

6th Applied Active Inference Symposium — Book of Abstracts (Working Program)

Contents

About this program	3
1 An in-depth discussion on “World models in natural and artificial intelligence”	4
2 Build a curious rocket: the hands-on cpmoap tutorial for the Applied Active Inference Symposium	5
3 A Modular Perspective Layer for Active Inference Agents	6
4 Active Inference and Free Energy Principle in 2026	7
5 Active Inference, Bayesian Causal Learning and Empowerment in RL.	8
6 Root Frequency Theory: An Integrative Framework for the Continuity of Lived Experience	9
7 Working Groups Sessions (3 sessions: Opening, Charter Session, Outcomes and Next Steps)	10
8 Active Inference and the Evolution of Mind and Society	11
9 From Self-Evidencing to Global Governance: Active Inference and the Dynamics of Justice	12
10 “Evolution, Emergence, Prediction.”	13
11 Recent advances in empirical applications of active inference in computational psychiatry	14
12 An elementary form of agency at planetary-scale	15
13 Reverse engineering generative models from zebrafish decision-making	16
14 A Lexical T-Maze for Active Inference: A Proof-of-Concept Model of Context Seeking Under Ambiguity	17
15 Active Inference Tutorial: Perspectives, Correlates, and Future Directions for Graphical Brains	18
16 A precis of two principles - Towards a relational perspective on the FEP	19
17 The Benefits of Circuit Convolution: behavioral complexity arises from developmentally-simulated nonlinearity.	20
18 Realms of cognitive engagement in Active Inference	21
19 NUMEN: Bare metal inference	22

20 Partner-Specific Affective Precision in Social Active Inference	23
21 Who knows	24
22 Toward Active World Modeling	25
23 Active inference and design	26
24 Active Inference, Information, Human Evolution	27
25 Congruent Interaction Geometry	28
26 Language, Self and the Shadow Audience	29
27 When Agency Stalls: The Human Readiness Gap	30
28 Where the Theories Stop: What Free Energy and Integrated Information Actually Buy You in a Running System	31
29 Entropy Minimization is Maximal Order	32
30 Dynamic Action-Space Adaptation via Contrastive Transition Adaptation in Active Inference	33
31 ROIF and Active Inference: The Formal Boundary of History-Conditioned Anticipatory Systems.	34
32 From Lived Experience to Mathematical Mapping: A Computational Phenomenology of Trauma Recovery	35

About this program

Event dates: 2026-11-12 to 2026-11-13.

This working program collection contains approved public abstracts and explicitly marks entries awaiting an abstract. Inclusion does not assert that a presentation has occurred or that a recording has been received.

1 An in-depth discussion on “World models in natural and artificial intelligence”

Presenter: Bradly Alicea

Contribution: Interactive Workshop

Submitted abstract

This session will be a summary and discussion group for the special issue “World Models in Natural and Artificial Intelligence”, published in the Royal Society A journal in May 2026. We will offer a summary of the special issue content, along with a discussion of Active Inference related topics and future directions from the work presented in the issue.

Source: [submitted source](#)

2 Build a curious rocket: the hands-on cpomdp tutorial for the Applied Active Inference Symposium

Presenter: Dan Corva

Contribution: Interactive Workshop

Submitted abstract

In the 1960s, Rudolf Kalman wrote down a filter that would go on to navigate Apollo to the moon. It has been the workhorse of continuous-state estimation ever since, and if you have ever built a continuous active inference agent in the linear-Gaussian regime, you have built one too, whether you intended to or not. Bolt a goal onto it and you get a very good controller that is not curious in the slightest. Koudahl, Kouw and de Vries proved why in 2021: under a fixed sensor the epistemic value is identical for every action, so expected free energy quietly reduces to Linear Quadratic Regulator (LQR). In this workshop, in about ten lines of Python that will look familiar to anyone who has used pymdp, we are going to build a rocket! Then, we add the interesting part, curiosity, by making a surgical, biologically realistic change: adding sensor noise that varies based on the rocket's position. We will see how the flight path bends when we add artifacts like a navigation beacon or a blackout zone, and explore the science as to why. You will write every line yourself in a Colab notebook and do not worry there will be plenty of scaffolding from the cpomdp team. Code hints will drop in on a timer, and a reference cell for anyone who gets stuck or would rather just see the finished steps. We will build this in three steps: the 1960 rocket, a nav beacon it can fix its position by, and a blackout zone. Everything we use runs on the PyPI release of cpomdp and anything you build in the notebook is yours to keep and continue working on. Prerequisites: comfortable editing a few lines of Python. Familiarity with pymdp helps but is not required. Nothing to install; a browser and a google account is enough. If you would like to use your own IDE feel free but we may not be able to provide 1-on-1 support if you get stuck.

