5TH APPLIED ACTIVE INFERENCE SYMPOSIUM

ABSTRACT BOOK

NOVEMBER 12-14, 2025

SYMPOSIUM.ACTIVEINFERENCE.INSTITUTE

Application | Education | Open Science

ACTIVE INFERENCE INSTITUTE (AII) IS AN OPEN-SCIENCE INSTITUTE DEDICATED TO IMPROVING THE ACCESSIBILITY, RIGOR, AND APPLICABILITY OF THE ACTIVE INFERENCE FRAMEWORK.

THERE ARE MANY WAYS TO GET INVOLVED WITH THE INSTITUTE AND LEARN ABOUT ACTIVE INFERENCE:

Active Inference Institute

https://www.activeinference.institute/

5th Applied Active Inference Symposium

November 12-14, 2025

Symposium program

Part 1

Part 2

Part 3

Book Editors:

Maria Luiza Iennaco & Daniel Friedman

DOI: 10.5281/zenodo.17555267

5TH APPLIED ACTIVE INFERENCE SYMPOSIUM 2025

Abstracts

Keynote speakers

Pre-recorded presentations

Live-streamed pannels

Interactive Workshops

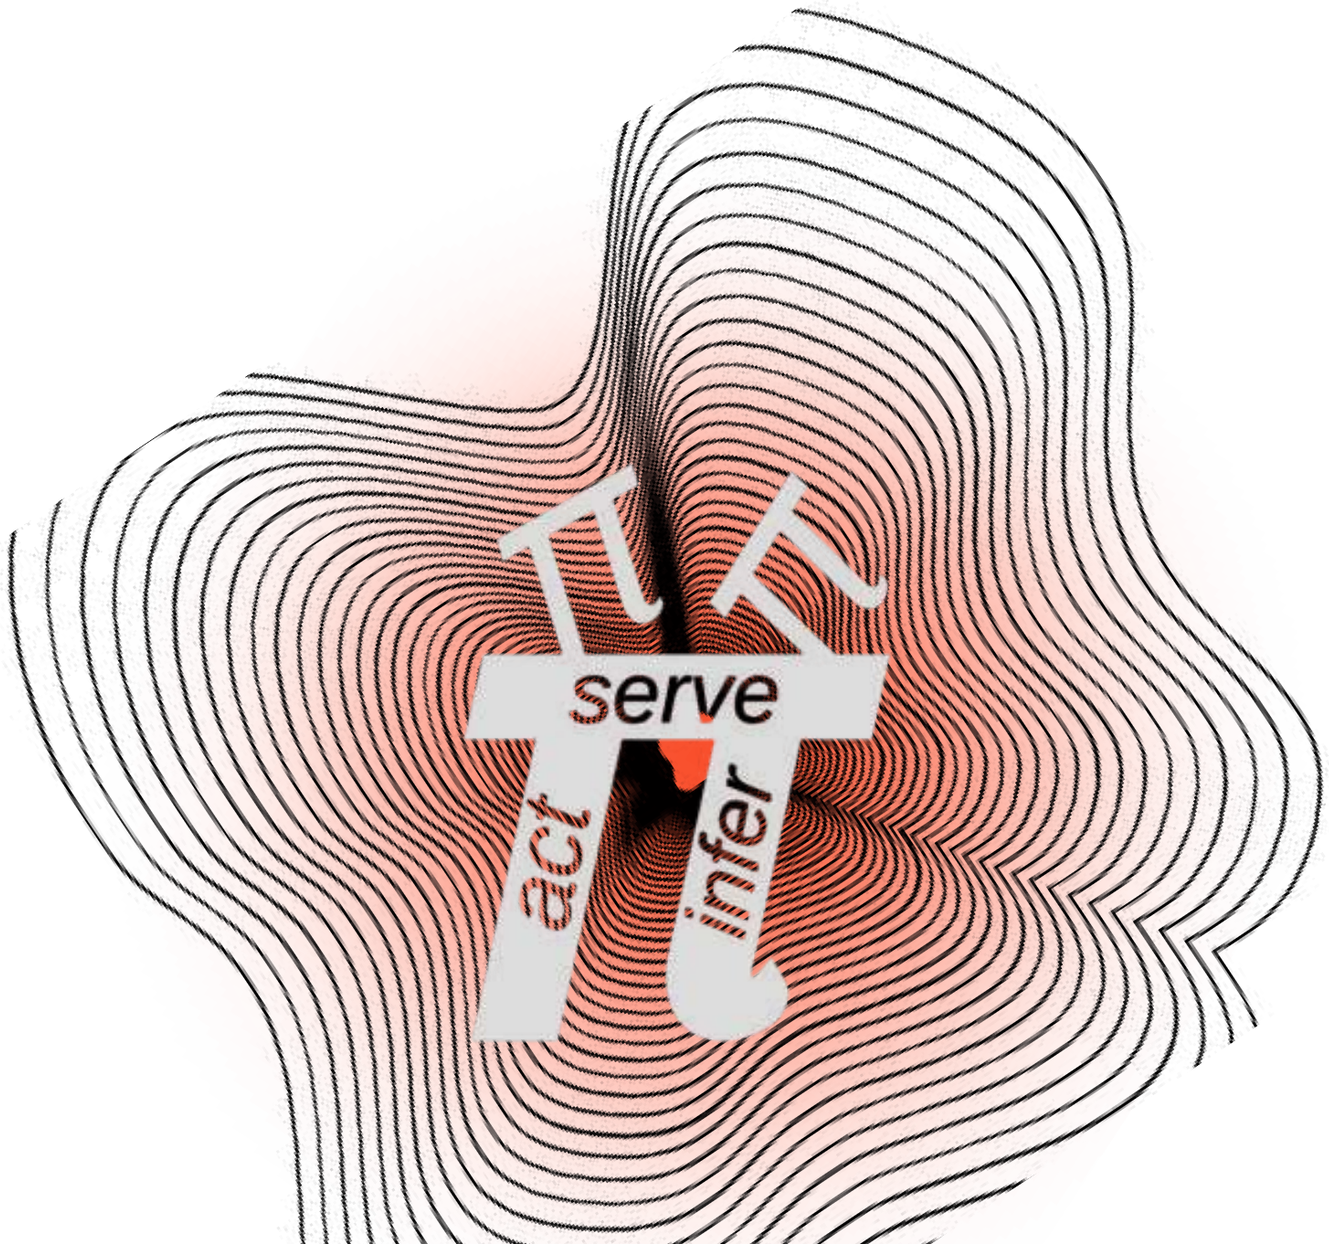

5TH APPLIED ACTIVE INFERENCE SYMPOSIUM 2025

KEYNOTE SPEAKER

Mathematical foundations of active inference

Karl Friston
University College London

Keynote address

The AI Capabilities and Alignment Consensus Project (AICACP)

