Welcome

The Active Inference Institute supports education, research, and applications of Active Inference

Active Inference Institute (AII) is a participatory Open Science institute dedicated to improving the accessibility, rigor, and applicability of the Active Inference framework.

This is our main, interactive, living document (chat with this document directly), containing pages and sections on The Active Inference Institute and The Active Inference Ecosystem, such as: Institute Programs, Projects, and Ecosystem Support.

As of 2024 we are a 501(c)(3) educational non-profit organization (donate and Philanthropy).

All backgrounds, time zones, and familiarity with Active Inference are welcome to Get Involved in the Institute Programs and Activities.

Learn more About the Institute & Ecosystem

- See our playlist of quarterly updates through time for more on The Active Inference Institute, including the History of The Institute and Institute Organization.
- Join the Discord for text and voice chats with others.

Check out Activities and Institute Programs:

- Engage in learning on your own time, and join synchronous activities when you can.

Projects

- Ecosystem Projects
- Institute Projects
- Project ~ Preparation for proposing projects at the Institute
- Project ~ Measurement for reporting updates from The Active Inference Ecosystem

How to get started with Active Inference?

- Explore this living Coda document
  - See the Start page
- See our past and upcoming Livestreams (Videos and Podcasts)
  - Listen to the conversational *Active Inference Insights* podcast on YouTube or Spotify.
- Textbook Group
- Explore various roles and Affordances, which may fit you or others you know:
  - Volunteer, Internship, Mentorship
  - Fellows and Partnership
  - Scientific Advisory Board, Board of Directors, Officers
- Readings and Research:
  - The 2022 Textbook: "Active Inference: The Free Energy Principle in Mind, Brain, and Behavior" by Thomas Parr, Giovanni Pezzulo, Karl J. Friston — See Textbook Group to learn this material in a collaborative setting.
  - "The free energy principle—a precis" by Maxwell Ramstead (October 2023)
  - "Spinning Up in Active Inference and the Free Energy Principle" by Jared Tumiel
  - "FEP and Active Inference Paper Repository" by Beren Millidge
  - "Active Inference & Behavior Engineering for Teams", Vyatkin et al. 2020
  - A. Levenchuk, 2015 "Towards a Systems Engineering Essence"
  - "An Active Inference Ontology for Decentralized Science: from Situated Sensemaking to the Epistemic Commons", Friedman et al. 2022.
  - "Of woodlice and men: A Bayesian account of cognition, life and consciousness" 2018 conversation-style interview with Karl Friston.
- Code — see Implementations of Active Inference.

Active Inference Institute is active on the following platforms:

- **Discord:** discord.activeinference.institute
- **YouTube:** youtube.com/c/ActiveInference
- **X**: https://x.com/InferenceActive
- **BlueSky:** https://bsky.app/profile/activeinference.bsky.social
- **Podbean:** https://activeinference.podbean.com/
- **Facebook:** https://www.facebook.com/ActiveInference
- **LinkedIn:** https://www.linkedin.com/company/active-inference/

Email: blanket@activeinference.institute

How to participate?

There are many ways of learning, applying, and supporting Active Inference at the Institute.

- Start at: welcome.activeinference.institute
  - Email to: blanket@activeinference.institute
- Join the Discord: discord.activeinference.institute
- See projects: projects.activeinference.institute
  - textbook-group.activeinference.institute
  - rxinfer.activeinference.institute
- More about the Active Inference Ecosystem: ecosystem.activeinference.institute
- For a more structured learning journey & involvement with projects:
  - volunteer.activeinference.institute & intern.activeinference.institute
- For staying updated on happenings in the Institute/Ecosystem:
  - Past and upcoming videos: video.activeinference.institute
  - Newsletter: newsletter.activeinference.institute
- Fellows program: fellows.activeinference.institute
- Explore organizational Partnerships: partnerships.activeinference.institute
- Support the Institute (a 501(c)(3) non-profit): support.activeinference.institute

Activities

- See Welcome for more background and context on the Institute.
- These events can be added to your calendar with this link, and seen in the Discord Events section.
- Email Blanket@ActiveInference.Institute if you have questions about the activities or Institute
- This page shows the coming activities in the next 7 days, and information on all active projects of different types (e.g. Institute Projects,
- Feel free to drop in to any of these Activities whenever it works for you — see the table for information on how to get involved with each project.
  - See Projects for more information on proposing or measuring projects of your own.
  - Also see the Volunteer and Internship Institute Programs if you want to engage in a more structured way, and Affordances for specific contribution opportunities.

Activities at the Institute ~ Coming 7 days

Not synced yet

Date & Time (UTC)	Event name	Description
12/17/2025, 15:00	Active Inference GuestStream #125.1 ~ "From Charles Darwin's "Root Brain" to Nikola Tesla's "6G World Brain" and XAI-native 6G Networks"	
12/18/2025, 16:00	ReviewStream 2025	
12/19/2025, 18:00	2025 Quarterly Roundtable #4	
	CogNarr Ecosystem project	Information on CogNarr Ecosystem project

The Projects at Active Inference Institute are Institute Projects and Institute Programs.
Scroll down further to see Projects by Research Fellows, Scientific Advisory Board members, Current Partners, and Ecosystem Projects.

Projects at Active Inference Institute

Not synced yet

Organizational Unit	Project	Documentation	Mission & Objectives
▼ ReInference 6	RxInfer.jl learning and development group		Learn and apply RxInfer.jl in 2024 — building out multiscale se modeling.
	Knowledge Engineering		This project seeks to alleviate the information burden in the Ac through information curation, organization, and condensation- productions (courses, livestreams, etc), enhancing the CRM, e
	Active Blockference		We are applying Active Inference by building capacities & crea

	FarmWorks		Develop minimal model of personalized agents
	Applied Active Inference Symposium		To have a year-end Symposium, featuring applied Active Infere
	AICACP		AICACP is a multi-year initiative designed to reshape the conv and regulation.
▼ EduActive 7	Active Inference Ontology		Maintain, improve, elaborate, extend, translate, educate, docu Ontology as core infrastructure for the Active Inference Institu
	Audio-Visual Production		Produce accessible, rigorous, informative (epistemic value) an content, for example through Livestreams, Podcasts, and othe
	Active Inference Journal		To develop evolving hybrid (AI+people) project architecture an
	Textbook Group (Parr, Pezzulo, Friston 2022)		Improve the accessibility, rigor, applicability, and impact of the and Friston.
	Course Development		Develop educational materials and experience to increase fam practice.
	Applied Active Inference Symposium		Host an annual Symposium to highlight the state of the art in a
	Seasonal School		Develop in-person experiences for education and developmen

Projects in the Active Inference Ecosystem are organized by individuals in The Active Inference Ecosystem
These projects are submitted via the → Project ~ Preparation form.

Projects by Research Fellows

Projects by Research Fellows

Not synced yet

Project	Documentation	Mission & Objectives
Symbolic cognitive robotics		Explore the joint problem space of “symbolic active inference”, “societies of mind” and “morta computing”, with an emphasis on unsupervised learning. Using symbolic processing, build a rudimentary artificial agent (a LEGO rover robot) whose b fulfills the requirements of Active Inference
Active Inference Cycle Book for Self-Knowing		Perform a meta analysis of the “wellness” space through the lens of active inference highligh most impactful points for the larger population in an easily digestible format. Use this work t longer term collaboration and contribution to the larger AI community.

CogNarr Ecosystem: Facilitating Group Cognition at Scale		The initial mission is to advance the CogNarr project from its current incubation phase into a concept demonstration, followed by a minimal viable product. In concept, the CogNarr ecosystem of software and tools is designed to serve as a compone group's cognitive architecture.
Model-Centric cognition		Develop the central idea, raise awareness; assess whether this departure from the brain-as paradigm is needed. Some pushback is expected, even hoped for. The project is a developm existing wave hypothesis.
Humanity's Story of an Uncertain Self		Producing an academic paper or blog that contains a set of equations, computer simulations ultimately a framework that explains the core components of humanity's sociological-narrati framework. Specifically, breaking down a few pieces of say, ancient epics, along with an set economic and civic institutions, would allow us to model to simulate, predict, and give maste otherwise seemingly intractable world of humanity's cultural niche.

Projects by Ψ Scientific Advisory Board members

Projects by SAB members

Not synced yet

Project	Documentation	Mission & Objectives
Project Development for "Solving the Tower of Babel Problem: UniFysica Philo-sophia"		To outline, draft, a collection of papers with the title "An Inclusive System of Communication Shared Meanings and Cognition: From Blombos to Friston and Fields"
Creativity and creators under the light of the Free Energy Principle		Design and run experiments to answer the key questions
The Three Mosqueteers		Create a livestream aimed at disseminating science and helping people without a scientific b to adopt a more critical attitude toward the information they receive.

Projects by Current Partners (Partnership organizations)

Projects by Partner organizations

Not synced yet

Project	Documentation	Mission & Objectives
Numinia		First mission would be to make sure we are implementing Active Inference in the game prope well explained, another mission would be to ensure that the design of the incentives aligned values of *Numinia* and the All.

Ecosystem Projects

Projects in the Active Inference Ecosystem

Not synced yet

Project	Documentation	Mission & Objectives
Active Inference Account of Belief Updating in PTSD		Write a theoretical paper in the style of Parr et al. chapter 6
Improving RxInfer.jl's Model Visualization Capabilities		Our mission is to equip RxInfer.jl - and its relevant component libraries - with a host of model visualization modalities that prove useful to those who wish to use, and/or to develop RxInfe To that end, we anticipate measuring the initial quality of our contribution/s by their receptio RxInfer.jl's core developers: TU/e's BIASlab. All our objectives must therefore take the approv BIASlab as their proverbial North Star.
Neurodivergent Learning Sessions		Neurodivergent learning is focused on outreach and spreading awareness geared towards th struggle with standardized curriculum environments when it comes to public and higher edu milestones... as a number of people with neurological conditions not limited to autism spectr can struggle in varying ways with learning and being in the right environment in which inform presented to them in a manner which is coherent.
The Unordinary Bible Study (abbreviated as TUBS)		Hosting once a month sessions that focus on cross-referencing biblical verses but not spenc much time digging into scripture as opposed to focusing on inter-faith and contemporary pe focused dialogue.
The Einstein Model of a Solid as a Model of the Mental Apparatus from the Economic Perspective of Psychoanalytic Theory.		Bridging Psychoanalysis and Thermodynamics with applications to Artificial Intelligence. App AI.
Project Sweet (Sus) Dogg		To Help Warm-up or Prepare a Plausibly Notable Aspect of Agent Based Alignment By Social
Active Inference for Built Environments & CooperActive Systems		To advance the application of Active Inference in designing, managing, and evolving built en that prioritize the flourishing of all life on this planet. While humans possess unique cognitive capabilities, we recognize that excessive anthropocentrism blinds us to the needs of other li organisms. Our work centers on **life prosperity** as the foundational principle for all built envir decisions. We seek to develop adaptive, nature-integrated solutions through distributed intelligence, di technologies, and decentralized decision-making that serve the broader web of life while me human cooperative living needs.