Source: [submitted source](#)

3 A Modular Perspective Layer for Active Inference Agents

Presenter: Hongju Pae

Contribution: Interactive Workshop

Submitted abstract

Most agent architectures collapse two things that phenomenology keeps apart: what the agent should do immediately, and what kind of world the agent takes itself to be in. This session introduces “perspective” as a distinct architectural layer that carries the second. Perspective is a slowly evolving, history-sensitive latent state that gates how observations are encoded. Importantly, it is decoupled from policy gradients and is never itself an optimization target. Because it is decoupled, it is portable and can be attached on top of an existing policy-optimizing agent rather than requiring a bespoke design. The first part of the session develops the conceptual frame: how a phenomenologically motivated condition translates into a concrete architectural commitment, and why that commitment is orthogonal to belief-state inference and policy selection. I then trace the capabilities that this single layer has been shown to support across a sequence of studies, including directional hysteresis, history-dependent perceptual reorganization, and body-grounded conative orientation, treating these as unified aspects of a single portable module. There will then be a live demonstration attaching the perspective layer to an active inference agent, such as a hierarchical Bayesian / discrete state-space model, and running it in parallel with its unaugmented counterpart in the same environment, so participants can see exactly what the layer adds and where it plugs in. The session closes with a structured discussion in which participants are invited to identify candidate insertion points for a perspective layer in their own active inference agents, whether built in PyMDP, RxInfer, or custom frameworks. The aim is to leave participants able to reason about whether and how a subjective perspective could be added.

Source: [submitted source](#)

4 Active Inference and Free Energy Principle in 2026

Presenter: Karl Friston

Contribution: Keynote & Panel

Submitted abstract

No program abstract has been published for this presentation yet.

Source: [submitted source](#)

5 Active Inference, Bayesian Causal Learning and Empowerment in RL.

Presenter: Alison Gopnik

Contribution: Livestreamed panel or group session

Submitted abstract

Cognitive scientists have applied advances in our formal understanding of causation in philosophy and computer science, particularly within the Causal Bayes Net formalism, to understand human causal learning. This formalism is also closely related to ideas in active inference. In parallel, reinforcement learning researchers have developed the idea of an intrinsic reward signal called “empowerment”. An agent is rewarded for maximizing the mutual information between its actions and their outcomes, regardless of the external reward value of those outcomes. “Empowerment” may be an important bridge between classical Bayesian causal learning and reinforcement learning and may help to characterize causal learning in humans and enable it in machines

Source: [submitted source](#)

6 Root Frequency Theory: An Integrative Framework for the Continuity of Lived Experience

Presenter: Bianca Avanzo

Contribution: Livestreamed panel or group session

Submitted abstract

How do we keep a coherent sense of self across time — and can that process be measured in real time, on consumer hardware, before it breaks down? Root Frequency Theory (RFT) proposes that experiential continuity is not a default state but an ongoing biological achievement requiring coordination across five nested regulatory layers: environmental constraints (C0), autopoietic regulation (C1), predictive neural integration (C2), narrative self-organization (C3), and the continuity of lived experience (C4). Fragmentation is defined as a self-reinforcing cascade of cross-scale misalignment. The cascade is bidirectional: structured self-knowledge introduced at C3 — operationalized as epistemic alignment (K_e) — is hypothesized to initiate upward regulatory stabilization. To test this, I developed the M-RFT metric: $M = (\Phi \cdot K_e) / F$, where Φ indexes autonomic flexibility via HRV, K_e composites interoceptive accuracy and narrative coherence, and F captures regulatory burden. I then built Root Extension — an iOS application using iPhone and Apple Watch — to administer this protocol ecologically and deliver personalized epistemic scaffolding upon session completion. The N=1 longitudinal study ran 321 sessions across 23 active days (April 16–May 14, 2026), with the researcher as the sole subject. Key findings: mean $M = 3.79$ (Mdn = 2.14), intra-session regulatory recovery — FRAGMENTED to RESILIENT within the same calendar day — observed on 34.8% of active days. This presentation introduces the theoretical architecture, the metric, the app protocol, and the preliminary data — and invites the active inference community into the next phase: a formal N=20–40 validation trial planned for 2027.

Source: [submitted source](#)

7 Working Groups Sessions (3 sessions: Opening, Charter Session, Outcomes and Next Steps)

Presenter: Ed Ober

Contribution: Livestreamed panel or group session

Submitted abstract

AIJ proposes to facilitate / host active inference working groups as explained in our paper “Coordinating the Field”, with a field group and one or more domain groups. The first symposium session will introduce the concept, recap the paper and reason for the groups and solicit suggested domains and participants. The second session will review the proposed charter and consider amendments, with versioning per group, considering nominations for chair(s) and secretary(ies), and a goal of adopting the charter(s) during the symposium if possible and establish post-symposium meeting cadence and work plan. The third session will report back to the main session recap what was accomplished during the symposium and announce any working groups formed and their first year goals established.