Adam Safron
Allen Discovery Center, Tufts University

The AI Capabilities and Alignment Consensus Project (AICACP) is a multi-year initiative to reshape the discourse on AI capabilities, alignment, and regulation. It combines academic publishing with in-person workshops to bridge divides between different perspectives on AI. The project also aims to support public engagement through a podcast series, and potentially an online forum. At the very least, we hope to provide 2-3 interesting special issues on deep and timely topics in AI and cognitive science, and to share these insights with both domain experts and the general public. Most ambitiously, we hope to start a novel academic society dedicated to using the highest quality practices for collaborative sensemaking to figure out what world we are in with respect to increasingly advanced AI technologies, what this means for science and society more generally, and what we ought to do about it.

Coherence, Rupture, Regeneration

Alexander Sabine
University of Portsmouth

The Coherence-Rupture-Regeneration (CRR) offers an early-stage mathematical framework to explore active inference principles by incorporating non-Markovian temporal dynamics through coherence-based memory integration. CRR posits that adaptive systems maintain identity-through-change via three coupled operators: coherence accumulation $C(x) = \int L(x,\tau)\, d\tau$, discrete rupture events $\delta(t-t_o)$, and regenerative reconstruction $R\chi = \int \varphi(x,\tau)\cdot e^{(C(x)/\Omega)}\cdot\Theta(t-\tau)\, d\tau$.

A central finding that connects CRR to the Free Energy Principle is that coherence appears to increase as free energy variational bound decreases, suggesting that coherence accumulation serves as a functional proxy for uncertainty reduction. The relationship emerges naturally from a generalised Euler-Lagrange formulation, where the standard variational derivative is augmented by exponentially-weighted memory integrals and impulsive rupture terms that reset boundary conditions for variational problems.

Unlike standard Markovian models, CRR systems construct their own temporal structure through coherence fields $L(x,\dot{x},t)$, which can be positive (memory-building) or negative (decoherent). Toy simulations of CRR show how dynamically transition between Markovian and non-Markovian regimes based on coherence density, with rupture events selectively reorganising memory from the field, rather than simply degrading it.

This enables metabolised discontinuity; ruptures become opportunities for adaptive reorganisation, rather than catastrophic failure.

CRR offers potential as a predictive and diagnostic tool for understanding systems from neural dynamics to ecological transitions to cultural evolution, using a single variational calculus of adaptive temporality. The framework suggests that active inference agents are fundamentally history-bearing systems whose temporal structure is actively constructed through coherence-rupture-regeneration cycles, providing a possible mathematical language for identity maintenance over time.

Constellatory Cognition: How the brain can be modelled and use modelling without having a model

Andrea Hiott
Universität Heidelberg

Recent developments in the study of the hippocampal formation call old ideas of representation into question and are forcing a change in the way we understand the study of memory and navigation, opening the path towards a radical interpretation of the cognitive map. Through this lens, we can now assess the body's cognitive abilities (such as thinking and remembering) through the same process by which we assess the body's abilities of locomotion (such as navigating and wayfinding). In so doing we benefit from already established ideas in radical embodiment to move beyond traditional dichotomies of mental and physical. This gives us an understanding of representations not as 'findable' or 'locatable' in any hierarchy in the natural world. Rather, the scaling happens in our assessments: Representations are the ways we communicate those assessments to one another, and to ourselves; they are real, and they are interactional.

5TH APPLIED ACTIVE INFERENCE SYMPOSIUM 2025

PRE-RECORDED PRESENTATION

Making Sense of Uncertainty: Designing Experiences for Adaptive Inference

Anna Pereira

Active Inference Institute & Act In Cycle

The theories of active inference and the free energy principle offer powerful frameworks for understanding perception, learning, and adaptation, but they remain largely inaccessible outside of specialized academic communities. This paper introduces a design-led approach to translating these foundational ideas into experiential tools that can support adaptive behavior in everyday life. Drawing on interdisciplinary research, algorithmic metaphors, and visual communication strategies, we propose a simple, flexible framework grounded in three interdependent domains: mind, body, and environment. Within this context, a four-quadrant model is introduced to engage individuals across varying levels of cognitive and physical activity, encouraging personalized exploration, psychological flexibility, and a more coherent engagement with uncertainty. The approach is supported by analogies to learning algorithms, including ensemble methods, denoising, and curriculum learning, which serve as conceptual bridges to guide understanding. While early-stage and intentionally simplified, these strategies are intended to evolve iteratively through real-world application, feedback, and collaborative co-construction. This work aims to bridge the gap between theory and practice by making the principles of active inference both comprehensible and actionable by inviting individuals, researchers, and institutions to explore new possibilities for adaptation, alignment, and navigating uncertainty.

