO_a_00699.
- [7] Philipp Schwartenbeck, Thomas FitzGerald, Christoph Mathys, Ray Dolan, Friedrich Wurst, Martin Kronbichler, and Karl Friston. “Optimal inference with suboptimal models: Addiction and active Bayesian inference”. In: *Medical Hypotheses* 84.2 (2015), pp. 109–117. DOI: 10.1016/j.mehy.2014.12.007.
- [8] Karl Friston, Philipp Schwartenbeck, Thomas FitzGerald, Michael Moutoussis, Timothy Behrens, and Raymond J. Dolan. “The anatomy of choice: dopamine and decision-making”. In: *Philosophical Transactions of the Royal Society B* 369.1655 (2014), p. 20130481. DOI: 10.1098/rstb.2013.0481.
- [9] Thomas FitzGerald, Raymond J. Dolan, and Karl Friston. “Dopamine, reward learning, and active inference”. In: *Frontiers in Computational Neuroscience* 9 (2015). DOI: 10.3389/fncom.2015.00136.
- [10] Philipp Schwartenbeck, Thomas FitzGerald, Christoph Mathys, Ray Dolan, and Karl Friston. “The Dopaminergic Midbrain Encodes the Expected Certainty about Desired Outcomes”. In: *Cerebral Cortex* 25.10 (2014), pp. 3434–3445. ISSN: 1047-3211. DOI: 10.1093/cercor/bhu159.
- [11] Philipp Schwartenbeck, Thomas FitzGerald, Christoph Mathys, Ray Dolan, Martin Kronbichler, and Karl Friston. “Evidence for surprise minimization over value maximization in choice behavior”. In: *Scientific Reports* 5.1 (2015), p. 16575. DOI: 10.1038/srep16575.
- [12] Michael Moutoussis, Nelson Trujillo-Barreto, Wael El-Dereby, Raymond Dolan, and Karl Friston. “A formal model of interpersonal inference”. In: *Frontiers in Human Neuroscience* 8 (2014). DOI: 10.3389/fnhum.2014.00160.
- [13] Karl Friston, Francesco Rigoli, Dimitri Ognibene, Christoph Mathys, Thomas FitzGerald, and Giovanni Pezzulo. “Active inference and epistemic value”. In: *Cognitive Neuroscience* 6.4 (2015), pp. 187–214. DOI: 10.1080/17588928.2015.1020053.
- [14] Karl Friston, Thomas FitzGerald, Francesco Rigoli, Philipp Schwartenbeck, John O’Doherty, and Giovanni Pezzulo. “Active inference and learning”. In: *Neuroscience & Biobehavioral Reviews* 68 (2016), pp. 862–879. DOI: 10.1016/j.neubiorev.2016.06.022.
- [15] Thomas H. FitzGerald, Raymond J. Dolan, and Karl Friston. “Model averaging, optimal inference, and habit formation”. In: *Frontiers in Human Neuroscience* 8 (2014), p. 457. DOI: 10.3389/fnhum.2014.00457.

- [16] Karl Friston and Gyorgy Buzsáki. “The Functional Anatomy of Time: What and When in the Brain”. In: *Trends in Cognitive Sciences* 20.7 (2016), pp. 500–511. DOI: 10.1016/j.tics.2016.05.001.
- [17] M. Berk Mirza, Rick A. Adams, Christoph D. Mathys, and Karl Friston. “Scene Construction, Visual Foraging, and Active Inference”. In: *Frontiers in Computational Neuroscience* 10 (2016). DOI: 10.3389/fncom.2016.00056.
- [18] Karl Friston, Thomas FitzGerald, Francesco Rigoli, Philipp Schwartenbeck, and Giovanni Pezzulo. “Active Inference: A Process Theory”. In: *Neural Computation* 29.1 (2017), pp. 1–49. DOI: 10.1162/NECO_a_00912.
- [19] Karl Friston, Marco Lin, Christopher D. Frith, Giovanni Pezzulo, J. Allan Hobson, and Sasha Ondobaka. “Active Inference, Curiosity and Insight”. In: *Neural Computation* 29.10 (2017), pp. 2633–2683. DOI: 10.1162/neco_a_00999.
- [20] Karl J. Friston, Richard Rosch, Thomas Parr, Cathy Price, and Howard Bowman. “Deep temporal models and active inference”. In: *Neuroscience & Biobehavioral Reviews* 90 (2018), pp. 486–501. DOI: 10.1016/j.neubiorev.2018.04.004.
- [21] Thomas Parr and Karl Friston. “Working memory, attention, and salience in active inference”. In: *Scientific Reports* 7.1 (2017), p. 14678. DOI: 10.1038/s41598-017-15249-0.