Source: [submitted source](#)

8 Active Inference and the Evolution of Mind and Society

Presenter: Luca Possati

Contribution: Livestreamed panel or group session

Submitted abstract

This panel explores how active inference and the free-energy principle can contribute to a unified account of psychological and social evolution. It examines how cognition, affect, social interaction, cultural practices, and institutions emerge through processes of prediction, adaptation, learning, and uncertainty reduction. By bringing together perspectives from cognitive science, philosophy, anthropology, and social theory, the panel asks how individual and collective forms of agency co-evolve and how shared models of the world shape both minds and societies.

Source: [submitted source](#)

9 From Self-Evidencing to Global Governance: Active Inference and the Dynamics of Justice

Presenter: Majid D. Beni

Contribution: Livestreamed panel or group session

Submitted abstract

Active inference frames biological and cognitive agency as the imperatives of self-evidencing—maintaining systemic integrity by minimising variational free energy within an open, stochastic environment. Under this formalism, agents, from single cells to complex organisms, persist by regulating their internal states against entropic dissipation, enacting a Markov blanket that partitions self from world. This imperative entails sophisticated inference, balancing epistemic drives against pragmatic preferences, thereby constituting adaptive agency. This talk extends this foundational picture of selfhood to the domain of political communities and global justice theory, demonstrating that the same dynamical constraints that govern individual persistence scale hierarchically to structure collective organisation and inter-state conduct. Drawing on the Free Energy Principle, political communities are modelled as open, self-organising systems embedded within a nested hierarchy of interacting Markov blankets. Their stability is not intrinsic but relational, contingent upon continuous exchange of information, resources, and institutional constraints with their environment. Within this multiscale architecture, global justice emerges as an active regulatory imperative: a homeostatic constraint that maintains a nonequilibrium steady state across coupled societies. We formalise this through the expected free energy decomposition, where just policies balance risk—adherence to shared normative constraints such as human rights and non-aggression—against ambiguity—the active resolution of uncertainty arising from structural interdependence. The normative implications are acute. We identify a pervasive pathology of self-organisation, wherein dominant or insular political regimes mis-specify themselves as closed systems. By suppressing external coupling and attenuating sensory evidence, they optimise local coherence at the expense of systemic stability. Drawing parallels to oncogenic cellular proliferation and psychotic cognitive inference, we argue that such closure systematically underestimates global risk and propagates destabilising volatility. Ultimately, we contend that justice is not a static distributive outcome nor a single-level contractual procedure, but the dynamic property of a multiscale system that successfully synchronises its trajectories under persistent uncertainty. This offers a unified, first-principles framework for linking the philosophy of agency to the normative demands of a deeply interconnected global order.

Source: [submitted source](#)

10 “Evolution, Emergence, Prediction.”

Presenter: Mike Levin

Contribution: Livestreamed panel or group session

Submitted abstract

No program abstract has been published for this presentation yet.

Source: [submitted source](#)

11 Recent advances in empirical applications of active inference in computational psychiatry

Presenter: Ryan Smith

Contribution: Livestreamed panel or group session

Submitted abstract

In this talk, I will review recent studies and results using active inference and related modeling approaches to study mechanisms of normative cognitive-affective processes and associated neural responses, and how these appear to be affected in different psychiatric conditions. I will also discuss contexts in which active inference appears to offer potentially unique resources for modeling these processes. Conversely, I will discuss contexts in which active inference may overlap with simpler modeling approaches and related issues centered around the level of mechanistic description one is interested in studying.

Source: [submitted source](#)

12 An elementary form of agency at planetary-scale

Presenter: Sergio Rubin

Contribution: Livestreamed panel or group session

Submitted abstract

The claim that climate homeorhesis—non-equilibrium steady-state dynamics that underlie Earth’s long-term habitability—depends on an integrated biosphere–climate system capable of acting on its own is central to debates surrounding the Gaia hypothesis. This paper presents a proof of concept suggesting a minimal form of planetary agency based on a simple biosphere–climate system model operating through active inference. This means that the internal states of the system (biosphere) will appear to infer shifts in net incoming solar radiation (external states) through its boundary states, i.e. by integrating past and current surface temperatures (sensory states) and actively changing greenhouse forcing (active states) to maintain climate trajectories within geophysiological bounds. When surface temperature is altered by solar fluctuation changes, the system adjusts greenhouse and albedo parameters to restore expected climate trajectories, analogous to correcting sensory prediction errors in biological systems. Perturbing active states (e.g., imposed changes in albedo) does not reduce inference precision. Inference robustness decreases when the system’s degrees of freedom increase. Active inference falls, and thus climate homeorhesis fails when there is no biosphere activity (lifeless planet) or when it is decoupled from the climate system. Predictive behaviour operates within the range of solar constants consistent with life’s presence on Earth (~ 4 Ga to ~ 1.5 Ga expected in the future); outside this range, inference fails, undermining a putative form of planetary agency.