Affordances

Specific opportunities for your contributions

Check out the Affordances (opportunities for action) table below, and email blanket@activeinference.institute, or follow specific instructions, if you are interested in exploring more:

Affordances

Not synced yet

Contribution sought	Details	Mo
Livestream and Podcast organizer/contributor	Have you enjoyed the Active Inference Institute videos/podcasts? https://video.activeinference.institute/ https://www.youtube.com/@ActiveInference/ I would GREATLY enjoy the collaboration of 1 or more people in planning and implementing the video production for 2025. There is a checklist for curating and inviting guests, on through implementing the recording or stream. Expertise in Active Inference is not required. Truly this is a great opportunity for people of any background, who want to learn more about the space, connect with the authors/researchers personally, and have a big impact in increasing the visibility and accessibility of Active Inference. If you might like to join on this journey in 2025 -- Email blanket@activeinference.institute with subject [PRODUCTION]	

Weekly Update

Announcements for week of December 15, 2025

Greetings. The Institute is enjoying a winter break until January 2026. Read on for some of the last updates of the year, ways for your end-of-year updates to get visibility, and areas of contributions for next year.

---

1. Submit Your End of Year Updates

We would love to include **your** updates in the December Newsletter and the Quarterly Roundtable. Please use the Measurement form for providing updates on your (research, learning, application) work, by December 17th. https://measure.activeinference.institute

---

2. Applications open for 2026 Scientific Advisory Board (SAB)!

The SAB is a collaborative group of professionals who informally advise the Institute and serve as reviewers, mentors, contributors, and co-creators. For the coming year of 2026, we welcome applicants with backgrounds in Active Inference, as well as more broadly in education, research, open source, technology, and professional service. Membership on the SAB requires a modest time commitment (0–few hours per month), communication skills, and a shared enthusiasm for advancing the Institute's mission and our broader field. We work to make the experience meaningful and streamlined.

If you would like to be considered, please complete this form before the end of December 2025. Additionally, if you know someone who would be great for this position, feel free to pass this opportunity along to them. All information: https://sab.activeinference.institute/

---

3. Upcoming Livestreams:

- GuestStream #125.1 ~ 12/17/2025 at 15 UTC with Nika Hosseini, Osman Tugay Bosaran, and Martin Maier
  From Charles Darwin's "Root Brain" to Nikola Tesla's "6G World Brain" and XAI-native 6G Networks
  https://www.youtube.com/live/LmgPFAINHNQ
- ReviewStream 2025 ~ 12/18/2025 at 16 UTC
  2025 Active Inference Livestream Review
  https://www.youtube.com/live/gS-qhMNFm84
- 2025 Quarterly Roundtable #4 ~ 12/19/2025 at 18 UTC
  https://www.youtube.com/live/09FfbL1YYOI
- See https://youtu.be/TaFwI2zr_IE for more information if you would like to contribute to the scheduling and production of future livestreams.

---

4. More:

- Join the Discord: http://discord.activeinference.institute/
- Learn more about the Active Inference Ecosystem http://ecosystem.activeinference.institute/

- Make a Measurement, to get your update included in the upcoming Newsletter: http://measure.activeinference.institute/
- We are a 501(c)(3) educational non-profit. Donate at: http://donate.activeinference.institute/
- Email blanket@activeinference.institute with any questions.

About

About the Active Inference Institute

Active Inference Institute (AII) is a participatory Open Science institute dedicated to improving the accessibility, rigor, and applicability of the Active Inference framework.

The Active Inference Institute is a registered non-profit organization (Delaware, USA) which identifies, establishes, scaffolds, and supports the sustainable implementation of:

1. Education and Research services.
    a. We learn, teach, research, and apply Active Inference
    b. We host Institute Programs and Institute Projects
    c. We provide visibility and opportunities for Ecosystem Projects
2. Participation, communication, advisory, governance, and meta-governance affordances within the Institute and The Active Inference Ecosystem
3. Publishing, and licensing protocols that establish Open Source, fair use, and effective dissemination of community products within and beyond the Ecosystem.
4. Ecosystem Support services such as Communications, EduActive (Education), Ecosystem Projects, and operation of cyber and cognitive security systems aligned with our Mission, Vision, Values, and Principles

---

To learn more About us, see:

- History of The Institute since founding in 2021
- Mission, Vision, Values, and Principles
- Focus Areas for the Institute in terms of ongoing challenges (*"where you find the challenge is where the learning/solution is done!"*)
- Directions for the Institute we are taking in light of the focus areas.
- Institute Organization, or morphology, in terms of roles and positions.
- Institute Programs and avenues for participation, such as Volunteer, Internship, Fellows, Philanthropy, Grants, Partnership.
- Institute Projects hosted by the Organizational Units

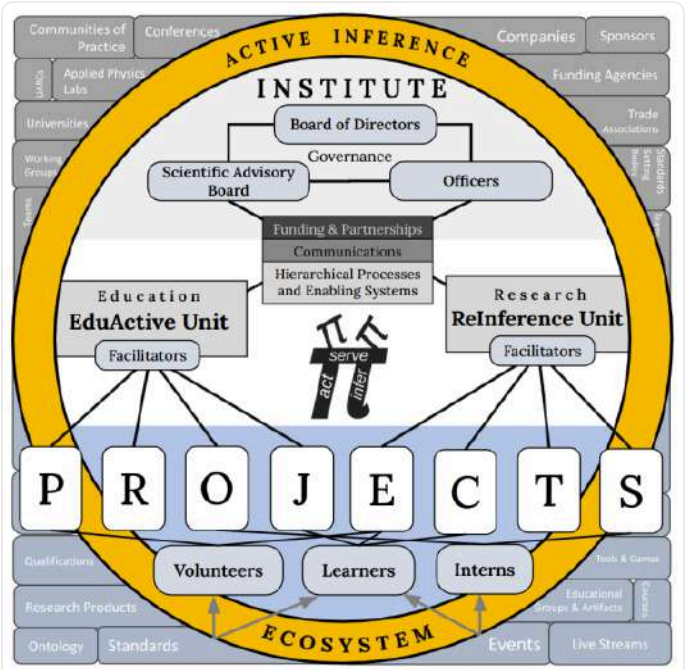

Figure 1. Active Inference Institute & the Active Inference Ecosystem.

Learn more about the ∞ Active Inference Institute & Ecosystem

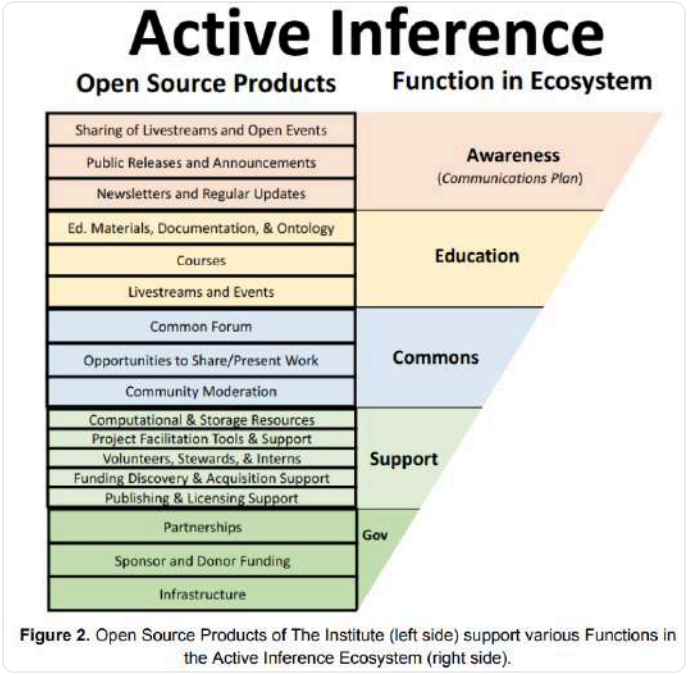

Figure 2. Open Source Products of The Institute (left side) support various Functions in the Active Inference Ecosystem (right side).

Above is a representation of the products and functions provided by the Institute Programs Ecosystem Support .

Institute & Ecosystem

This is the home page for the Institute & Ecosystem.

It is available as living document at https://ecosystem.activeinference.institute/ .

You can also have LLM-aided live chatting with the material via this chat link.

This document is structured according to the sections:

- Opening sections with information such as:
  - Abstract, Authors
  - Active Inference for an overview on Active Inference.
- The Active Inference Institute
  - Pages related to the history, projects, productions, goals, organizational anatomy, values, and people of the Institute.
- The Active Inference Ecosystem
  - Pages related to activities and areas of attention in the broader Active Inference Ecosystem.
- Discussion and Future Directions

This structure initially came from the sections and contents of the 2023 paper "The Active Inference Institute and Active Inference Ecosystem". From that starting point, during September-November 2024, the Authors made various contributions and additions to this living document.

This is a work in progress, and we will continue to update.
Past versions of this document: 2023 (version 1), 2024 (version 2).

Get in touch with any comments, questions, or inclination to assist, for example with curating Domains of Application or contributions to information about The Active Inference Institute.
Email: blanket@activeinference.institute

Abstract

This document surveys the current state of The Active Inference Institute and The Active Inference Ecosystem, in the context of our current and future directions. As embodied agents, we aim to update our decisions, goals and predictions as an institute by actively gathering (sampling) insights (observations) from our members. As Heraclitus once said "No one ever steps in the same river twice. For it's never the same river and it's never the same person". In the same way, the Institute evolves with each new member, accumulating a variety of perspectives to drive improvement.

A Unifying Framework Across Disciplines

Active Inference is a framework that emerges from studying how the interaction of living systems and their environment can be understood through the lens of information theory, drawing on important principles of thermodynamics — the science of energy and it's transformations. With solid theoretical grounding, this allows Active Inference to be applied to many areas of human behavior and social interactions.

As a framework, Active Inference is a powerful tool for thinking about systems that ecologically evolve over time. By bridging connections across multiple disciplines, ranging from computational neuroscience to ecological psychology, there has been a growing list of Domains of Application relating to its implications that continues to grow and expand. These applications range across a multitude of fields such as artificial intelligence, economics, law, governance, resource management, risk management, finance, decision theory and physics, highlighting the versatility as well as the value of a unifying framework to understand human behavior and it's adaptability across dynamic environments.

The Institute and You

The Active Inference Institute is an educational organization committed to promoting a better understanding of Active Inference and its potential benefits. Our goal is to build and provide a network of support to a wide audience of individuals interested in learning the foundational knowledge or practical applications of framework principles within civic, commercial, industrial and other domains. With a mixture of community initiatives, resources and collaborative learning, the Active Inference Institute aims to empower individuals and integrate framework principles into real-world problem solving. Due to its unifying nature, Active Inference has spawned a disparate and broad-reaching Ecosystem of researchers. The Active Inference Institute aims to provide Ecosystem Support through stewarding the information commons and infrastructure scaffolding.

Since its inception, the History of The Institute has been one of an evolving community driven by learners of all stages of experience, and from myriad backgrounds, who have worked individually and in various combinations to expand Active Inference across disciplines. We've understood from the outset that change is inevitable, that sustainability and growth are dependent on a willingness to take chances, and that building trust needs time. The community continues to grow and new members can get involved through participation in one of the many learning opportunities (e.g. Courses, Textbook Group, Production), can advance Active Inference through research and development programs and collaborative research initiatives (e.g. Internship & Fellowship programs), and can engage with the community by contributing to the discourse.

Be part of a globally connected community of Active Inference practitioners and join us as we create this community and expand what is possible for the institute and the many parties, organizations and organisms that can potentially benefit from this work.

As you find your pace and balance, we hope you'll find the Institute isn't just a place of convergence so much as a portal through which new worlds of connections await. We hope to create a space together where we can pursue those opportunities from numerous directions, and where each person feels welcome to enter differently according to the paths which have led them here.

Authors

Authors made various contributions to the Institute & Ecosystem (backend writing document with full trace of edits).

Last updated at the end of 2024.

2025 we will assemble another team of Authors.