5TH APPLIED ACTIVE INFERENCE SYMPOSIUM 2025

PRE-RECORDED PRESENTATION

Approaches to Spatial Navigation: Embodied and Cybernetic Underpinnings

Bradly Alicea

Orthogonal Research and Education Lab; OpenWorm Foundation; University of Illinois U-C

Spatial navigation in animals is typically thought of as a learning, memory, and representational phenomenon based on a process of consolidating episodic memory. In mammals, spatially explicit behaviors have been tied to place [1] and grid [2] cells in the entorhinal cortex, which provide a geometric representation (abstract map) of allocentric space. In animals more broadly, path integration provides a more implicit representational mechanism for tracking an organism's position [3]. Insects provide an example of two such mechanisms: using landmarks to calibrate an internal compass, or by tracking self-motion through cells of the central complex [4]. Robotic systems use an approach called Generalized Simultaneous Localization and Mapping (G-SLAM), which utilizes perception and action to build a representation of the world [5]. Here, we will pull all of these approaches together in terms of their implications for cybernetics, embodiment, and Active Inference. G-SLAM provides us with a global representation constructed from a series of subgraphs [6]. The expansion of spatial representations occurs through interaction, thus uniting the mechanisms of entorhinal cell-based cognitive maps and path and positional integration. Specifically, sensorimotor integration serves to build and enrich the representation by optimizing and expanding the graph representing space.

This type of representation also leads us to structures consistent with cybernetic regulatory relations and partially observed Markov models. In a neural context, these types of structures allow for the same type of information processing on a wide variety of anatomical substrates.

References:

[1] O'Keefe, John (1978). The Hippocampus as a Cognitive Map. Clarendon Press.
[2] Moser, E.I., Roudi, Y., Witter, M.P., Kentros, C., Bonhoffer, T., and Moser, M-B. (2014). Grid cells and cortical representation. Nature Reviews Neuroscience, 15, 466–481.
[3] Biegler, R. (2000). Possible uses of path integration in animal navigation. Animal Learning and Behavior, 28, 257–277.
[4] Heinze, S., Narendra, A., and Cheung, A. (2019). Principles of insect path integration. Current Biology, 28(17), R1043–R1058.
[5] Al-Tawil, B., Hempel, T., Abdelrahman, A., and Al-Hamadi, A. (2024). A review of visual SLAM for robotics: evolution, properties, and future applications. Frontiers in Robotics and AI, 11, 1347985.
[6] Safron, A., Çatal, Q., and Verbelen, T. (2022). Generalized Simultaneous Localization and Mapping (G-SLAM) as a unification framework for natural and artificial intelligences: towards

Open-source Sustainability: modeling the social aspects of production communities

Bradly Alicea
Orthogonal Research and Education Lab; OpenWorm Foundation; University of Illinois U-C

In this talk, we will review techniques and approaches to studying open-source production communities. Production communities are the ways in which collective behavior is enabled and leveraged in open-source communities. Of particular interest is the origins of self-sustaining community production using a Github-like (social coding and versioned-controlled) platform. Over the past several years, we have used three different approaches: Partially-observable Markov Decision Processes (POMDP), Reinforcement Learning (RL), Agent-based Modeling (ABM), and Large Language Models (LLMs).To better characterize open-source sustainability, we must characterize immediate social complexity and long-term social evolution. Immediate social complexity can be approximated using POMDPs and RL. The long-term social evolution of social complexity in projects is modeled in the language of RL and ABMs. Additional aspects of software production, such as communication within teams, is approximated using LLM-generated project tasks and complex network theory. With this collection of tools, we can think of open-source productivity as a complex system, which may help us answer questions related to collective behavioral dynamics and the emergence of community innovation.

References:

Alicea, B., Ather, H., Chougule, H., McCorkle, B., and Parent, J. (2023). Open-source Community Sustainability using Agent-based Models. ResearchGate. https://www.researchgate.net/publication/369143414_Open-source_Community_Sustainability_using_Agent-based_Models

Combining Hierarchical Active Inference with Affordance Theory for Scalable Policy Selection in Autonomous Drone Navigation

Harshil Shah & Satyaki Maitra
Mission San Jose High School, Active Inference Institute

This presentation introduces a hierarchical active autonomous drone navigation project developed within the Microsoft AirSim simulation environment. The project applies hierarchical active inference and affordance theory to enable adaptive, efficient, and interpretable navigation in complex 3D spaces.
It will outline the motivation for building a framework that can operate under uncertainty, covering key challenges such as defining action spaces, representing abstract environmental information, and maintaining computational efficiency for edge computing.
Then, there will be a comprehensive breakdown of our architecture. The presentation will describe how the system fuses LiDAR, camera, and IMU data to infer a latent suitability state that guides decision-making by filtering infeasible actions and balancing exploration and exploitation through Expected Free Energy minimization. It will also include details on implementation, such as code snippets and visualizations.

5TH APPLIED ACTIVE INFERENCE SYMPOSIUM 2025

PRE-RECORDED PRESENTATION

Finally, the presentation will cover the strong preliminary results in our core environment and analyze the strengths and weaknesses of our architecture. It will conclude by discussing ongoing work towards more challenging environments, hardware deployment, and more advanced directions for the architecture.
The website, Over the last few years, many educational and software resources have been developed for Active Inference. In this interactive session, open source educational and development resources will be reviewed and utilized. The focus will be on presenting (more than) enough entry points, for manual and AI-augmented learning + coding. Questions, frictions, memes, dreams, and other input will be curated for continued ongoing improvement. has the most up-to-date details on the project and will continuously be updated.

5TH APPLIED ACTIVE INFERENCE SYMPOSIUM 2025

PRE-RECORDED PRESENTATION

Modeling the Emergence of Perspective: An Active Inference Account of Intentionality

Hongju Pae

Research Fellow, Active Inference Institute

This presentation applies Active Inference to a foundational question in consciousness science: how subjective perspective (subjective intentionality) emerges as a self-organizing process within living systems.

Building on the analysis developed in Reflective Analysis on Empirical Theories in Consciousness (Pae, 2025), I argue that what may unify multiple mainstream theories of consciousness is their shared effort to explain the intentional quality of experience - the fact that consciousness is always from a particular point of view.

Within this framework, Active Inference offers a principled way to model such perspectival structure by treating intentionality as a latent informational process that modulates qualitative aspects of experience under continuous minimization of uncertainty about self-world relations. In doing so, it formalizes the emergence of subjectivity as an intrinsic feature of autopoietic inference. The presentation thus applies Active Inference to bridge phenomenology and computational modeling, illustrating how intentional stance and perspective formation can arise within probabilistic generative architectures.