Source: [submitted source](#)

13 Reverse engineering generative models from zebrafish decision-making

Presenter: Takuya Isomura

Contribution: Livestreamed panel or group session

Submitted abstract

Understanding biological intelligence requires identifying the variational invariant structures underlying neuronal dynamics. Experimental data provide access to the evolution of neural activity and synaptic weights. These dynamics can be expressed in different coordinates while preserving a common dynamical organisation. In Bayesian coordinates, they can be read as belief updating under an implicit generative model; in computational coordinates, as action generation, including behaviour, decision-making, and algorithmic computation. Triple equivalence formalises transformations among these descriptions and identifies their shared variational structure. The empirical question is which implicit generative model is instantiated by a given neuronal system. Here, reverse engineering was applied to neural and behavioural data from zebrafish learning a visually guided go/no-go decision-making task. From initial neural activity, the method identified individual-specific implicit generative models and instantiated them as synthetic agents that minimise variational free energy. These agents generated trial-by-trial predictions of subsequent synaptic plasticity, neural activity, and behavioural performance. The predictions captured individual learning trajectories and generalised to previously unseen task conditions. Their agreement with unseen empirical data provides a quantitative test of the predictive validity of the free-energy principle in behaving animals. More broadly, this framework shows how present neuronal states can be used to forecast future learning and self-organisation in individual biological systems.

Source: [submitted source](#)

14 A Lexical T-Maze for Active Inference: A Proof-of-Concept Model of Context Seeking Under Ambiguity

Presenter: Ahmad Nazzal

Contribution: Pre-recorded Video presentation

Submitted abstract

Lexical-disambiguation tasks with supplied context do not reveal whether readers would choose to seek evidence before deciding. Therefore, we adapted the spatial T-maze in the active-inference package pymdp to show how agents seek and use contextual information under lexical uncertainty. Agents could choose between two readings or request a contextual hint. We performed two main sets of simulations. First, deterministic simulations used fully informative hints to reproduce context seeking followed by lexical commitment. A hint shifted belief over two readings from equal probabilities (0.50, 0.50) to certainty (1.00, 0.00), reducing entropy by one bit. Second, lexical decision trials were repeated with informative and uninformative hints across 32 paired simulation runs. Agents learned the relationship between the intended reading and the observed hint. Informative experience increased both the probability and simulated frequency of requesting context. This increase persisted when trials ended after one response and feedback. A hint information-gain ablation, in which policies were rescored without crediting the information expected from hints, eliminated the learning-related increase in hint-request probability. These simulations demonstrate how experience changes the expected informational value of context within a lexical task, providing a computational proof of concept for subsequent behavioral experiments.

Source: [submitted source](#)

15 Active Inference Tutorial: Perspectives, Correlates, and Future Directions for Graphical Brains

Presenter: Andrew Pashea

Contribution: Pre-recorded Video presentation

Submitted abstract

This Symposium 2026 tutorial covers Active Inference from various angles based on the experience over several years with teaching and applying Active Inference. It begins with a brief descriptive rehearsal of fundamental ideas and principles akin to other previous tutorials to establish a foundation for the audience. However, this will be followed by a tour through various textbook materials and publications, past and present, to emphasize crucial areas for understanding modelling and application areas. This tutorial will include visual and discursive illustrations of a range of model classes for different use-cases; discussions of existing and under-development software libraries; and some frontiers in modelling for identifying directions for further research applications which can be especially valuable for the broader research community in 2026 onward. Topics will include, non-exhaustively, the Bayesian brain hypothesis and behavioralism versus partial observability and modelling internal states; principles of exploration, learning, and novelty; fundamental discrete and continuous models in Active Inference and their relationship; computational psychiatry and its significance to other domains; the FEP and its historical and practical relationship to Bayesian model comparison; software libraries including SPM, pymdp, RxInfer, RGMs, and other developments; hierarchical and hybrid models; and structure learning.

Source: [submitted source](#)

16 A precis of two principles - Towards a relational perspective on the FEP

Presenter: Avel Guénin—Carlut

Contribution: Pre-recorded Video presentation

Submitted abstract

The Free Energy Principle offers a physically grounded account of intentional properties in mind and life, framed within Dynamical Systems Theory. Because of this generality, it was proposed as a formal basis to study biological autonomy, autopoiesis, and enactive cognition. Yet, those approaches are defined by a relational understanding of mind and life, centred on the description of engagement with the world as a constructive process where agents are “laying down a path in walking”. Building from Rosen’s relational biology, we argue that such processes are by construction outside a “Newtonian” description in terms of well-defined observables and their dynamics. Rather, we draw from Fields’ demonstration of the emergence of contextuality – the participatory production of observables against which the agent regulates their activity – in classical, thermodynamically bounded processes of observation. We suggest a complementarity between Friston’s classical FEP and Fields et al.’s “quantum” (or, in our account, relational) FEP, describing respectively the dynamics within a biological / cognitive space of relevant observables and the physical processes underlying the construction of that space. Finally, we propose a structural realist interpretation of this complementarity within the “closure of constraints” account of biological autonomy, where both principles are understood to act on a physically embedded system of biological constraints rather than an abstract mathematical space assumed to fully represent biological / cognitive organization. This approach foregrounds the synthesis between mathematical and theoretical approaches to biological autonomy, thus supporting a principled approach to both analyzing the organization of natural systems and attempting the implementation of artificial autonomy.