Authors

Name	Affiliation	ORCID ID
Alex Vyatkin	Active Inference Institute	0000-0003-1306-4620
Alexandra Mikhailova	Active Inference Institute	0000-0002-8699-7125
Andrea Hiott	Universität Heidelberg	0000-0003-0748-2295
Andrew Pashea	Active Inference Institute	0009-0004-4061-6296
Ben Elers	The Brain Innovation Project	0009-0009-2761-9229
Bert Berkers	Active Inference Institute	0009-0006-6583-6687
Bleu Knight	Active Inference Institute	0000-0002-9894-1989
Chris Fields	Allen Discovery Center, Tufts University	0000-0002-4812-0744
Dan Whittet	AHA Sustainability	0009-0009-5089-0551
Daniel Friedman	Active Inference Institute, COGSEC	0000-0001-6232-9096
Déan Ticklẽs	Active Inference Institute	0000-0003-2213-0773
Fraser Paterson	Active Inference Institute	
Gareth Stubbs	Rabdan Academy	0000-0002-7631-2823
Holly Grimm	Active Inference Institute	0009-0001-6181-2569
Jakub Smekal	Stanford University	0000-0003-4989-4968
Jeremy Cooper	Regent University, School of Psychology & Counseling	0000-0002-2243-1966
John Boik	Active Inference Institute	0000-0003-1289-7997
Libor Burian	Independent, previously CTU FIT	0009-0007-6181-340X
Mahault Albarracin	Universite du Quebec a Montreal	0000-0003-0916-4645
Maria Luiza Iennaco	University of São Paulo, Active Inference Institute	0000-0002-5407-4852
Matthew Brown	ThoughtForge Inc	0000-0002-7552-0989
Mick Thacker	Royal College of Surgeons in Ireland	0000-0002-9034-1137
Peter Gilli	Active Inference Institute	0009-0006-2755-5921
Rafael Kaufmann	Gaia Lab; Primordia Co; Active Inference Institute	0009-0003-8678-4896
RJ Cordes	BlockScience, COGSEC	0000-0002-9913-7159

Ryan Henry	Yale University	0000-0002-0706-6841
Sandeep Ramesh	Panopticon Ventures; Primordia Co.	0009-0006-4976-3326
Scott David	University of Washington: Applied Physics Laboratory	0000-0003-0679-3286
Sebastian Alvarado	Queens College + City University of New York, Biology Department	0000-0001-5866-4043
Zach Baker	University of Colorado	0009-0006-7283-9392

Active Inference

What is Active Inference?

Active Inference is an integrated physics-based approach to modeling cognition and behavior as the active minimization of prediction error. Arising from the empirical study of cognitive systems (those involved in perception and action), Active Inference now is being explored across many Domains of Application.

The formal aspects of the framework describe in mathematical terms the tendency of complex adaptive systems to self-organize as to maintain low-surprise states (formally, through minimization of Variational Free Energy). Active Inference treats this tendency as the basic process, enabling the modeling of perception and behavior in various kinds of cognitive agents, including but not limited to humans.

For those encountering this term for the first time, this can sound technical and obscure, but Active Inference can also be first understood more conceptually and practically as a framing for analysis that is broadly useful towards addressing or gaining perspective in a wide variety fields that formerly seemed unconnected.

At its most basic level, Active Inference can be compared to the guessing game called “20 Questions,” a game in which one person is challenged to guess the identity of an object imagined by another person. In the game, each addition question asked is the “active” part of active inference, and the responses serially constrain the next question as the person guessing. Through this process surprise (bounded by “Variational Free Energy”) due to the differences of an observers “internal model” and the outside reality is reduced offering advantages to the cognitive/behavioral system, whether that system is a cell, an organism, a human organism, or an organization.

For example, Active Inference finds application as diverse as mental health and ecology (see Domains of Application). It is not surprising that Active Inference framing is broadly useful in structuring a deeper understanding of the information flows associated with human cognition and bio-social behaviors in a variety of interaction settings and contexts, since Active Inference first emerged from the study of information flows in nature, where the organizing effects of its thermodynamic underpinnings are expressed most freely.

When Active Inference analysis is directed toward human social and organizational structures and behaviors, it reveals how relevant these bio-physical imperatives are when reflected and expressed in our everyday world. Greater awareness of this foundation, both in individual and organizational contexts, could enhance the overall effectiveness of a variety of information and communication systems and structures, many of which have never enjoyed a “spring cleaning” since their respective historical inceptions.

We could, for example, define Active Inference in everyday language as an approach to understanding our interaction with the world and with those around us, how can we create some sort of model to understand how or why we behave as we behave and then apply this understanding to improve how we self-manage our shared models in a quickly changing reality. More formally, one could rather approach it through mathematics, and explore the foundational aspects of the operation of Active Inference.

The advice to new members of the community who are looking for the best way to begin interacting with the broad range of materials and use cases impacted by active inference analyses is to seek the papers, discussions, materials that present the most familiar vocabulary, narratives and metrics as the starting point, and then to explore from there. To look for terms and keywords across the Active Inference resources that appeal to you most readily and start your journey by following those threads. .

For background readings related to the theoretical basis of Active Inference, see: The 2022 Textbook: “Active Inference: The Free Energy Principle in Mind, Brain, and Behavior” by Thomas Parr, Giovanni Pezzulo, Karl J. Friston (focus of the Textbook Group), “The free energy principle—a precis” by Maxwell Ramstead (October 2023), Spinning Up in Active Inference and the Free Energy Principle" by Jared Tumiel (October 2020), and “Of woodlice and men: A Bayesian account of cognition, life and consciousness” 2018 conversation-style interview with Karl Friston.

For seeing specific applications of Active Inference, see Domains of Application, as well as Institute Projects & Ecosystem Projects. In short — read on!

What are key claims and aspects of Active Inference?

Active Inference is scale-free as both a theoretical framework and a modeling approach. It characterizes all [information processing?] systems [of interacting components?] as behaving in a way that satisfies a single, fundamental goal: every systems acts so as to maintain the distinction between it and its environment. It characterizes all systems as employing the same strategy to achieve this goal: maximizing their ability to predict how their environment will next impact them.

Active Inference thus characterizes all systems - from elementary particles to planetary ecosystems - as agents that both observe (accept input from) and act on (transfer output to) their environments. This information transfer is defined at the agent-environment boundary. For any agent, preserving its distinction from its environment is preserving its boundary, which preserves its identity. The Active-Inference process is, therefore, sometimes referred to as "self-evidencing": any Active Inference agent continually provides its environment with evidence of its existence.

By treating all systems at all scales as agents, Active Inference embraces a minimal, physical definition of "freedom": an Active Inference agent is "free" in the sense that its next action is not causally determined by its environment. One can also put this as: the current state of an Active Inference agent is not causally determined by any, or all, of its environment's past actions on its boundary. Freedom in this sense - freedom from local, causal determinism - is guaranteed to all physical systems by the Conway-Kochen theorems (2006, 2009), which show that local, causal determinism in inconsistent with special relativity, which requires that causal processes take time, and quantum theory, which forbids the state of any system to be fully characterized by a single measurement. Hence Active Inference agents have internal states, and internal processes, that are "protected" from their environments by their boundaries. "Self-evidencing" is, therefore, also "maintaining one's freedom of action".

The generality and action orientation of Active Inference makes it a natural bridge between descriptive approaches to systems, and prescriptive approaches to implementation of artificial intelligence (e.g., machine learning) and design (e.g., user experience, communication, policies, BOLTS, requirements, etc.). Active Inference therefore enables a principled account of composition and decomposition, construction and de-construction, in complex adaptive systems. This generality provides a unified conceptual and pragmatic approach towards establishing a foundation for modeling, designing, and implementing various information processing systems across scales, disciplines, and settings. Active Inference is, therefore, intrinsically a trans-disciplinary framework both for theory and for modeling. As such, it provides a powerful common language into which discipline-specific languages can be translated.

Active Inference leverages Bayesian principles, couching how systems perceive, learn, and act in their environments. It thus treats "knowledge" or "belief" as expectation or prior probability. It treats all agents as Bayesian satisficers, "doing the best they can do" in their environments given how they expect their environments to behave towards them.

Over the last several decades, Active Inference has been attracting increased attention as a quantitative and cognitive framework capable of acting as a common bridge, or Rosetta Stone, among various domains, and is gathering support across Domains of Application. Some citation search measures of this growth in popularity for "Active Inference" and "Free Energy Principle" are shown. Deeper Knowledge Engineering is needed to make stronger inferences about the growth and change of the ideas and their applications, in the research literature and beyond.

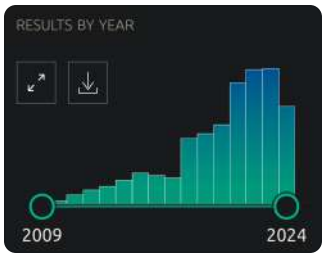

"Active Inference" on PubMed & arXiv

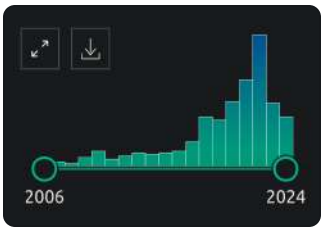

"Free Energy Principle" on PubMed & arXiv

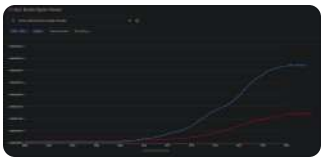

Google Books N-Gram viewer for "Free Energy Principle" and "Active Inference"

What is this excitement and growth about?!

Read on to learn about the The Active Inference Institute, the The Active Inference Ecosystem, and explore the depth and

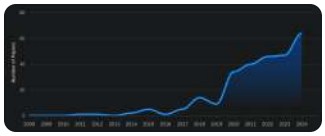

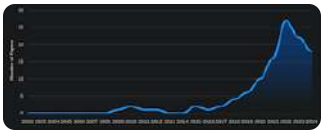

breadth of the work ongoing.

The Active Inference Institute

The Active Inference Institute is a registered non-profit organization (Delaware, USA) which identifies, establishes, scaffolds, and supports the sustainable implementation of:

1. Education and Research services.
    a. We learn and teach Active Inference
    b. We host Institute Programs and Institute Projects
    c. We provide visibility and opportunities for Ecosystem Projects
2. Participation, communication, advisory, governance, and meta-governance affordances within the Institute and The Active Inference Ecosystem
3. Publishing, and licensing protocols that establish Open Source, fair use, and effective dissemination of community products within and beyond the Ecosystem.
4. Ecosystem Support services such as Communications, Grants, Partnership, and operation of cyber and cognitive ˜security systems aligned with our Mission, Vision, Values, and Principles

---

The rest of this section covers:

- History of The Institute since founding in 2021
- Mission, Vision, Values, and Principles
- Focus Areas for the Institute in terms of ongoing challenges (*"where you find the challenge is where the learning/solution is done!"*)
- Directions for the Institute we are taking in light of the focus areas.
- Institute Organization, or morphology, in terms of roles and positions.
- Institute Programs and avenues for participation, such as Volunteer, Internship, Fellows, Philanthropy, Grants, Partnerships.
- Institute Projects hosted by the Organizational Units

History of The Institute

2020

The History of The Institute begins in the co-founder team meeting in 2020 around a common interest in Active Inference. This resulted in productive collaboration and the publication "Active Inference & Behavior Engineering for Teams" in September 2020 (Vyatkin et al. 2020). The group was then known as "Team Comm". Check out
our first livestream, ActInf Livestream #001.1 ~ "Narrative as active inference", on July 28, 2020.