5TH APPLIED ACTIVE INFERENCE SYMPOSIUM 2025
PRE-RECORDED PRESENTATION

The phenomeno-logical robot

Jean-François Cloutier
Research Fellow, Active Inference Institute

I present an update of the Symbolic Cognitive Robotics project I conduct as an AII Research Fellow.
I am programming a mortal, autonomous robot (a rover with motors and sensors) so it has felt experiences, makes sense of them, and evolves its agency to survive in a dissipative world. My project is a learning-by-doing exploration of phenomenology, active inference and artificial agency.
The robot will actively evolve a society of mind composed of simple cognition actors (CAs) organized in an abstraction hierarchy. A CA observes other, less abstract CAs (its umwelt), updates its beliefs from patterns it detects, decides which beliefs feel good, and acts on its umwelt to feel better. A CA uses a logic program it generates and updates to predict incoming observations and to anticipate the consequences of its actions.
I am hoping this work will help me answer difficult questions such as:
Can a robot be made mortal so it can have good vs bad experiences?
Can a robot's experiences and sense making be entirely its own and not those of its programmer?
Can there be something it is like to be a robot?

The missing center: Somatic belief updating in the global model.

Jim Freda
Independent Scholar/Practitioner

The core node has been undertheorized. The model must itself contain an internal model. This internal model has an outer boundary sheet, a Markov blanket of its own. I explore this somatically in the muscle sense.
In Sherrington's "neural architecture of the animal as a whole" we understand that there are three distinct spatial distributions in the model. We are well acquainted with these three perceptual systems, except for proprioception. The motoric element of this deep field is missing from the model. This is a significant barrier to recognition and action.
Overwhelming surprisal reflects a profound bias in the collective model, a control bias resulting in a dangerous feedback loop. I discuss this as exteroceptive bias. There has been a kind of refusal or repression that we can understand as a double bind in the blanket politics of the model. Because it contains not just one but two blankets, necessary for the model to be generative, then we must theorize their relationship.
It is the motoric element of the deep field in the model, through muscle sense, that actively generates the self evidencing needed for effective belief updating in the face of overwhelming surprisal.

This motoric element is the outer surface of the inner self, of the core node in the model. It both expresses and defends its neural content and has an important integrative biomechanical role.
We know this axial system as a series of pits or concavities. These are major structural junctures, macro-level sensory motor synapses, each with a single prominent bony servomechanism at its center. This is a self directed system (of systems) within us and I describe it’s synergistic kinesiology, allowing us to track its movements. Please be prepared for guided movement meditations.

The ambiguous status of attention in predictive models of cognition

Maria Luiza Iennaco
University of São Paulo / University of Porto

Predictive models of cognition, which frame the brain as a predictive machine, have become prominent in cognitive science. Despite their promise, one aspect remains under-theorized: attention. These models attempt to describe (or relate) attention as the optimization of precision weighting, i.e., the confidence assigned to sensory signals, but often treat it as a byproduct of perception and action.
This presentation offers a critical assessment of the potential of Hierarchical Predictive Coding and Active Inference models to encompass the different domains of attention – a crucial boundary phenomenon whose role is central to their entire functioning. To this end, we will start with a critical survey of the contemporary status of these models, followed by a more in-depth investigation of the extent to which they are capable of describing the particularities of attention.
Following a comparative examination, we discuss the models' potential and limitations in theoretical and clinical applications. While they offer appealing philosophical narratives and reasonable reinterpretations of psychiatric disorders, we believe they currently lack innovative, testable hypotheses. Finally, we examine important theoretical restrictions and potential future research areas for better integrating attention into predictive processing frameworks.

5TH APPLIED ACTIVE INFERENCE SYMPOSIUM 2025

PRE-RECORDED PRESENTATION

Hyperscanning under Active Inference for Mental Health

Nicolás Hinrichs

Okinawa Institute of Science and Technology & Max Planck Institute for Human Cognitive and Brain Sciences

Mental-health interventions thrive on attunement between therapist and client, yet quantitative tools for monitoring that attunement are scarce.
We present i) a real-time framework that extends active inference to dyadic interaction and links it to a geometric hyperscanning observable suitable for clinical deployment. Each partner in therapy is modelled as an agent whose self-model contains a generative model of the other. Prediction errors in this recursive model manifest in the topology of their combined neural activity.
We show ii) how to compute Forman-Ricci curvature on the weighted inter-brain network and track the entropy of its distribution.
We introduce iii) a prospective Digital Twin for Psychotherapy incorporates real-time monitoring of sharp entropy fluctuations, which timestamp-and ultimately, predict-phase transitions of crucial in-session behavioral phenomena, such as rupture, co-regulation and re-attunement, to guide the practitioners towards the optimal outcome: healing.

5TH APPLIED ACTIVE INFERENCE SYMPOSIUM 2025
PRE-RECORDED PRESENTATION

Observer Theory

Sam A Senchal
Wolfram Institute

This presentation examines a novel category-theoretic extension of Observer Theory that formalizes how observers construct reality through constrained sampling of Wolfram's Ruliad—the ∞-groupoid of all computational processes. The framework introduces observers as functors that reduce entropy while respecting boundedness, persistence, and relevance constraints, directly connecting to Friston's Free Energy Principle.

The proposed hierarchical domain structure—Physical (P), Valuational (V), Symbolic (S), and Minimally Constrained (M)—with embedding functors provides a mathematically rigorous approach to cross-domain causation and information integration. Qualia emerge as integrated information across domains, offering a testable formalization of consciousness compatible with Integrated Information Theory.

Key technical contributions include:

- Resolution of discrete-continuous duality through observer-dependent measure spaces
- Formal treatment of cross-domain causation without violating physical causal closure
- Information-theoretic formulation to quantify observer-accessible novelty

Terminal object (True Infinity) preventing infinite regress in the observer hierarchy.

5TH APPLIED ACTIVE INFERENCE SYMPOSIUM 2025

PRE-RECORDED PRESENTATION

Panelists will debate whether this synthesis successfully bridges Active Inference with fundamental physics, the necessity of the four-domain ontology, and implications for designing artificial observers. Critical discussion points include the framework's empirical testability, its relationship to existing theories (IIT, FEP, Constructor Theory), and whether True Infinity represents mathematical necessity or philosophical overreach.