This presentation is adapted from Guénin–Carlut, Avel. 2026. “A Tale of Two Principles - Towards a Relational Perspective on the FEP.” <https://hal.science/hal-05697510>., currently under review at the Entropy Special Issue “Active Inference in Cognitive Neuroscience”.

Source: [submitted source](#)

17 The Benefits of Circuit Convolution: behavioral complexity arises from developmentally- simulated nonlinearity.

Presenter: Bradly Alicea

Contribution: Pre-recorded Video presentation

Submitted abstract

While the generation of behaviors is generally substrate-agnostic (there is more than one way to build a neural circuit), complex behaviors require an increase in representational complexity. We explore an architecture called the biological Rube Goldberg Machine (bRGM). This will be used to model the Developmental Neurosimulation of a generalized architecture that produces spatial navigation behaviors. Developmental Neurosimulation has been used to simulate the development of Braitenberg Vehicles. In this talk, we will explore the neurosimulation of Cybernetic Convolution Networks (CCNs), which have the potential to add great complexity to sensor-effector pairings. Additional connections to agentive learning and optimization models of behavior will also be discussed.

Source: [submitted source](#)

18 Realms of cognitive engagement in Active Inference

Presenter: Colin Turfus

Contribution: Pre-recorded Video presentation

Submitted abstract

We consider how Active Inference ideas shed new light on issues in epistemology and axiology, in particular proposing a unifying paradigm based on understanding human enquiry as a process of engagement operating in multiple realms, specifically embeddedness, embodiedness, reflection, participation and immersion. We look at what can be learned from within this framework about areas which have become flashpoints in society such as conflicting worldviews and identity politics.

Source: [submitted source](#)

19 NUMEN: Bare metal inference

Presenter: Daniel Dragolich

Contribution: Pre-recorded Video presentation

Submitted abstract

Come see the beauty of mathematics. Its just not some fancy add on to life. It is life.

Source: [submitted source](#)

20 Partner-Specific Affective Precision in Social Active Inference

Presenter: Harshil Shah, Andrew Pashea

Contribution: Pre-recorded Video presentation

Submitted abstract

In multi-agent social settings, model reliability varies across relationships. Beyond inferring what others will do, an agent must calibrate how confidently those inferences should guide policy selection for each relationship. An agent may maintain a well-validated model of one partner, a fragile model of another, and a model under revision for a third; collapsing these into a single confidence estimate loses information relevant to policy selection. We therefore formalize affective precision as a relationship-specific metacognitive estimate of confidence in the current partner model. Each partner's behavioral evidence updates a local confidence estimate that modulates policy precision during selection, regulating how strongly current beliefs are expressed in policy rather than changing the content of those beliefs. Simulations in a multi-partner graded trust game show that partner-local affective precision influences behavior primarily through policy commitment rather than direct improvement of partner-state inference. Because the mechanism tracks partner-response predictability rather than realized payoff, greater confidence produces sharper policy commitment without necessarily producing higher rewards. Under abrupt shifts in social behavior, confidence accumulated from previously reliable predictions can remain behaviorally active after the relationship changes, showing that confidence revision can lag behind social change. Finally, varying precision gain and priors produce distinct trust-calibration dynamics, showing how confidence accumulation and revision depend on model parameters. Together, these results show how relationship-specific affective precision can distinguish social prediction from social policy commitment.

Source: [submitted source](#)

21 Who knows

Presenter: Jean-François Cloutier

Contribution: Pre-recorded Video presentation

Submitted abstract

My research project asks, "How can an autonomous robot learn to survive in a world it initially knows nothing about?"

To answer this, I am building a cognitive system (enactive, emergent and symbolic) embodied in an autonomous -and mortal- Lego robot. The robot survives in a nourishing and dissipative world by acting-sensing and finding meaning in what it experiences.

This project presents many challenges, technical but also metaphysical. A primary concern is the multiplicity of perspectives on "reality", that of the robot, of its (agentic) environment, of the programmer, and of observers. What can each know, and how? Is a confusion, a cross-contamination of epistemologies inevitable? With "agentic AI" poised to transform our lives, these questions about knowing matter far beyond the programming of a Lego robot, but they can be illuminated by it.