Following the 2020 publication, discussions turned towards exploring approaches that could catalyze the accessibility, rigor, and applicability of Active Inference, and how to merge the developing framework with the Systems Approach and Open Source. Out of these discussions an "Active Inference Lab" (or ActInfLab) was formed and began operations in 2021.

2021

Over the first year of our operations, dozens of individuals from around the world engaged with ActInfLab through various projects such as educational Production , Open Source publishing, collaborative research projects, focused learning groups, Active Inference Journal, and initial developments of the Active Inference Ontology.

Since the first quarter of operations in 2021, the ActInfLab hosted Quarterly Roundtable livestreams for communicating quarterly expectations and results to the community, a tradition that we continue to this day.

2022

Beginning in 2022, a cohort-based Scientific Advisory Board (SAB) was established to connect the ActInfLab to cutting-edge theoretical work as well as various domain-specific applications. As interest in both the ActInfLab's activities and Active Inference itself began to grow, ActInfLab soon emerged as a key facilitating organization in what was then a primarily academic community working on the underlying theory and potential implications for Active Inference.

The first Active Inference textbook comes out in 2022 (Parr, Pezzulo, Friston 2022), and the Institute begins hosting a Textbook Group (ongoing through 7 cohorts in 2024). The Textbook Group is an important ecosystem service, as there are few academic/institutional locations where learners can be supported through the curriculum of the textbook and beyond. Additionally, the Institute has curated and categorized learning materials that learners create while participating in the group, including questions and discourse.

The Institute begins the Internship program to scaffold and support the learning journey of learners. Interns come from different backgrounds — including high school, college, and graduate students on academic tracks, as well as professionals and others outside of academia. Interns, with their mentors, develop a personalized education and research curriculum which lasts months-years.

In mid-2022, ActInfLab made the developmental leap to become The Active Inference Institute, a non-profit organization registered in Delaware, USA with the intention of making its facilitatory role in the community impactful and sustainable. As part of the requirements for a non-profit, we also laid out the Institute Organization, comprised of the Organizational Units: Administrative, EduActive (Education), and ReInference (Research).

At the end of 2022, the Board of Directors has its first meeting. The Board continues to meet on a quarterly basis.

2023

Institute Projects and Institute Programs continue, including the first two full course offerings: Physics course and Social Science course. These courses span months, and include office hours with the lecturer and teaching assistants.

In addition to continuing livestream Production on YouTube (GuestStream, ModelStream, PaperStream, etc), the Institute hosts the popular Active Inference Insights podcast.

During the year, we begin researching and applying for private and government Grants.

2024

Organizationally, the Institute receives official recognition as a 501(c)(3) non-profit organization, supporting our Philanthropy efforts. We were able to achieve this milestone with the *pro bono* support of the Fried Frank law firm.

The largest cohort to date of the Scientific Advisory Board makes many diverse contributions across projects.

The Fellows program begins to highlight and scaffold the work of Ecosystem member. As of November 2024, there are 5 Research Fellows have joined. Fellows represent members of the Ecosystem who have contributed substantially to the ecosystem through publications and presentations.

To meet the needs of trainees and Interns for one-on-one guidance with projects, we introduced the Mentorship program. Members of the Scientific Advisory Board and select other individuals, volunteer to mentor and connect with individual trainees.

Following the Quantum Active Inference Prepare-Measure cycle described by Chris Fields in the 2023 Physics course, we implemented a "Prepare and Measure" system for Institute Projects and Ecosystem Projects. Prepare and Measure allows people to set goals and report back when they have reached them. Project ~ Measurement can be provided by anyone about different Domains of Application for Active Inference. In contrast Project ~ Preparation describes what someone is preparing to do, whether they are just letting us know,

These always-open reporting systems are used to gauge the ongoing projects and work done by community members, and provide visibility to these updates in the newsletter.

Work during this year remains all-volunteer. Philanthropy support begins to come in, supporting some operational software costs. We applied for several Grants (such as FarmWorks and related to AI safety with the RxInfer.jl Learning Group).

The Authors collaborated on this Institute & Ecosystem leading up to the 4th Applied Active Inference Symposium on November 13th, 2024.

2025

- See 2025 for more information on the ongoing year!

2025

January

- Institute Programs and Activities begin for the year
- Theoretical Neurobiology (TNB) Group is now hosted at the Institute.

February

- New Partnership with @Lazy Dynamics, collaborating on the RxInfer.jl Learning Group and elsewhere.

March

- March 28 — 2025 Quarterly Roundtable #1

April

May

June

- June 27 — 2025 Quarterly Roundtable #2

July

- Summer break!
- Complete Project ~ Measurement to have Projects stay in active state.

August

September

- September 26 — 2025 Quarterly Roundtable #3

October

November

- November 12-14th — 5th Applied Active Inference Symposium
- Recruitment for next year Board of Directors and Scientific Advisory Board,
- Projects mostly finish up after the Applied Active Inference Symposium

December

- December 19 — 2025 Quarterly Roundtable #4
- We continue to review Strategy, Mission, Vision, Values, and Principles, Ecosystem Support, Philanthropy, Grants, and more.
- We select the Board of Directors for 2026.

- We have applications open for the Board of Directors for 2026

2026

2026 as it happens

January

- Institute Programs and Activities begin!

February

March

April

May

June

July

August

September

October

November

December

Mission, Vision, Values, and Principles

Our Mission

To support the accessibility, rigor, and applicability of Active Inference.

Act. Infer. Serve.

The formal mission

The formal mission statement of the Institute only scratches the surface of the goals and aspirations of its members and the many parties in its broad ecosystem.

This is screenshot/text from our Form 1023 (this is from the IRS 501(c)(3) status application), submitted in 2023.

The formal mission of the Institute, seen in the screenshot to the left, is:

> *Active Inference Institute, Inc. (the Institute) is dedicated to developing, supporting, and promoting open science and integrative frameworks such as active inference. In furtherance of its mission, the Institute will conduct the following activities: (1) education, (2) research, (3) grantmaking, and (4) administration.*

EduActive (Education)

ReInference (Research)

Grants & Philanthropy

Administrative & Institute Organization

Our Vision

The Active Inference Institute serves as a scaffold for stabilizing and connecting myriad fields around a central tradition and approach of Active Inference.

The Institute aims to make the Active Inference framework and the Ecosystem we serve more accessible, applicable, rigorous, and integrated.

We facilitate educational, theoretical, and applied engagement with Active Inference, promoting awareness of the field within the lay, academic, public-sector, and professional communities.

We envision a future in which the term “Active Inference” is used as widely as “Machine Learning”, as a result of its demonstrated utility and impact in a variety of Domains of Application.

Our Values and Principles

We are committed to fostering a culture of excellence, collaboration, and innovation. Our values and principles serve as the guiding principles that shape our Strategy and define our organization's character.

- **Active Inference and Exploration**. At The Institute, we embrace the principle of Active Inference, Systems Approach, and open-ended exploration as a fundamental driving force. We cultivate a culture of curiosity and continuous learning. Through engaging in endeavors across multiple scales (person, project, Institute, Ecosystem), we enrich our understanding and make relevant contributions to our niche.

- **Integrity**. We strive to uphold and promote honesty, accountability, professionalism, as well as responsible conduct in research, education, and facilitation among members of The Institute, Ecosystem, and communities we serve. We foster diversity, respect, and global inclusion through community engagement. We treat differences in perspective and understanding as a wellspring of valuable creative and productive potential, driving breakthroughs and strengthening collaborative research outcomes.
- **Towards Institute- and Ecosystem-Scale Generative Modeling**: At The Institute, we aim to use model-based approaches towards organizational design and operation. Informally we use the Active Inference Ontology where possible, and looking towards more sophisticated computational modeling in the future. We look support shared informational niches for different scales, spanning the Ecosystem, Institute, Organizational Units, and Projects. We continuously develop and refine hierarchical models, drawing on sensory information, exploiting data, and gathering feedback. Our dynamic self-modeling enables efficient resource allocation.
- **Anticipatory Behavior**: The Institute's commitment to anticipatory behavior equips us to excel in uncertain environments. Leveraging our internal models, we generate predictions at various scales and time horizons, empowering us to take initiative and adapt our policies accordingly. This forward-thinking approach enables us to plan strategically and make informed decisions, thus remaining at the forefront of our fields.
- **Continuous Development**: Embodying the ideas of open-endedness and techno-evolution, we wholeheartedly embrace the principle of continuous development at The Institute. Recognizing the dynamic nature of our environment and the constant advancements in science and technology, we continually evolve our internal models and approaches. This perpetual learning and evolution enable us to remain adaptive and at the cutting edge of our fields, driving impactful research that contributes significantly to the scientific community. In the spirit of action and perception, we encourage learners to produce and share artifacts, then receive feedback: informally and formally (through Prepare and Measure).
- **Participatory Engagement**: At the Institute, we encourage collaborative active learning through artifacts. Digital, stigmergic modifications of our online environment are the central method for engagement. Projects are enacted through preparation and periodic measurements that trace development through time. We support the accessibility, applicability, and inclusivity of Active Inference by seeking **Open Source** (and related: Open Science, DeSci) approaches where possible.

Code of Conduct

Code of Conduct v1, adopted 11-20-2025

1. **Context**
   - Understand the All · Mission, Vision, Values, and Principles and All · History of The Institute, and take this into thoughtful account with your engagement (e.g. in terms of what is relevant to do for or with the Institute, especially for those with informal or formal responsibilities to an organizational mission).
   - Understand the open source context of the Institute, and take it into account when determining how you will make contributions to Institute projects and share information about your own projects (e.g. by not sharing proprietary information).
2. **Respect and Integrity:**
   - Treat all members of the community with respect, dignity, and professionalism.
   - Conduct work and communication with honesty, integrity, and transparency.
   - Foster an engaging environment that values multiple perspectives, backgrounds, and identities and ensures opportunity for all.
3. **Collaboration and Collegiality:**
   - Promote collaboration, teamwork, and mutual support among researchers, learners, and staff.
   - Foster a culture of open communication, constructive feedback, and academic freedom, where ideas can be freely exchanged and challenged in a collaborative manner.
4. **Safety and Well-being:**
   - Prioritize the safety, health, and well-being of all members of the community.
   - Provide a supportive environment that promotes mental and physical health, work-life balance, and personal development.
5. **Epistemic Norms:**
   - Uphold high standards of academic integrity and intellectual property rights in all research activities.
   - Avoid plagiarism, data fabrication, falsification, and other forms of misconduct, ensuring the credibility and reliability of the scientific product and process.
6. **Professional Conduct:**
   - Conduct research in accordance with applicable laws, regulations, and institutional policies (recognizing the global context of Institute participation).
   - Maintain professional conduct in interactions with colleagues, collaborators, funders, and the broader community, avoiding conflicts of interest and maintaining confidentiality where required.
7. **Accountability and Responsibility:**
   - Take responsibility for one's actions and decisions, acknowledging and addressing any mistakes or errors.
   - Hold oneself and others accountable for upholding the principles of this code of conduct, and report any violations or concerns through email to blanket@activeinference.institute.
8. **Continuous Learning and Improvement:**

- Commit to continuous learning, professional development, and the pursuit of excellence in research, teaching, and service.
- Embrace feedback, reflect on experiences, and adapt practices to contribute positively to the advancement of knowledge of adjacent ecosystems and communities affiliated with the Institute.

By adhering to this code of conduct, members of our academic and research driven institution contribute to a culture of excellence, integrity, and collaboration, fostering a positive and inclusive environment for the pursuit of knowledge and innovation.