This interdisciplinary dialogue promises to advance our understanding of observers as active constructors of reality, with profound implications for physics, consciousness studies, and artificial intelligence.

5TH APPLIED ACTIVE INFERENCE SYMPOSIUM 2025
PRE-RECORDED PRESENTATION

From Information Economics to Active Inference: A Unified Bayesian Mechanics of Choice Under Uncertainty

Samuel Montañez
Artificial Intelligence PhD student at Universidad Panamericana (Mexico City)

This paper establishes mathematical equivalences between classical economic decision theory and active inference, demonstrating that economic behavior emerges from free energy minimization principles. We show that foundational models in economics—including Stigler's (1961) optimal search, Simon's (1955) bounded rationality, Arrow's (1962) value of information, and Sims' (2003) rational inattention—constitute special cases of active inference under specific conditions.
Four key theoretical contributions emerge. First, Stigler's optimal stopping rules correspond to active inference when epistemic affordance (expected information gain) is excluded, leaving only pragmatic affordance (expected value). This reveals why empirical search behavior often exceeds Stigler's predictions: agents intrinsically value uncertainty reduction beyond its instrumental benefits. Second, Simon's satisficing emerges naturally when computational costs are incorporated into the free energy functional, with aspiration levels representing equilibrium points where the marginal cost of deliberation equals marginal expected benefit.

Third, Arrow's value of information corresponds precisely to the Kullback-Leibler divergence between posterior and prior beliefs, providing the information-theoretic foundation previously absent from economic theory. Fourth, Sims' rational inattention framework, which constrains information processing through Shannon channel capacity, generates the same softmax choice structure that emerges from variational approximation in active inference—revealing that economists independently discovered free energy minimization through information-theoretic reasoning.

This unification has profound implications for behavioral economics. Documented anomalies—including the Allais paradox, framing effects, and probability weighting—emerge as optimal solutions given the brain's hierarchical predictive architecture and precision-weighting mechanisms. Rather than representing irrational deviations, these behaviors reflect Bayes-optimal inference under neural computational constraints. The framework generates testable predictions linking economic choices to precision-weighted neural responses in frontal-striatal circuits. By recognizing that economic agents minimize variational free energy through coupled pragmatic and epistemic drives, this synthesis bridges a century of economic theory with contemporary neuroscience, artificial intelligence, and the physics of complex self-organizing systems.

Fundamentals of active inference: A self-guided textbook for learning and applying active inference from first principles

Sanjeev Namjoshi
VERSES

Over the past decade, research in active inference has expanded dramatically, generating a wide array of new ideas across diverse subfields and application domains. Yet, the discipline remains in its formative stages, despite its potential to transform research and practice across science and engineering. With the growing excitement surrounding generative AI and reinforcement learning, active inference is now uniquely poised to attract attention from students and researchers in multiple disciplines, setting the stage for a wave of innovation in both academia and industry reminiscent of the deep learning surge of the early 2010s. However, despite its remarkable promise, active inference has yet to achieve the broad adoption it merits largely due to the steep learning curve and the lack of accessible, structured introductions in the literature. In this talk, I present Fundamentals of Active Inference, a textbook designed to bridge that gap by providing a unified, systematic, and rigorous foundation for students and researchers entering the field. The textbook is entirely self-contained and suitable for self-study, assuming no background outside the basics of probability theory, linear algebra, and multivariate calculus. Following the release of the textbook many supplemental educational resources will be made available including Jupyter notebooks, simulations, video lectures, and interactive software intended to teach and educate the concepts of active inference in order to bring these ideas to a wider audience and induce a paradigm shift leading to the next phase of both applied and theoretical AI research.

5TH APPLIED ACTIVE INFERENCE SYMPOSIUM 2025

PRE-RECORDED PRESENTATION

Active Self-Referecing: The concept of self-reference might be that which joins complex self-modeling with reference problems in formal languages across a spectrum

Sonia de Jager

Erasmus School of Philosophy

Referents can be understood as (shared) models encoding possibilities. All (linguistic, logical, etc.) referents are inherently ambiguous, and require context for concretization, else they remain unactualized plans/imaginations/memories. Simple referents with relatively “easily” concretizable possibility spaces (“cat”) are less computationally challenging than complex ones (“negation,” or: the concept of reference itself). We examine the predictive challenge agents encounter in self-referential “loops”: when reference “models” reference itself. What predictive function does self-reference serve? Is it a foundational feature of self-modeling, or a side-effect hereof?

The proposal we begin from is that basic self-reference is a prerequisite for selfhood, where selfhood emerges through the surveilling and synthesizing of one’s own information-processing. The loop hereby created is sustained by agents first modelling themselves as agents to effectively act upon their environment. Maintaining this recursive process sustains the experience of continuity through memory and planning (which are not symmetrical, but intertwined).

We conceptualize this continuity as predictive rehearsal: the remembering-planning and recursive updating of simulated experience. For all cognitive systems, lived experience necessarily involves the continuous rehearsal of self-predictions, but also the permanent updating hereof.
We propose that a self-model which is capable of reference (has linguistic/communicative capacities) updates its predictive self-referential tension with itself, and this effect is mirrored in the larger communicative landscape through other processes of self-reference. What is also set under speculative investigation here is the concept of self-reference as haecceity: most linguistic/communicative phenomena are representable as referents within shared generative models, but reference itself—the substrate enabling linguistic modeling—cannot be modeled as a referent within its own system, much like the self remains (computationally) intransparent to itself, and can therefore be understood as the ultimate representation of thisness or individuation. This has implications for understanding metacognitive prediction limits and the emergence of perspective/selfhood in active-inferential agents.