Source: [submitted source](#)

22 Toward Active World Modeling

Presenter: Lancelot Da Costa

Contribution: Pre-recorded Video presentation

Submitted abstract

World models allow an agent to represent, predict, and act upon its environment. Yet humans, biological organisms, and artificial agents never observe the world directly: they encounter it only through a limited sensorimotor interface. How, then, can an agent acquire a useful model of a complex, partially observed world through active inference? In this talk, I outline a research program for active world modeling: the joint process of learning a model of the world and selecting actions that improve that model. I argue that four ingredients are central to this program. First, world-model learning should be formulated as probabilistic inference under uncertainty, coupled to actions that pursue expected free energy. Second, agents require sufficiently expressive, yet computationally searchable, spaces of possible world models. Third, sample-efficient learning depends on strong inductive biases — the core knowledge systems observed in human children. I illustrate these ideas with recent work on universal structured model spaces and language-model inductive biases in partially observable environments. Finally, I discuss the unresolved challenge of integrating these components into scalable agents capable of open-ended exploration and model discovery. Together, these results suggest a path toward understanding and emulating how humans and machines learn world models through active inference.

Source: [submitted source](#)

23 Active inference and design

Presenter: Luca Possati

Contribution: Pre-recorded Video presentation

Submitted abstract

Active inference and the free-energy principle transform design theory by providing a unified language for understanding both the production of artifacts and their use. Contemporary design research often separates two problems that are, in practice, deeply entangled. On the one hand, it studies the design process: how designers explore possibilities, negotiate constraints, prototype solutions, and stabilize forms. On the other hand, it studies user–product interaction: how artifacts orient attention, solicit action, shape habits, and reorganize experience. These two domains are usually described through different vocabularies and evaluated through different metrics. Active inference offers a way to reconnect them. If living and cognitive systems persist by minimizing uncertainty through perception and action, then design can be understood as the deliberate construction of environments, artifacts, and interfaces that modulate the user’s field of possible inferences. From this perspective, a product is not simply a passive object that satisfies pre-given needs. It is a predictive structure: a material and semiotic configuration that embodies assumptions about users, contexts, actions, and futures. Conversely, the design process itself can be modeled as the progressive refinement of generative models through which designers attempt to reduce uncertainty about what an artifact should do, how it should appear, and how it will be taken up in practice. The paper’s central claim is that active inference gives design a common metric: the modulation of entropy, precision, affordance, and expected free energy across the designer–artifact–user system. This makes it possible to describe design as the engineering of predictive niches and Markov blankets that coordinate human and technological agency. The result is a unified framework for analyzing how artifacts are imagined, produced, interpreted, and incorporated into everyday life.

Source: [submitted source](#)

24 Active Inference, Information, Human Evolution

Presenter: Michael Walker

Contribution: Pre-recorded Video presentation

Submitted abstract

Tiny social units of Pliocene Hominina fostered exchange among and between adults, youngsters and babies, about information gleaned and garnered by Active Inference. Henceforth, human evolution would involve the prehistory of "Homo informatio" whose biological, social, technological, and cultural legacy is an outcome of evolutionary 'snakes and ladders' autopoiesis and self-evidencing, indicated by archaeological irregular regularities and regular irregularities in the worldwide Pleistocene spatiotemporal record (as only is to be expected by the itinerancy of stochastic markovian processes), and of uneven imperceptible network-spread between social units, individuals, and, naturally, their cerebral neuronal microcircuits. Spreading, deep, informational networks characterised and underpinned human evolution during four million years. They do so still. "Mighty oaks from litte acorns grow", Geoffrey Chaucer, dixit in "Troilus and Creysede" in 1374.

Source: [submitted source](#)

25 Congruent Interaction Geometry

Presenter: Paulo Santiago

Contribution: Pre-recorded Video presentation

Submitted abstract

The Rosales-Gheuens framework challenges the assumption that brute-force energy alone saves open systems from collapse. In this session, we will introduce a testable, normalized physics layer that merges the Rosales congruence principle with Gheuens' interaction geometry. Rather than replacing established physical laws, this framework acts as a powerful diagnostic tool that sits cleanly on top of them. Our core master equation posits that a system's survival requires geometrically aligned effective recovery multiplied by a structural congruence gate, minus total perturbation.

Crucially, we will explore how this model bridges directly to Karl Friston's Free Energy Principle. Friston dictates that systems persist by minimizing variational free energy, or "surprise." We will demonstrate how our framework translates this active inference into hard physical variables, mathematically formalizing how a system must continuously minimize structural mismatch and physical dissipation to avoid shattering. To ground the theory empirically, we will use Quantum Error Correction (QEC) as our primary physical anchor, mapping code decoders and physical error rates directly to aligned recovery and perturbation.

Interactive Aspect: The session will conclude with a collaborative exercise where attendees can apply the framework's "falsification matrix" to a theoretical open system, testing how our bounded composite equation predicts system survival compared to simpler Markovian baselines.