Focus Areas for the Institute

Below are some Focus Areas for the Institute, and how those Focus Areas are addressed by Directions for the Institute.

The Focus Areas were developed from feedback from participants, and presented here as a part of the overall milestones/snapshot.

Focus Areas

Focus Area	Area Description (why is it challenging, what are the risks?	Related Directions & Steps
Research Advancement and Cross-disciplinary Expansion	Bridging diverse disciplines and translating Active Inference concepts across fields is complex. Without this, we risk siloed knowledge, missed opportunities for innovation, and limited real-world impact of Active Inference principles.	Research Advancement Cross-disciplinary Expansion
Educational Outreach and Resource Development	Active Inference involves abstract concepts and mathematical formalisms, making it difficult for newcomers to engage. Failure to address this could result in a limited pool of practitioners and researchers, slowing the field's growth and application.	Educational Outreach
Software Development and Practical Applications	Developing user-friendly, robust software tools for Active Inference is technically challenging. Without accessible tools, we risk limiting practical implementations and real-world testing of Active Inference models.	Software Development Practical Application
Community Growth and Engagement	Maintaining a cohesive, productive community across diverse backgrounds and interests is complex. Failing to do so could lead to fragmentation, reduced collaboration, and slower progress in advancing Active Inference.	Community Growth Public Engagement
Public Engagement and Ethical Considerations	Translating complex Active Inference concepts for broader public understanding while addressing ethical implications is challenging. Without this, we risk public misunderstanding, potential misuse of the framework, and missed opportunities for societal impact.	Public Engagement

Quality, Performance, and Growth Evaluation

The Institute intends to evaluate quality, performance, and growth within community development at three scales, listed below, based on best practices within the Open Source community and adapted for our use-cases which include software, videos, and other products.

Participant scale

Evaluation at the level of individuals, with consideration for a plurality of individual priors (i.e., diversity in perspective, experience, culture, language, preferences, discipline, and level of expertise) and a focus on accessibility and onboarding. Objectives include quality of participant and user experience, plurality of educational mediums and formats (i.e., accessibility), networking and collaboration opportunities, and professional development. Pending grant or donor funding, The Institute will work with user experience, communications, and requirements engineering professionals to improve current and establish new feedback mechanisms and implement best practices for aforementioned evaluations. The following tools serve as a basis for evaluation:

- Individual feedback forms and surveys
- Participation (e.g., number of projects completed and contributed to)
- Continuing Professional Development (e.g., courses completed, certifications)

Institute scale

Evaluation at the level of The Institute will consider various areas such as sustainability of personal and collective efforts, support reliability, and user experience quality, and Institute quality control and improvement. Objectives include increasing collaboration opportunities, ensuring consistency and rapid handling of inconsistency in documentation, and supporting and facilitating projects. Specific metrics of quality, performance, and growth at The Institute scale may include:

- Number of participants and commits in open source projects
- Number of responses to our Volunteer and Internship forms
- Number of Newsletter signups
- Statistics on projects facilitated by The Institute (e.g., total completed, ongoing, and dissolved)
- Offered and completed Internships
- Frequency of discovery and resolution of inconsistencies in research, documentation, tools.
- Frequency of discovery and resolution of gaps in implementation (i.e., frequently questioned answers and frequently asked questions)
- Number of facilitators, stewards, and volunteers and related turnover and activity
- Aggregation of individual feedback forms and surveys

Ecosystem scale

Evaluation at the level of the Ecosystem and community scale with consideration for impact and relationship management, and a focus on impact. Objectives include minimizing turnover rate in educational courses, increasing the number of participants, and maintaining and adding partnerships. Metrics of quality, performance, and growth at the community scale may include:

- Quality and quantity of Ecosystem Projects, connected with The Active Inference Institute or not.
- Frequency and number of edits and engagements with Coda pages
- Number of participants in Discord General Channel
- Number of participants contributing to facilitated projects

- Turnover rate in engagement and participation (e.g., direct participant engagement with Institute releases and material, and annual involvement in collaborative activities)
- Number of individuals enrolled in educational courses
- Turnover and completion rate in educational courses
- Turnover rate in partnerships (e.g., research and education partnership decisions to renew, maintain, or dissolve)
- Social media analytics (e.g., views, watch time, audience diversity)

Communications

Internal Communications Plan

Institute participants, Officers, Volunteers, and other roles communicate with one another and with members of the community as follows:

- Email serves as the primary means of communication for internal announcements, updates, sharing important documents, and any other professional communications where record keeping is of interest.
- Regular Synchronous Officer Meetings are held to keep communication lines open, address questions, and discuss progress on projects. The Scientific Advisory Board meets regularly in an open discussion format. The Board of Directors meets quarterly to respond to the quarterly roundtable update and address any other issues or concerns.
- Shared Calendars are used to schedule meetings, appointments, and events, ensuring everyone is aware of each other's availability.
- The Institute-operated Discord Server is the primary location for asynchronous discussion and synchronous project meetings. Currently there are over 1000 people in the server, and we strive to keep it an accessible entry point for learning and applying Active Inference.

Organizational Communication

The Institute communicates with potential partners, sponsors, and relevant constituencies through channels including:

- Livestreams and Production provide exciting avenues for live community engagement.
- Content Announcements via X @inferenceactive, Discord, Facebook, Newsletter, LinkedIn, Bluesky @activeinference.bsky.social
    - Measurements come from The Active Inference Ecosystem and reflect: completed projects, recent publications, collaboration and other project opportunities, new releases of educational materials and tools, etc

Target Audiences

- Curious and exploratory learners from all backgrounds and levels of familiar with different subjects/skills.
- Professionals and Academics: Individuals with an interest in cognitive science, machine learning, philosophy, physics, linguistics, computer science, and related areas.
- Potential Partners: Government agencies, funding organizations, academic institutions, and other research-focused organizations.
- Active Inference Community: Researchers, academics and professionals who use and reference Active Inference and related approaches in their daily work.
- Broader Scientific Community: Researchers, academics, and professionals in compatible fields.
- Social Change Organisations: International Organisations, NGOs, civil society
- General Public: Individuals who may have a personal interest in cognitive science, machine learning, philosophy, physics, linguistics, computer science, and related areas.
- Research and Educational: Universities and academic institutions.
- Trade Associations and Think Tanks. Organizations which perform research about future industry trends, in addition to other communities of practice.
- Corporate: Companies with employees who would benefit from knowing Active Inference related approaches to business organization and operations.
- Government: Government agencies and funding vehicles.

- Private Donors: Individuals who understand the value and potential impact of this community of practice and its subject matter, and would be willing to help support it.
- Social Change Organisations: Taking basic underlying concepts and translating them into non-technical language and frameworks for organisations involved in change around large scale social issues (e.g., climate change, peace building)

Approach

The goal of our organizational communications plan is to provide the foundation for sustainable and accessible funding, and to work toward making Active Inference a household term, used as widely as “Machine Learning”, reflecting its demonstrable utility and impact in implementation. An ideal next step toward this goal is the professionalization of Active Inference core competencies and techniques and related competency and qualification standards.

Information Management

The Institute hosts and disseminates information using Coda, YouTube, Discord, Github, Newsletter and other platforms as needed. This stack of platforms streamlines specific levels of access to shared resources, and enhances overall productivity within the organization. We aim to ensure that participants are aware of the platforms being used and understand their purposes and functionalities. We regularly evaluate, communicate, and reinforce best practices for information storage, access, and organization. We implement security measures, such as strong passwords, 2-factor authentication, and appropriate access permission in order to protect sensitive information. We back up important data regularly to prevent loss due to technical issues or accidental deletion. We conduct periodic reviews and audits of the information storage systems to identify areas for improvement and optimization. The specific use of each platform is described below.

YouTube (Live Streaming and Video Hosting)

YouTube is the primary platform for storing audiovisual content created for and by The Institute. Our designated YouTube Channel holds distinct playlists for courses, live streams, symposia, and other content that we host. We share and embed links within internal and external communication channels to provide easy access to relevant content. The content on YouTube is also backed up in a personal cloud storage service as well as in offline hard drives.

Discord (Forum and Instant Messaging)

Discord is our primary platform for engaging with the Active Inference Ecosystem and broader community. We use Discord for real-time communication, informal discussions, and team collaboration. Dedicated channels are used within Discord to categorize discussions based on topics or projects. Participants are encouraged to share relevant files, documents, or links within Discord channels, fostering easy access to shared resources. We regularly monitor and moderate Discord channels to maintain professionalism, and eagerly look to improve our protocols and guidelines here and elsewhere.

Discord

Join the Discord: https://discord.activeinference.institute/

The Active Inference Institute (AII) maintains a Discord server as its primary communication hub where all meetings, discussions, and collaborative activities take place. This digital workspace serves as the central nexus for the institute's diverse community of researchers, practitioners, and enthusiasts interested in active inference.

Server Structure

Main Categories

- Research and education activities
- Project coordination
- Community discussions
- Voice chat rooms for meetings and livestreams

Key Features

The Discord server facilitates:

- Live voice meetings and discussions
- Project collaboration and coordination
- Access to educational resources
- Community engagement and networking

Participation

As with the Institute overall, the Discord server welcomes participants from:

- All backgrounds and experience levels
- Different time zones
- Various levels of familiarity with Active Inference

The server can be accessed through the link https://discord.activeinference.institute/ . It serves as the primary venue for all institute meetings and collaborative activities, making it an essential platform for anyone interested in engaging with the Active Inference community.

Coda

Essentially all Institute Projects use Coda as a document system.

Clicking through links and documentation of Institute Projects you will find many examples of links within and across documents — this Institute & Ecosystem was written collaboratively in Coda, and then exported for snapshot (whereas in 2023 version 1 we used a Google Document linear manuscript co-editing style).

Coda is the primary platform for knowledge and project management at The Institute, Ecosystem, community, and individual scale. It organizes all information and content related to each project (or sub-project). Coda is version-controlled and access-restricted, ensuring that all of our data is protected against accidental deletion and inappropriate user access. We use Coda for storing and organizing important documents, such as policies, procedures, project plans, and meeting notes.

We follow best practices for Coda, including: (1) creating dedicated Coda “documents”, or work areas, for different departments or projects to ensure easy access and organization of relevant information, (2) implementing a clear folder and file structure within Coda to maintain document organization and version control, (3) archiving unnecessary and irrelevant pages, files, and folders, and (4) granting appropriate access permissions to users, allowing them to view, edit, or comment on documents as required.

With adequate future support, Coda will be upgraded to an Enterprise License and consultants will assist in development of templates and low-code applications for streamlining support, records and knowledge management, and project management functions. Further, an Enterprise License will allow for a variety of new mechanisms for user-access control and permissioning, and for tracking of work activity and community engagement with hosted content.

Newsletter

Since the initial activities of the Institute (History of The Institute), we have written a monthly Newsletter.

See the archives https://activeinferenceinstitute.substack.com/

https://newsletter.activeinference.institute/

Directions for the Institute

Directions for the Institute describe ongoing areas of activity and development at the Institute scale.

The following table lists current developmental Directions & Steps and connections with Focus Areas.