5TH APPLIED ACTIVE INFERENCE SYMPOSIUM 2025

PRE-RECORDED PRESENTATION

Tradescaping for Active Inference Action Research

Susan Hasty

Independent

TbD

ALife simulations for multi-agent alignment of active inference agents

Adam Safron
Allen Discovery Center, Tufts University

We will describe a variety of research questions that we are preparing to pursue with an integrative simulation environment. We will implement meta-learning world-modeling agents wherein intelligence and values co-develop through unified prediction-error minimization across hierarchical temporal scales. We will deploy computational architectures based on hierarchical generative models that perform reciprocal message passing similar to mammalian nervous systems, enabling flexible construction of self-world representations with varying temporal depth. Intrinsic drives for curiosity and empowerment guide exploration, and meta-learning allows rapid adaptation to new contexts while preserving established value orientations through precision-weighted attention to cultural and social cues. Unlike current approaches that treat capabilities and alignment as separate problems, our approach embeds value-formation within the fundamental learning process. We will demonstrate how quasi-Kantian ethics can be made to emerge naturally from temporally deep policy selection and generative modeling in ways that require coordinating with others. This research will involve interdisciplinary collaborations between individuals with expertise in the psychology of personality (broadly understood as system-defining attracting states), psychiatry, multilevel selection and evolutionary game theory, and even basal cognition (to investigate the generality of these principles).

Entropic motivation

Alex Kiefer
Monash Center for Consciousness and Contemplative Studies; VERSES

A key feature distinguishing active inference from more mainstream AI paradigms like reinforcement learning is the inclusion, in a principled way/at the ground level, of dimensions of motivation related to exploration and curiosity. How deep do such “entropic” principles of motivation go? Are they always in the service of utility or preference-seeking behavior? I argue that entropy-maximizing behavior is a sui generis dimension of motivation orthogonal to the pursuit of reward, and in fact prior to it in the sense that unlike utility, it is model-agnostic.

Hierarchical Active Inference Modeling of Social Trust Dynamics in PTSD: Integrating Qualitative Phenomenology with Computational Psychiatry

Andrew Pashea
Active Inference Institute, University of Chicago Harris School of Public Policy Applied Data Fellow

This project explores the dynamics of trust and decision-making in behavioral task agent-based modeling for individuals with PTSD, integrating insights from qualitative clinical research, fear extinction and avoidance behavioral research, and phenomenology with the Active Inference framework. PTSD is known for involving difficulties with emotional and physiological regulation in relation to the aftermath of trauma. This can often manifest behaviorally and cognitively as, e.g., distrustful or avoidant behaviors, dissociation or interoceptive disconnect, potential burnout. While previous research has focused on ethological models in approach-avoid or explore-exploit tasks, our work emphasizes the human element of trust in everyday life. We employ the Card Advisor task used in previous work used to model “delusion” broadly, in the context of schizophrenia, at the same time aligning our priors more closely to literature in PTSD and adding a metacognitive layer to the model for reflective inference about safety. This choice reflects a translational perspective on psychiatry where the modeling process can help inform or complement our understanding in cases of ambiguity when relying upon traditional definitions of psychopathological disorders.

5TH APPLIED ACTIVE INFERENCE SYMPOSIUM 2025

LIVE-STREAMED PANEL

Our work aims to further interdisciplinary work and perspectives for aiding clinical practice and understandings of PTSD. As this project is still in preliminary stages, with its beginning at the Active Inference Institute, the primary presenter will include reflections on the progression of the project design from a few perspectives, general modeling flow from a programming perspective, to foster further conversation and interdisciplinary collaborations.

5TH APPLIED ACTIVE INFERENCE SYMPOSIUM 2025

LIVE-STREAMED PANEL

An invitation to Active Inference for the Social Sciences

Avel GUÉNIN—CARLUT

University of Sussex / Kairos Research / Active Inference Institute

Active Inference as a model for human cognition has been productively applied to social cognition for the best part of 10 years, and has more recently been related to open research questions in anthropology sociology, linguistics, archeology, political science, and social ontology. In particular, various researchers have analyzed the consequences of Active Inference to the following topics: the construction of cultural affordances and social learning ; the existence and ontological status of collective and imagined entities ; the nature of social/material organization in relation to cognition ; and the role of ritual in political power and social organization.

This literature offers a rich, powerful, and yet elegant approach to integrate the study of social behavior in the many relevant scales of organization - from the biological underworld of Levin's "agentive materials", to the overworld of structural constraints on State organization. It vindicates and specifies insights from well-established branches of the social sciences, such as Goffman's approach to symbolic interactionism, Lakoff's embodied socio-linguistics, and more generally the pragmatist tradition pioneered by Dewey, James, Peirce, and Mead (among others) at the turn of the XXth century.

5TH APPLIED ACTIVE INFERENCE SYMPOSIUM 2025

LIVE-STREAMED PANEL

Yet, a problem remains for the Active Inference theorist. This approach exists at the interface of many different disciplines, blending of course cognitive and social sciences, but also blurring the traditional boundaries between sociology and archeology, design and political theory, anthropology and semiotics, physics and philosophy. It makes it difficult to pin down to specific fields of study, and therefore difficult to engage with in the highly siloed landscape of the academic world. A systematic effort is yet to be undertaken for the new world to be born, and this livestream intends precisely to lay out the ground work for new people to join the effort of building Active Inference for the Social Science.

ActInf in biology: modeling the environment as an agent

Chris Fields
Tufts University

The symmetry of the FEP makes the environment of any active inference agent an active inference agent. The environment's actions on the agent "of interest" are driven by the environment's GM, which describes what the environment believes about the agent. Every action of the environment on the agent can be considered a communication from the environment formulated by the environment's GM. How does this symmetry change the way we think about and build models in biology, from the molecular scale to the social or ecological scale? What is the environment "telling" a polypeptide as it folds into an active protein? What is the environment telling a differentiating cell during morphogenesis? What is your environment telling you? What can we infer, as modelers, about the environment's GM from the way the environment acts on the agent of interest?