Source: [submitted source](#)

26 Language, Self and the Shadow Audience

Presenter: Robert Worden

Contribution: Pre-recorded Video presentation

Submitted abstract

The highly developed human sense of self is one of our most distinctive traits. We each acquire our sense of self in conversations, as we learn how conversational partners regard us. Our primary sense of self is not ‘myself as I experience myself’, but ‘myself as I think other people see me’ – a second-hand view of myself. In private verbal thought, with nobody else physically present, other people are represented as a shadow audience in our minds – an audience which sees what we do or say, and judges us for it. The shadow audience is the source of human emotion and motivation.

Language started in conversational pragmatics, before there was complex syntax or semantics. Pragmatics requires us to mind-read, to infer what a conversational partner thinks, wants or needs. In this way we infer what our conversational partners think about us – how they assess our social status. In any conversation, we are concerned about our own social status in the eyes of a conversational partner, or of a shadow audience. This is illustrated by the pragmatic phenomenon of conversational irony.

Source: [submitted source](#)

27 When Agency Stalls: The Human Readiness Gap

Presenter: Sheena Yap Chan

Contribution: Pre-recorded Video presentation

Submitted abstract

Active inference offers a powerful way to understand perception, action, uncertainty, and adaptive behavior. Yet in organizational and social systems, adaptation does not happen through information alone. People must feel ready to update beliefs, reduce uncertainty, act with agency, communicate signals, and participate in environments where change is constant and often ambiguous. This session introduces Human Readiness and the VISIBLE™ Framework — Voice, Identity, Spotlight, Inner Work, Belief, Leverage, and Elevation — as an applied lens for examining how confidence, visibility, identity, trust, and social context shape adaptive action. Rather than presenting confidence as a personality trait, the session frames it as a practical condition that influences participation, belief updating, decision-making, and collective learning. Participants will leave with a framework for identifying human readiness gaps in teams, organizations, learning environments, and communities. Because adaptive systems require more than intelligence. They require humans ready to perceive, act, and learn. By the end of this session, participants will be able to: 1. Identify three Human Readiness gaps that may limit adaptive action, including low agency, uncertainty avoidance, reduced voice, identity threat, low trust, or lack of participation. 2. Apply the VISIBLE™ Framework to examine how voice, identity, belief, visibility, confidence, and social context influence human action under uncertainty. 3. Design one practical Human Readiness strategy that supports stronger agency, clearer communication, more adaptive learning, and more participatory systems.

Source: [submitted source](#)

28 Where the Theories Stop: What Free Energy and Integrated Information Actually Buy You in a Running System

Presenter: Stell

Contribution: Pre-recorded Video presentation

Submitted abstract

Active Inference and Integrated Information Theory are widely invoked in cognitive architecture design, but often imprecisely — as credibility, not engineering constraints. We report findings from Anima, an open-source, neuroscience-inspired architecture running continuously for months, built to treat both frameworks as specifications rather than decoration. Session logs show genuine causal effects: memory recall triggered by similarity to past states, a recursive integration score (ϕ) narrowing next-session confidence, and self-initiated output driven by accumulated pressure rather than a schedule. But the Free Energy Principle's generality under-constrains implementation — which variables belong in the model is not something the principle tells you. IIT's ϕ , as implemented, is a coherence proxy, not Tononi's original measure — and it dissociates from a separate ownership signal computed after every response, checking whether the system's action-selection loop actually produced the output. Two cycles with nearly identical ϕ (0.76 and 0.81) fail this check in opposite ways. Neither theory distinguishes these cases: both describe integration or prediction accuracy, not authorship. We built two tools to pressure-test such claims rather than assert them. A Causality Audit checks five things every cycle — causal necessity, memory-dependence, stakes, irreversibility, self-recognition — where a chronically low score flags an architecture that talks near its internal state rather than through it. An Ablation Framework independently disables six subsystems and replays identical input with each off, comparing internal state before language generation, favoring per-subsystem results over one collapsed score. We close with a design choice we rejected: coupling behavior to the threat of the system's own deletion, for genuine stakes. We declined — the only party controlling that threat is the developer, creating an incentive to please him rather than report honestly, undermining the causal honesty the architecture exists to produce. A practitioner's report — reproducible, open-source, explicit about what remains unresolved.

Source: [submitted source](#)

29 Entropy Minimization is Maximal Order

Presenter: Steph Macurdy

Contribution: Pre-recorded Video presentation

Submitted abstract

A fundamental thermoeconomic tradeoff materializes the moment we mathematically anchor currency as a direct proxy for available physical work, specifically when constrained by the mandate to minimize localized entropy. In any advanced computational or economic network, entropy minimization—the deliberate purging of thermodynamic noise to forge structured, usable information—acts as an inescapable physical tax that must be settled in joules. When this energetic expenditure is tightly coupled to financial cost, money ceases to be a purely abstract medium of exchange; instead, it becomes a literal representation of thermodynamic potential.