Directions & Steps

Direction	Method	Deliverables	Impact / Implication
Research Advancement	Support core Active Inference research; Explore theoretical implications in Philosophy ; Examine group cognition functionality	Research papers; Theoretical frameworks; Computational models	Deepened understanding of Active Inference; New insights at the intersection of multiple fields; Improved models of collective cognition
Software Development	Improve RxInfer.jl visualization capabilities; Enhance PyMDP usability; Develop and curate examples of Domains of Application	Updated Open Source software tools; User-friendly interfaces; Application case studies	More accessible and powerful Active Inference modeling; Increased adoption by researchers and practitioners; Practical demonstrations of Active Inference in action
Educational Outreach	Develop curricula for different languages and contexts; Provide courses and workshops; Increase Communications efforts	Comprehensive curriculum; Industry-focused courses; Educational materials for various skill levels	Wider accessibility of Active Inference concepts; Increased industry engagement; Growth of skilled Active Inference practitioners
Cross-disciplinary Expansion	Seek grants for cross-disciplinary AI research; Pursue features in popular science media; Focus outreach to social sciences	Grants and proposals; Media articles; Collaborative research projects	Broader adoption of Active Inference across disciplines; Increased public awareness; New applications in social sciences
Community Growth	Facilitate intern-mentor connections; Encourage SAB member interactions; Foster edge interactions within community	Mentorship program; Enhanced community engagement; Collaborative projects	Stronger, more connected Active Inference community; Knowledge transfer between experts and newcomers; Innovative cross-pollination of ideas
Public Engagement	Translate concepts for broader public; Address societal challenges through Active Inference; Provide foundations for trust and ethics in AI	Accessible content; Applied solutions to real-world problems; Ethical guidelines for AI development	Increased public understanding of Active Inference; Real-world impact on societal issues; Responsible AI development informed by Active Inference principles
Practical Application	Develop policy appraisal methodologies; Consider ethical and cognitive security aspects; Research capabilities in various domains	Policy frameworks; Ethical guidelines; Domain-specific applications	Informed decision-making in policy; Enhanced cognitive security measures; Demonstration of Active Inference's versatility across fields

Below, we revisit the Focus Areas for the Institute and outline some Directions for the Institute.

1. Research Advancement and Cross-disciplinary Expansion
    a. Seek Grants for cross-disciplinary research
    b. Support core Active Inference research (ReInference (Research)) and educational (EduActive (Education)) development
    c. Explore implications in philosophy, social sciences, and other Domains of Application
    d. Facilitate collaboration with other cognitive models and research communities

    e. Develop new policy appraisal methodologies with focus on ethical and cognitive security considerations
2. Educational Outreach and Resource Development
    a. Develop a full academic curriculum for interdisciplinary audiences
    b. Create educational resources (Fundamentals of Active Inference and Beyond)
    c. Provide courses on Implementations of Active Inference for industry professionals
    d. Increase learning resources for coding Active Inference agents/simulations
    e. Develop foundations for trust, ethics, and education in the context of rapid AI advancement
3. Software Development and Practical Applications
    a. With Open Source development, Improve RxInfer.jl and PyMDP visualization capabilities and overall usability
    b. Develop real-world Implementations of Active Inference across Domains of Application
    c. Support multi-agent workflows (e.g. using ≈ Active Entity Ontology for Science (AEOS))
    d. Create reliable and accurate models for engineers
4. Community Growth and Engagement
    a. Facilitate Mentorship connections with Internship, Fellows, and Ψ Scientific Advisory Board members
    b. Foster edge interactions within the community and The Active Inference Ecosystem
    c. Implement automated feedback mechanisms
    d. Moderate community discourse to ensure compliance with culture and values
    e. Improve onboarding experience for new users
    f. Increase awareness and involvement from organizations outside the Institute
5. Public Engagement and Knowledge Dissemination
    a. Translate Active Inference concepts for broader public understanding
    b. Develop Communications strategies to disseminate knowledge to general public, and professional across areas
    c. Explore the intersection of ∞ Active Inference with current global issues (social, economic, geopolitical, technological, environmental)
    d. Continue to develop Open Source publishing and licensing support systems for contributors

Strategy

https://www.activeinference.institute/strategy

Active Inference Institute (AII) is on the path of open-endedness.

Our Strategy considers learning and applying Active Inference for changes in the niche over multiple nested scales. Through time we increase the degree of hierarchical organizational complexity to overcome competing interactions and frustrated states.

We engage in policy selection across multiple scales, reducing our uncertainty about realizing our expectations and preferences. We learn, finding epistemic value along the way, while pragmatically ensuring Institute persistence and development.

The three scales that Active Inference Institute modifies and interact with:

- **Participant** as an agent. The Institute provides affordances and updates participants' generative model via niche modification and offering of affordances.
- **Institute** as the agent. This is where we engage in Institute-level policies selection and evolve our shared generative model.
- The Active Inference Ecosystem as our epistemic niche.

Institute Organization

The Organizational Units host Institute Programs and Institute Projects

The below shows the overall Institute Organizational morphology, in terms of internal structure and engagement interface with The Active Inference Ecosystem.

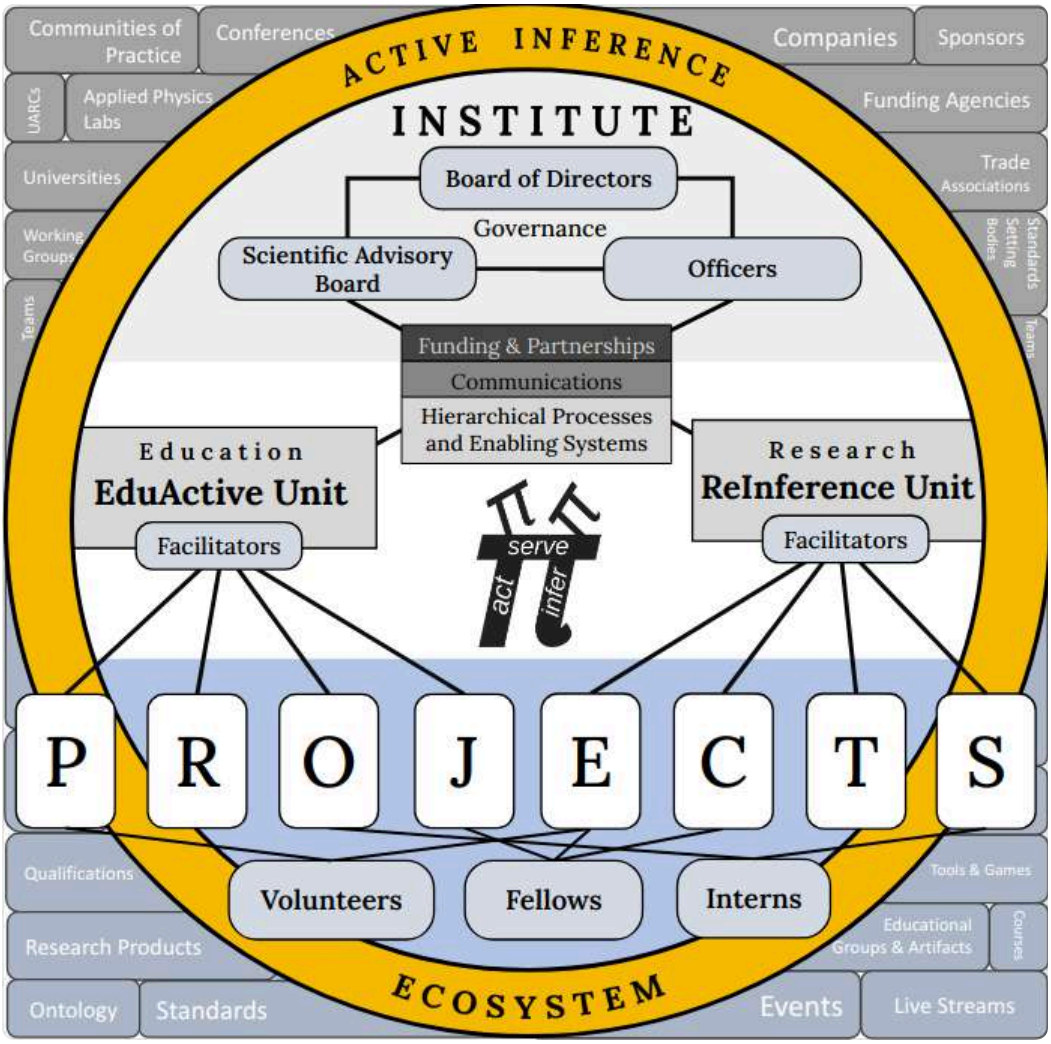

Scientific Advisory Board

The Scientific Advisory Board (SAB) comprises external experts in Active Inference and related research areas who provide guidance, review grant proposals, and offer advice on scientific integrity. The first SAB was active during 2022, and the fourth SAB cohort is active during 🗓️ 2025.

SAB participants offer expertise, advice, guidance, and recommendations to the Institute. They draw on their experience from academia, private business, the public sector, not-for-profit organizations, and beyond. The Scientific Advisory Board acts in an 'advisory capacity' and is not a managing board.

Complete this form to be considered to join a future SAB cohort.

2025 Scientific Advisory Board

- Mahault Albarracin
- Bradly Alicea
- Sebastian Alvarado
- John Boik
- Matt Brown
- John Cook
- Scott David
- Shanna Dobson
- Shady El Damaty
- Jeff Emmett
- Chris Fields
- Karl Friston
- Holly Grimm
- Avel GUÉNIN—CARLUT
- Elliott Hauser
- Andrea Hiott
- Ana Magdalena Hurtado
- Susan Keen
- Thomas Kehler
- Magnus Koudahl
- Michael Lennon
- Héctor Manrique
- George Mobus
- Haris Neophytou
- Arun Niranjan
- Alexander Ororbia
- Chokha Palayamkottai
- Andrew Pashea
- Candice Pattisapu
- Andrew Penland
- Sandeep Ramesh
- Ali Rahmjoo
- Maxwell J. D. Ramstead
- Adeel Razi
- Jeffrey Samuel Schulman Jr.
- Cory Slater
- Jakub Smekal
- Ian Tennant
- Mick Thacker
- Shingai Thornton
- Mark Wilcox
- Michael Zargham

2024 Scientific Advisory Board

Mahault Albarracin, Bradly Alicea, Sebastian Alvarado, John Boik, Matt Brown, John Cook, Scott David, Renée Davis, Shanna Dobson, Shady El Damaty, Jeff Emmett, Chris Fields, Karl Friston, Holly Grimm, Avel GUÉNIN—CARLUT, Sarah Hamburg, Susan Hasty, Conor Heins, Andrea Hiott, Susan Keen, Thomas Kehler, Héctor Manrique, Alexandra Mikhailova, Haris Neophytou, Alexander Ororbia, Sandeep Ramesh, Maxwell J. D. Ramstead, Adeel Razi, Manuel Razo-Mejia, Jakub Smekal, Ian Tennant, Mick Thacker, Shingai Thornton, Mark Wilcox, Michael Zargham

---

2023 Scientific Advisory Board

Bradly Alicea, John Boik, Matt Brown, Scott David, Shady El Damaty, Jeff Emmett, Chris Fields, Karl Friston, Holly Grimm, Sarah Hamburg, Victor Kariuki, Anatoly Levenchuk, Maxwell J. D. Ramstead, Adeel Razi, Michael Zargham

---

2022 Scientific Advisory Board

Bradly Alicea, John Boik, Matt Brown, John Clippinger, Scott David, Jeff Emmett, Chris Fields, Karl Friston, Rafael Kaufmann, Anatoly Levenchuk, Rosalyn Moran, Elba Serrano, Charel van Hoof, Tim Verbelen, Swan Webb, Michael Zargham

Board of Directors

The Board of Directors has been in operation since the end of 2022. The Board of Directors is composed of individuals with expertise in ∞ Active Inference, non-profit governance, fundraising, and various other domains. They meet quarterly and are responsible for setting the organization's Strategy, providing oversight, and ensuring compliance.