- [22] David A. Benrimoh, Thomas Parr, Peter Vincent, Rick A. Adams, and Karl Friston. “Active Inference and Auditory Hallucinations”. In: *Computational Psychiatry* 2 (2018), pp. 183–204. DOI: 10.1162/CPSY_a_00022.
- [23] Thomas Parr, David A. Benrimoh, Peter Vincent, and Karl Friston. “Precision and False Perceptual Inference”. In: *Frontiers in Integrative Neuroscience* 12 (2018). DOI: 10.3389/fnint.2018.00039.
- [24] Thomas Parr and Karl Friston. “The Computational Anatomy of Visual Neglect”. In: *Cerebral Cortex* 28.2 (2017), pp. 777–790. DOI: 10.1093/cercor/bhx316.
- [25] Thomas Parr, Geraint Rees, and Karl Friston. “Computational Neuropsychology and Bayesian Inference”. In: *Frontiers in Human Neuroscience* 12 (2018). DOI: 10.3389/fnhum.2018.00061.
- [26] Thomas Parr, Rajeev Vijay Rikhye, Michael M. Halassa, and Karl Friston. “Prefrontal Computation as Active Inference”. In: *Cerebral Cortex* 30.2 (2020), pp. 682–695. DOI: 10.1093/cercor/bhz118.
- [27] Thomas Parr and Karl J. Friston. “Uncertainty, epistemics and active inference”. In: *Journal of The Royal Society Interface* 14.136 (2017), p. 20170376. DOI: 10.1098/rsif.2017.0376.
- [28] Thomas Parr and Karl Friston. “The computational pharmacology of oculomotion”. In: *Psychopharmacology* 236.8 (2019), pp. 2473–2484. DOI: 10.1007/s00213-019-05240-0.

- [29] Anna C. Sales, Karl Friston, Matthew W. Jones, Anthony E. Pickering, and Rosalyn J. Moran. “Locus Coeruleus tracking of prediction errors optimises cognitive flexibility: An Active Inference model”. In: *PLOS Computational Biology* 15.1 (Jan. 2019), pp. 1–24. DOI: 10.1371/journal.pcbi.1006267.
- [30] Peter Vincent, Thomas Parr, David Benrimoh, and Karl J Friston. “With an eye on uncertainty: Modelling pupillary responses to environmental volatility”. In: *PLOS Computational Biology* 15.7 (2019), pp. 1–22. DOI: 10.1371/journal.pcbi.1007126.
- [31] Karl Friston, Thomas Parr, and Bert de Vries. “The graphical brain: Belief propagation and active inference”. In: *Network Neuroscience* 1.4 (2017), pp. 381–414. DOI: 10.1162/NETN_a_00018.
- [32] Thomas Parr and Karl Friston. “The Discrete and Continuous Brain: From Decisions to Movement—And Back Again”. In: *Neural Computation* 30.9 (2018), pp. 2319–2347. DOI: 10.1162/neco_a_01102.
- [33] Jelle Bruineberg, Erik Rietveld, Thomas Parr, Leendert van Maanen, and Karl Friston. “Free-energy minimization in joint agent-environment systems: A niche construction perspective”. In: *Journal of Theoretical Biology* 455 (2018), pp. 161–178. DOI: 10.1016/j.jtbi.2018.07.002.
- [34] Axel Constant, Maxwell J.D. Ramstead, Samuel P.L. Veissière, John O. Campbell, and Karl Friston. “A variational approach to niche construction”. In: *Journal of The Royal Society Interface* 15.141 (2018), p. 20170685. DOI: 10.1098/rsif.2017.0685.
- [35] Raphael Kaplan and Karl Friston. “Planning and navigation as active inference”. In: *Biological Cybernetics* 112.4 (2018), pp. 323–343. DOI: 10.1007/s00422-018-0753-2.
- [36] Maell Cullen, Ben Davey, Karl Friston, and Rosalyn J. Moran. “Active Inference in OpenAI Gym: A Paradigm for Computational Investigations Into Psychiatric Illness”. In: *Biological Psychiatry: Cognitive Neuroscience and Neuroimaging* 3.9 (2018), pp. 809–818. DOI: 10.1016/j.bpsc.2018.06.010.
- [37] Alexander Tschantz, Manuel Baltieri, Anil K. Seth, and Christopher L. Buckley. “Scaling Active Inference”. In: *2020 International Joint Conference on Neural Networks (IJCNN)*. 2020, pp. 1–8. DOI: 10.1109/IJCNN48605.2020.9207382.