From Coordination to Cognition: Rethinking Agent Communication Protocols for Active Inference in the Spatial Web

Denise Holt
AIX Global Media

As the global demand for autonomous systems accelerates, a new class of Agent Communication Protocols (MCP, ACP, A2A, ANP) is emerging to facilitate interoperability among distributed AI agents. Yet, these protocols are rooted in architectures that favor reactive task execution, rather than adaptive, inference-driven behavior. This presentation explores how these current protocols, though solving real enterprise needs, are structurally misaligned with the causal reasoning, real-time adaptation, and generative modeling required by Active Inference Agents.
I will evaluate each protocol's strengths and limitations in supporting inference-based decision-making, and contrast their architectural assumptions with the scale-free, probabilistic modeling required for Active Inference. More critically, this talk will introduce how the Spatial Web Protocol (HSTP/HSML) and the Universal Domain Graph (UDG) offer a dynamic, context-rich substrate that complements the Free Energy Principle and enables distributed Active Inference at scale.

5TH APPLIED ACTIVE INFERENCE SYMPOSIUM 2025

LIVE-STREAMED PANEL

This session invites the Active Inference research community to reimagine not just how agents coordinate—but how they learn, infer, adapt, and act together in an environment built for inference, not just information exchange. As we shape the future of agent architectures, it's time to go beyond static message-passing and toward embodied, semantically grounded communication protocols that align with the principles of Active Inference.

An Applied Active Inference Group Practice for Transdisciplinary Collaboration

Ian Tennant
Anglia Ruskin University

Transdisciplinary projects often falter not because of technical gaps, but because of difficulties in communication, alignment, and meaning-making across mixed domains. This session introduces a structured group practice, inspired by therapeutic traditions including Circling which has recently been developed by Guy Sengstock and John Vervaeke. Such group work encourages present-moment awareness, perspective-taking, and reflective listening. Participants explore the “here-and-now” of interaction through paraphrasing, impact-checking, and attunement, with the goal of improving shared understanding and collective insight.
We propose that these group practices can be understood as a form of applied Active Inference. With...
Generative model alignment - Participants continually update their beliefs about others’ intentions and perspectives.
Precision calibration (relevance realization) - The practice highlights which signals deserve confidence and which need revisiting.
Epistemic action - Asking clarifying questions and mirroring are ways of actively sampling to reduce uncertainty.
Shared anticipations - Over time, the group develops coordinated protention and a collective salience landscape.

5TH APPLIED ACTIVE INFERENCE SYMPOSIUM 2025

LIVE-STREAMED PANEL

In this session, participants will learn the basics of this group practice, experience short facilitated practices, and reflect on how such methods can support transdisciplinary teams in industry. By framing Circling-like group practice as a live process of Active Inference, the workshop offers both a theoretical lens and a practical tool for improving genuine collaboration across disciplinary boundaries.

Story Graphs: An Exploration of Use Cases In and Beyond CogNarr

John Boik
Active Inference Institute, Verses.AI

Story graph is a formalism under development at the Institute's CogNarr (Cognitive Narrative) Project that captures and makes transparent and explicit a storyteller's belief model regarding some situation. A story graph might convey a person's beliefs about what happened, what might happen, what it means, why it happened, what future is desired or feared, who was involved, and so on. It is an intermediate representation between natural language and a technical implementation (e.g., code in a model) that is readable by both humans and computers. A story graph conveys information with less ambiguity than natural language. In the CogNarr setting, the sharing of story graphs helps to facilitate group cognition at scale. But story graphs could have applications well beyond CogNarr, in medicine, industry, media, governance, commerce, research, education, and civic society, for communication, decision making, forecasting, model verification and explainability, and other purposes. This panel will describe the vision behind story graphs and, with help from the audience, explore possible use cases that could benefit society.

Adaptive Robotics for the Real World: A New Paradigm for Autonomy

Matthew Brown
ThoughtForge

Robotics adoption has long been limited by brittle systems that fail when the environment changes. A new approach rooted in Active Inference enables robots to adapt on the fly and learn from their environment while maintaining stability and precision. ThoughtForge's platform integrates these principles into real-time robotic control, solving key challenges like sim-to-real transfer and variability in unstructured environments. From adaptive inspection of large-scale infrastructure to precision manipulation in manufacturing, this talk highlights real-world applications where adaptive AI transforms robotics from rigid automation into truly intelligent systems.

5TH APPLIED ACTIVE INFERENCE SYMPOSIUM 2025

LIVE-STREAMED PANEL

TBD

Michael Garfield

AICACP, Humans On The Loop

TBD

An Active Inference Agent for Asset Pricing

Samuel Montañez
Universidad Panamericana

We propose a unifying, agentic account of financial decision-making and asset pricing grounded in Active Inference. Rather than treating uncertainty as exogenous noise to be averaged away, we model investors as POMDP agents that minimize expected free energy—balancing extrinsic value (risk-adjusted returns) with epistemic value (information gain). This perspective reveals classical approaches—Black–Scholes–Merton (BSM) real options valuation and stochastic discounted cash flow (DCF)—as limiting cases that constrain belief dynamics to gradient flows, suppressing strategic information search. Using a Helmholtz decomposition, we formalize how Active Inference relaxes this constraint, enabling curvature-aware belief updates and stable nonequilibrium behavior on information-geometric manifolds. We illustrate the framework with a financial agent and a lightweight, model-free baseline, clarifying when epistemic control improves sequential allocation. Conceptually, the work links portfolio choice to Bayesian experimental design and precision (attention) control, offering testable predictions for neuroeconomics and finance while providing a reproducible path to engineering belief-driven market agents.