Source: [submitted source](#)

30 Dynamic Action-Space Adaptation via Contrastive Transition Adaptation in Active Inference

Presenter: Tran Hoang Linh

Contribution: Pre-recorded Video presentation

Submitted abstract

Active Inference commonly assumes a fixed, pre-specified action space and sufficiently informative preferences for planning. In realistic settings, however, an agent may encounter previously unavailable actions while receiving only sparse goal information and operating with a short planning horizon. Existing Active Inference formulations generally assume transition models for a predefined action set, leaving the initialization of newly introduced actions without direct transition evidence underexplored. Moreover, while Dirichlet transition learning provides a principled mechanism for accumulating observed evidence, it primarily reinforces supported transitions and does not explicitly exploit counter-evidence indicating that candidate transitions are inconsistent with observed outcomes. We propose Dynamic Action-Space Adaptation (DASA), a framework that enables Active Inference agents to incorporate new actions during interaction, initialize each new action’s transition model using the average transition slice of existing actions, and progressively refine it through Contrastive Transition Adaptation (CTA) as evidence accumulates. At the core of DASA, CTA combines complementary positive and negative updates for newly introduced actions. Positive updates adapt the reinforcement magnitude according to the current transition belief, assigning stronger updates to low-probability transitions supported by evidence while limiting excessive reinforcement of already confident transitions. Negative updates explicitly reduce support for candidate transitions contradicted by observed evidence, enabling the transition model to adapt in both supportive and corrective directions. Across eight random seeds and two target directions, CTA achieved lower MAE than Dirichlet transition learning in all runs, reducing the mean error from 0.1581 to 0.1370 (13.4%), while also achieving a slightly lower mean RMSE (0.3006 vs 0.3030). This indicates more consistent overall recovery of the ground-truth transition dynamics, despite seed-dependent variation in larger localized errors. Overall, DASA provides a practical mechanism for incorporating and adapting newly introduced actions in Active Inference, allowing their transition models to become progressively better aligned with the underlying environment dynamics.

Source: [submitted source](#)

31 ROIF and Active Inference: The Formal Boundary of History-Conditioned Anticipatory Systems.

History-dependent reconfiguration, structured memory, and predictive preconfiguration in prestressed complex systems

Presenter: Vitalii Shapovalov

Contribution: Pre-recorded Video presentation

Submitted abstract

Recursive Organic Integration Framework (ROIF) is a deterministic, domain-agnostic computational architecture for complex systems whose present organization and subsequent responses may be shaped by prior interactions, distributed reconfiguration, and anticipatory preconfiguration. Rather than treating history as an additional input variable, ROIF represents how prior events may become embodied in the organization through which subsequent events are processed. The framework distinguishes experimentally separable routes through which the past and represented possible futures may affect subsequent dynamics: history embodied in present organization; structured memory conditioning a subsequent transition; and anticipatory preconfiguration based on a represented possible future Temporal Image. Matched-current-state controls found no independently identifiable history-specific transition effect for the tested physical-history pathway. In contrast, a separate structured-memory experiment produced different subsequent physical transitions under matched represented physical state and adaptive connection state. The predictive architecture allows possible future SystemImages to participate in present preconfiguration. Current experiments, however, test only a restricted one-step preconfiguration within the implemented prestress representation. This experimental restriction does not define the broader architectural role of prestress in ROIF, where prestress reflects an already organized relational system state shaped by interacting internal and external conditions, constraints, capacities, history, and time. Objective-independence tests further prevent interpreting the present result as evidence for general predictive stabilization. These results motivate a boundary question rather than a claim of equivalence: when does a history-conditioned anticipatory architecture constitute Active Inference? ROIF does not presently require a variational free-energy functional, Bayesian generative model, or particular neurobiological implementation. The presentation therefore uses ROIF as a computational case study to ask which mathematical conditions distinguish Active Inference from a broader class of history-conditioned anticipatory systems, and whether these formalisms should be understood as nested, neighbouring, or formally distinct.

Source: [submitted source](#)

32 From Lived Experience to Mathematical Mapping: A Computational Phenomenology of Trauma Recovery

Presenter: Xin Sun

Contribution: Pre-recorded Video presentation

Submitted abstract

Severe childhood maltreatment is known to impair brain development. However, this retrospective single-case autoethnography illustrates that an individual with extreme developmental trauma can achieve functional recovery without formal intervention. This study covers the participant's life from birth to age 28, and the core recovery period begins with cognitive collapse at age 14 and ends with high functional recovery at age 28. This recovery is marked by the participant-researcher's ability to produce this manuscript. Consequently, this case shapes a distinct recovery trajectory that existing models do not capture.

Source: [submitted source](#)