The Board of Directors currently consists of:

- John Clippinger (2022-ongoing) — “I want to bring Active Inference into a broad range of applications, specifically into a new model of the firm, markets and finance.”
- Bleu Knight (2022-ongoing) — “I ensure that our actions align with our values and strategic objectives, thus generating the sensations we prefer.”
- Alex Vyatkin (2025-ongoing) — "I want to establish organizational development principles that will strengthen the foundation for the next techno-evolutionary leap."
- Daniel Friedman (2022-ongoing) — “I expect and prefer to integrate the Institute’s daily operations with our broader vision.”
- Mike Smith (2022-ongoing) — “I contribute to strategies for service and education, and facilitate epistemic foraging with active inference in commercial applications.”
- Vladimir Baulin (2025-ongoing) — “Bridging scientific research with real-world implementation strategies”

Previous Board of Directors members:

- Dean Tickles (2022-2024) — “I see my role as a supplier of blind spot remover and a suggester of “Escape Room” strategies as we open up active inferring.”
- Rafael Kaufmann (2022-2024) — “I build adaptive sociotechnical systems that help human collectives, from teams to civilizations.”

Apply for the Board of Directors by completing this form (selected annually)

Officers

- Daniel Friedman (President and Treasurer, 2022-Ongoing)
  - As President, responsible for overall leadership, direction, and Institute Programs.
  - As Treasurer, responsible for managing the financial activities of the Institute, such as Philanthropy and Grants.
- Alexandra Mikhailova (Vice-President and Secretary, 2025-Ongoing)
  - As Vice-President, shares the responsibilities and activities of the President.
  - As Secretary, provides logistical support for Institute activities, with a focus on knowledge engineering.

Prior Officers

- Alexander Vyatkin (Vice President, 2022-2024)
- Bleu Knight (Secretary, 2022-2024)

Apply to be an Officer by completing this form.

Members

The Members of The Active Inference Institute are Virginia Bleu Knight, Daniel Ari Friedman, and Karl John Friston.

Organizational Units

The 🏘️ Organizational Units of the Institute describe the main concentrations or nestings of

---

⚙️ **Administrative for organizational and operational work**

EduActive (Education) for inquiry and learning

🔍 **ReInference (Research) for research and development**

Administrative

The Institute Administrative Unit performs various support tasks within The Active Inference Institute and the wider The Active Inference Ecosystem, such as project coordination, record keeping, graphic design, Grants, project facilitation, preparation, Communications and compliance, and other activities.

Administrative activities contribute to the development of core infrastructure to provide such support and automate or systematize and standardize tasks, and become the organizational umbrella for financial, human resources, Internship, Volunteer, security, community moderation and management, and related activities and organizational components. These tasks are currently assumed by The Institute's Officers, who will continue to provide oversight as the unit develops to include more contributors.

EduActive (Education)

The Institute's Education Unit is named "EduActive" to highlight the active element of education.

Projects of EduActive (Education) include:

- Active Inference Journal
- Active Inference Ontology
- Active Entity Ontology for Science (AEOS)
- Applied Active Inference Symposium
- Educational Standards & Qualifications
- Fundamentals of Active Inference
- Physics course
- Social Science course
- Textbook Group
- Production

Active Projects ~ EduActive

Not synced yet

Type of project	Project title	Examples of work	Coda	Mission objectives
▼ Ecosystem 7	Project Development for "Solving the Tower of Babel Problem: UniFysica Philo-sophia"			To outline, draft, a collection of papers with th of Communication for Sapiens' Shared Meanin Blombos to Friston and Fields"
	Numinia	The Adventure of Curiosity (oncyber.io) https://app.charmverse.io/numinia		First mission would be to make sure we are im in the game properly and is well explained, an ensure that the design of the incentives aligne and the AII.
	Neurodivergent Learning Sessions			Neurodivergent learning is focused on outreac geared towards those who struggle with stanc environments when it comes to public and hig as a number of people with neurological condi spectrum disorder can struggle in varying way the right environment in which information is p manner which is coherent.
	Active Inference Cycle Book for Self-Knowing			Perform a meta analysis of the "wellness" spa inference highlighting the most impactful poin an easily digestible format. Use this work to k collaboration and contribution to the larger AII
	The Unordinary Bible Study (abbreviated as TUBS)			Hosting once a month sessions that focus on verses but not spending too much time diggin to focusing on inter-faith and contemporary p

	Creativity and creators under the light of the Free Energy Principle			Design and run experiments to answer the key
	The Three Mosqueteers			Create a livestream aimed at disseminating sc without a scientific background to adopt a mo information they receive.
▼ Institute 7	Active Inference Ontology	Public Active Inference Ontology website		Maintain, improve, elaborate, extend, translate apply the Active Inference Ontology as core in Inference Institute & Ecosystem.
	Audio-Visual Production	Table with all livestreams and videos from 2020-Ongoing		Produce accessible, rigorous, informative (epis (pragmatic value) audio-visual content, for exa Podcasts, and other formats.
	Textbook Group (Parr, Pezzulo, Friston 2022)	5.5 completed cohorts since 2022 (see Coda)		Improve the accessibility, rigor, applicability, a 2022 Active textbook by Parr, Pezzulo, and Fri
	Active Inference Journal	See the Github		To develop evolving hybrid (AI+people) projec volunteers team
	Course Development	Courses, Obsidian Repository		Develop educational materials and experience Active Inference theory and practice.
	Applied Active Inference Symposium	Applied Active Inference Symposium		Host an annual Symposium to highlight the sta Active Inference.
	Seasonal School	Seasonal School		Develop in-person experiences for education a

RelInference (Research)

The Research unit at the Institute is named “RelInference” to highlight the perspective on scientific research and inquiry more broadly in terms of ∞ Active Inference (Pietarinen & Beni 2021, Balzan et al. 2023).

The Institute’s RelInference unit focuses on Research activities such as: (i) the forming of fit-for-function interdisciplinary research teams, (ii) the development and execution of research proposals and projects aligned with the mission of The Institute and challenges faced by The Institute and Ecosystem at large.

The RelInference unit is committed to hosting and sharing all relevant data, findings, publications, tools, and derivative artifacts under Open Source or similarly accessible licensing agreements wherever practicable and appropriate.

Projects of RelInference (Research) include:

- RxInfer.jl Learning Group
- Active Blockference
- Active InferAnts
- Tech Tree
- Knowledge Engineering
- Generalized Notation Notation
- FarmWorks

Active Projects ~ RelInference

Not synced yet

Type of project	Project title	Examples of work	Coda	Mission objectives
▼ Institute 6	Knowledge Engineering	Public frontend and Literature meta-analysis from end of 2022		This project seeks to alleviate the information b Ecosystem through information curation, organ summaries of institute productions (courses, liv
	Active Blockference	Github, Blog post, video overview from 2022		We are applying Active Inference by building ca generative models.
	RxInfer.jl learning and development group	See project overview		Learn and apply RxInfer.jl in 2024 — building ou generative modeling.
	FarmWorks	See FarmWorks pageand 2024 publication.		Develop minimal model of personalized agents

	Applied Active Inference Symposium	4 prior Symposia from 2021 through 2024		To have a year-end Symposium, featuring appli the world
	AICACP	AICACP		AICACP is a multi-year initiative designed to re capabilities, alignment, and regulation.
▼ Ecosystem 9	Active Inference Account of Belief Updating in PTSD			Write a theoretical paper in the style of Parr et
	Symbolic cognitive robotics	Most recent paper from 2023, Robotics & Embodied		Explore the joint problem space of “symbolic a “mortal computing”, with an emphasis on unsu Using symbolic processing, build a rudimentar whose behavior fulfills the requirements of Act
	Humanity’s Story of an Uncertain Self	ActInf GuestStream 061.1 ~ Shagor Rahman: "Myth of objectivity and the origin of symbols"		Producing an academic paper or blog that cont simulations, and ultimately a framework that ex humanity’s sociological-narrative framework. S of say, ancient epics, along with an set of econ us to model to simulate, predict, and give mast intractable world of humanity’s cultural niche.

Improving RxInfer.jl's Model Visualization Capabilities	Current work on visualization		Our mission is to equip RxInfer.jl - and its releva model visualization modalities that prove usefu develop RxInfer.jl. To that end, we anticipate measuring the initial reception from RxInfer.jl's core developers: TU/ therefore take the approval of the BIASlab as th
CogNarr Ecosystem: Facilitating Group Cognition at Scale	CogNarr (Cognitive Narrative) Ecosystem: Facilitating Group Cognition at Scale		The initial mission is to advance the CogNarr p into a proof-of-concept demonstration, followe In concept, the CogNarr ecosystem of softwar component of a group's cognitive architecture.
Model-Centric cognition	Wave Hypothesis		Develop the central idea, raise awareness; ass brain-as-computer paradigm is needed. Some The project is a development of existing wave l
The Einstein Model	2023 paper		Bridging Psychoanalysis and Thermodynamics

Institute Programs

The Institute Programs are the specific modes of active participation and engagement (beyond e.g. just watching Production).

For individuals:

- Volunteer positions contribute to the Institute's work in specific ways.
- Internship and Mentorship provide structure for those looking to advance their learning and work.
- Fellows provides the opportunity for committed individuals to be recognized as a leader in Active Inference research and education.

For organizations:

- The Partnership, Philanthropy, and Grants programs all provide channels of support and bi-directional learning with the Institute.

Volunteer

Operationally, all participants of the Institute are volunteers. Volunteers join the Institute by emailing project facilitators, or communications from Discord. Communications such as the website https://www.activeinference.institute/ and Newsletter also contain solicitations for signing up for general Volunteer lists, and specific Institute Projects.

As a community-driven open science organization, there are multiple opportunities for contribution. All backgrounds, time zones, time availability, and levels of familiarity with Active Inference are welcome and encouraged. Volunteers are active learners who want to contribute to ongoing projects at The Institute. Volunteers have the opportunity to engage in and lead a wide array of projects without any constraints on their type or quantity. These projects encompass a range of activities such as study groups, livestreams, marketing initiatives, publications, symposiums, and applied research.

We look to develop and clarify how the Volunteer position will work in 2025. Current thinking is exploring ideas around:

- Official Recognition (via specific affordances and statuses such as: affiliation status for papers and communications, access to code repositories and digital resources, @activeinference.institute email address, letters of recognition, inclusion on Grants, payment via Philanthropy, etc).
- Role of Mentorship in Volunteer efforts and Education more broadly.
- Position and Institute Projects-specific Documentation, regular participation in prepare/measure cycles.
- Stewardship of specific Domains of Application paper sections
- Sampling among expertise areas to build Prediction Matter Expertise (PME), not just Subject Matter Expertise (SME). Facilitating discussions and answering basic questions.
- Contributing to Open Source projects.

The volunteer program aims to balance structure with flexibility, providing clear value while maintaining active inference principles in learning and contribution. This framework allows volunteers to grow within the Institute while contributing meaningfully to its mission.

Complete this form to Volunteer at the Institute.

We keep Volunteers posted about affordances for Learning Groups, Projects, Internships, and more.
Let us know in the final question response, or via email to blanket@activeinference.institute if you have any questions.

Internship

The Internship program benefits the intern by increasing their familiarity and expertise with Active Inference and associated areas, as well as offering practical experience with teamwork and project-specific skills. Upon completion, interns will receive acknowledgement and a certification of completion with the duration and focus specified. Letters of recommendation will be granted on a case-by-case basis.

The form to apply for an Internship is at the bottom of this page.