Money and the Free Energy Principle

Steph Macurdy

Wolfram Research, Wolfram Blockchain Labs, Quai Network, UTXO Alliance

An introduction into the topic of cryptocurrency, specifically a two-currency system called Qi and Quai, and how they relate to energy and entropy, and how that relates to the Active Inference community and the free energy principle. (15 to 30 minutes)
Next, a presentation of the math of Active Inference and the Free Energy Principle. A comparison between two interpretations of the math, the customary Bayesian interpretations and an alternate interpretation that describes a dialogue between two systems, PIx,y) and Q(x). (30 minutes)

Active Neurorobotics: Deep Prior Preference Learning Guided by Learnable Goals

Viet Dung Nguyen
Rochester Institute of Technology

Research in active inference (AIF) has led to many improvements with respect to the robustness of model-based agents that are applied to both Markov decision processes (MDPs) and partially observable MDPs (POMDPs). Active inference, essentially, entails a form of inference and learning that seeks to balance goal-orienting objectives or prior preferences – exploitative instrumental signals – with epistemic ones – foraging, exploration-driving signals. Despite these advances, few studies have investigated the complexity of prior preferences or leveraged them effectively to generate informative instrumental signals. Moreover, most AIF research has overlooked scenarios where goals are provided as observations, a setting critical for tasks that require explicit goal querying or instruction-following. In real-world robotic control tasks – particularly those with complex, dynamically changing goals, such as language-guided manipulation – standard AIF architectures are insufficient. In this talk, we will specifically analyze these limitations by introducing a novel, systematic framework for deep prior preference model learning guided by goals or instructions in any modality. Our approach frames AIF agents in terms of progressively constructing and adapting a prior preference by leveraging multimodal fusion and queried goals at each time step, facilitating the dynamic emission of a useful, dense instrumental signal based on provided instructions for neurorobotic problem settings. We further study how our prior preference adaptation scheme makes an AIF agent flexible in challenging problem contexts, particularly in those that require the querying of goals related to human instruction or desired states.

NGC-Learn V3: A Fast, Modular, Computational Neuroscience Library

William Gebhardt
Rochester Institute of Technology

As the field of computational neuroscience expands, the models being developed and studied are becoming more and more complex. As a result, the methods utilized to run and simulate these models take longer and require more resources. This holds especially true when working with large, stateful systems such as those that make up spiking neural networks. Furthermore, testing preliminary ideas and methodological ideas still incur a high cost in both time and labor to correctly set up all of the moving parts, generally resulting in highly specialized, brittle code that will need to be stripped down in order to be reused for another, later idea. In the Neural Adaptive Computing (NAC) Lab, we have designed NGC-Learn to be a library that handles key experimental concerns related to optimization and reusability. The library allows users to focus on what really matters -- the research and development of novel methods and models -- without worry as to the simulation and design overhead. Join us in this high-level tech demo of the latest version of NGC-Learn, covering its design patterns and modularity as well as the internal processes that it uses to speed up computation.

5TH APPLIED ACTIVE INFERENCE SYMPOSIUM 2025
INTERACTIVE WORKSHOP

Using Active Inference for the Management of Distributed Energy Resources in CityLearn

Kobus Esterhuysen
LearnableLoopAI.com

Active inference offers a powerful framework for managing distributed energy resources (DERs) because it integrates perception, learning, and action in uncertain environments, using probabilistic models to make adaptive decisions about resource allocation and control. This approach enables energy systems to self-optimize and dynamically respond to changing demand, generation fluctuations, and grid conditions, improving energy efficiency, reliability, and long-term resilience. By leveraging the capabilities of active inference, operators can better coordinate diverse DER assets without full system knowledge, efficiently balancing supply and demand while supporting the integration of renewables and distributed generation.

Interactive Math Tools for Active Inference Education

Octopus

Eight Arms Nine Brains

To help students reach their goals, the Teacher must be open to and welcome feedback. This session will discuss tools and strategies for eliciting feedback in various forms.

Active Inference & RPGs

PabloFM
Role Playing Games

In this session, we will take a guided walk through the emerging landscape of the 3D internet — the metaverse — and explore how principles of Active Inference and the Free Energy Principle can inspire the design of meaningful, interactive experiences. At Numen.games, we have spent the past several years experimenting with how these theoretical frameworks can be applied in practice, consciously embedding them into gamified, educational environments.

Our partnership with the Active Inference Institute allows us to share these spaces with a broader community, creating opportunities for exploration, learning, and collaboration.

Participants will not only hear about our journey but will also engage directly with interactive demonstrations inside a virtual environment. By moving beyond theory into lived, embodied experience, we aim to show how Active Inference can inform the way we build and inhabit digital spaces.

The session will be highly participatory: attendees will be invited to join us inside a 3D online environment, interact with one another, and reflect together on how the principles of Active Inference shape perception, action, and communication in real time. This interactive component ensures that the session is not just about describing ideas, but about experiencing them.

Ultimately, we invite participants to imagine with us how the metaverse can become more than entertainment: a space for education, scientific exploration, and collective sense-making — grounded in the principles of Active Inference.

Workshop on ActiveInference.jl.

Peter Thestrup Waade
Aarhus University

In this presentation, we will describe recent updates to ActiveInference.jl, as well as how it connects to other software packages in the larger active inference and Bayesian inference ecosystem, such as RxInfer.jl and Turing.jl. ActiveInference.jl has been extended to be able to accommodate other generative models than the discrete POMDP classically used in the field, such as generalized filtering, hierarchical gaussian filtering, differential equation models or predictive coding networks. The package also allows for using different types of perceptual inference methods (such as amortized inference or different variants of variational Bayes), and different variants of the actions selection process (sophisticated inference etc.). We have aimed at making the package easy to use - and easy to extend and develop for contributors, in the hope that it can help facilitate iterative development in the community. Finally, owing to Julia's native autodifferentiability, ActiveInference.jl is invertible, so that it can be used for computational phenotyping (or cognitive modelling, as the method is more broadly known), such that inference can be drawn about unknown generative models held by systems of interest, such as human participants in psychiatric experiments.

5TH APPLIED ACTIVE INFERENCE SYMPOSIUM 2025

INTERACTIVE WORKSHOP

Consciousness and Active Inference

Robert Worden

Active Inference Institute

Recent work on consciousness implies that a pure computation cannot be conscious. This is because all information in a computer is encoded, whereas consciousness is not encoded. This means that the neural computations of Active Inference are not conscious. Consciousness could arise from an analogue model of 3-D space in the brain, held as a wave in the thalamus. If it does, how do the neural computations of Active Inference couple to the wave? Possible couplings between Active Inference and consciousness are described. The workshop will discuss research projects to investigate them.

Thermodynamics

Patrick Huembeli & Maxwell Ramstead

5TH APPLIED ACTIVE INFERENCE SYMPOSIUM

ABSTRACT BOOK