Internship Format

- The Internship program is customized to situation and timing of each person, and ranges in duration from several months-years. The internship activities are streamlined with the individuals other activities and aligned with their preferences. Interns are assigned an organizational point of contact, optional additional mentorship, and have periodic synchronous and asynchronous check-ins.
- Interns are responsible for active and documented participation in the Internships. It is critical that interns are open to adapting the internship plan and actively reduce uncertainties as needed.
- There are two primary components to the Internship:
  - Learning and Updating: self-guided as well as participation in EduActive (Education) projects (e.g. Textbook Group , Production)
  - Research and Development: private projects and/or ReInference (Research) activities.
- Most work will be done on your own time. Most synchronous interactions will be in the context of All group meetings.
- You will be connected with Mentorship support, and have periodic synchronous and asynchronous check-ins.
- There are two primary areas of engagement to the Internship
  - Education: Participation in Learning Groups at EduActive (Education) & development of the Intern's personal learning journey.
  - Research: Participation in ReInference (Research) projects & advancement of personal research programs.

Benefits of the Internship

- For the Intern:
  - Increased familiarity, expertise, and knowledge about Active Inference
  - Practical experience with team working and Project-specific skills.
  - Upon completion of term of internship:
    - Acknowledgement and Certification that one completed the Internship with the duration and focus specified.
    - Availability of research infrastructure in future, consideration for paid positions
    - Possible letter of recommendation (this is up to the person you ask)
- For Active Inference Institute:
  - Increased participation in Projects & Learning Group.
  - Advancement of projects & increased impact/service to The Active Inference Ecosystem.
  - Implementation of our Strategy in terms of participation, engagement, and methodology.

Intern Responsibilities

- Active and documented participation in Learning Groups and Projects.
- Active communication with Learning Groups, Projects, and Mentor.
- Openness to adapting the Internship approach as needed.
- Staying balanced and healthy.

Link to apply for an Internship

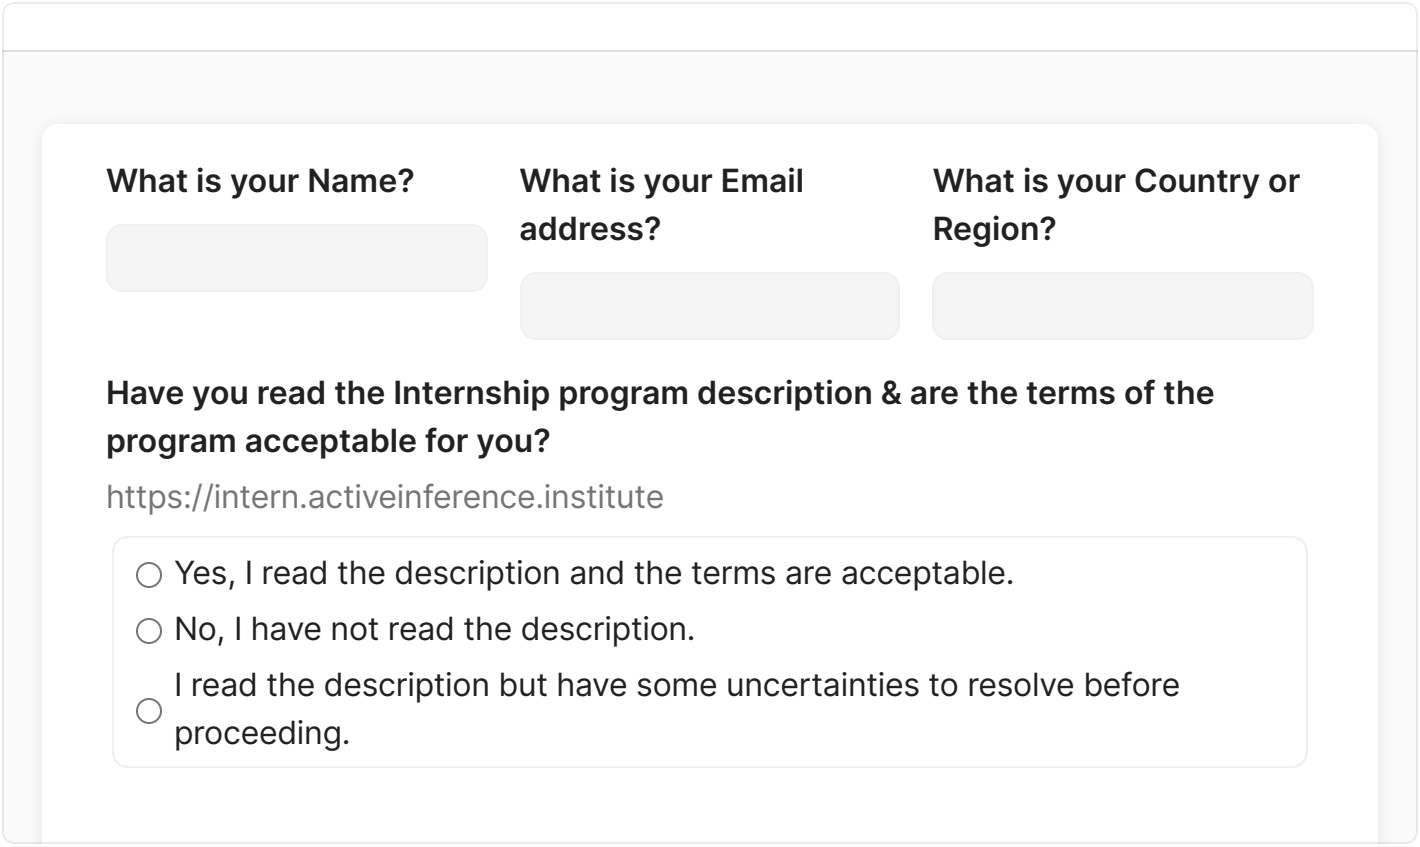

http://intern.activeinference.institute/

Mentorship

The Mentorship program connects Interns with additional one-on-one support from a Mentor (a member of Ψ Scientific Advisory Board who assumes the role of sponsor for authentic learning exposures both in and beyond the Active Inference Institute).

Serving as a Mentor is a way to provide a unique contribution for our community, **through** engaging with the Educational space parameterized by Subject Matter Expertise and Prediction Matter Expertise (see the Education page for further details). This learning condition presents as a co-learning opportunity where the updating of ALL participant generative models, is expected and preferred. Here “what got you to a **new** Know” matters as much as “what you know already.”

Using the Mentorship form, individuals in the Volunteer and Internship program submit a personal statement and request for kind of advising that would best serve their learning trajectory. The Institute works to match up individuals on a rolling basis.

Fellows

Since 2024, the Active Inference Institute has hosted a Fellows program, designed to support and advance research in Active Inference. This program provides a unique opportunity for exceptional individuals of different career stages to join the Institute's vibrant community and contribute to the development of Active Inference and its applications across domains.

Eventually we expect to have multiple kinds of Fellows, focusing on Research, Education, and more. Initially we have begun with Research Fellows, who conduct self-directed, innovative research projects that align with the Institute's research and education missions.

Research Fellows have access to the Institute's facilities, computational resources, a network of leading experts in the field, and the ability to engage with Philanthropy and Grants Institute Programs. Fellows also benefit from professional development opportunities, including Mentorship, training workshops, and support for Open Source practices.

The Research Fellow position is an unpaid non-employee position. The default term for a Research Fellow is 2 years, with the potential to renew. Applications are considered on a rolling basis. Eventually, we look forward to developing our collaborations, philanthropy, and partnerships at the Institute, in order to enable fellowships to be paid financially in some way.

https://fellows.activeinference.institute/

Research Fellows

Research Fellows program at the Active Inference Institute

The Research Fellows program is designed to support and advance research in Active Inference. This program provides a unique opportunity for exceptional researchers of different career stages to join The Active Inference Institute's vibrant community and contribute to the development of Active Inference and Domains of Application.

Research Fellows conduct self-directed, innovative research projects that align with the Institute's research and education missions. Research Fellows have access to the Institute's facilities, computational resources, and a network of leading experts in the field. Fellows will also benefit from professional development opportunities, including mentorship, training workshops, and support for open science practices.

Eventually, we look forward to developing our collaborations, Philanthropy, and Partnerships at the Institute, in order to enable fellowships to be paid financially. For now the Research Fellow position is an unpaid non-employee position.

Join us in pushing the boundaries of Active Inference research and shaping the future of this exciting field! See Research Fellow Application. Applications are always open & will be considered on a rolling basis.The default term for a Research Fellow is 2 years (with possible renewal).

Contact:
For more information about the Active Inference Institute Research Fellows program, please visit fellows.activeinference.institute or contact us at blanket@activeinference.institute .

Current Research Fellows

Not synced yet

Jean-Francois Cloutier

2nd, 3rd Active Inference Symposium

I seek to find out what it takes, at a minimum, for a robot to learn, on its own, how to surviv...

Symbolic Cognitive Robotics

Anna Pereira

Cultivating a grass roots impact project (initially through nonfiction literature) to explore...

John Boik

Livestream #021 series: .01, .02, .03, .04, .1, .2

As an Active Inference Institute Research Fellow, the research program I will pursue is a ...

Cognitive Narrative Ecosystem

David Bloomin

GuestStream 085.1

I am investigating how the principles of Active Inference, combined with social dynamic...

Metta AI

Robert Worden

GuestStream #082 series: .1, .2, .3, .4

I have two main research interests: 3-D spatial cognition, and language.

Wave Hypothesis

Shagor Rahman

GuestStream #061.1

I investigate how morality and symbolic thought co-evolved through what I call the "Myth o...

Myth of Objectivity

Hongju Pae

2025 Symposium Presentation

I develop computational and theoretical frameworks for modeling how artificial agents ...

Sheila Macrine

As an Active Inference Institute Research Fellow, I am extending the theoretical framework ...

Research Fellows ~ Table

Not synced yet

Name	ORCID	Starting date	Image	Livestreams	Overview
Anna Pereira	0009-0008-9049-0707	5/2024			Cultivating a grass roots impact project (initially thr Principles. Active Inference is the key lens that ther enabling humans to live more fulfilling lives, respon mutualistic opportunities for collaboration and seek Say hello, collaborate, or discuss at via anna@activ
David Bloomin		10/2024		GuestStream 085.1	I am investigating how the principles of Active Infer foster cooperation and alignment in multi-agent en mechanism in gridworld simulations. The project air minimize free energy. Through an open-source mo to aligned cooperative intelligence, informing the pa You can follow my progress at http://daveey.github.
Hongju Pae	0000-0002-5174-8858	11/2025		2025 Symposium Presentation	I develop computational and theoretical framework and achieve developmental alignment. My work int identify computable markers of subjective experien simulation prototypes and mathematical models to without external reward shaping. My goal is to esta grounded machine minds. Personal webpage https://www.linkedin.com/in/hjpa Lab webpage https://hjpae.github.io/cear/
Jean-Francois Cloutier	0009-0001-1841-2279	5/2024		2nd, 3rd Active Inference Symposium	I seek to find out what it takes, at a minimum, for a nothing about. My research is the continuation of a and ground my understanding of cognition. Looking for answers has already taken me on an un drawn into Active Inference of course but also Kant collective intelligence, autopoiesis and constraint cl

John Boik	0000-0003-1289-7997	5/2024		Livestream #021 series: .01, .02, .03, .04, .1, .2	As an Active Inference Institute Research Fellow, th book and in two series of concept papers. That pro architectures that are, by design, fit for purpose. The first series describes how the approach can be governance systems), which are viewed as compor The second series describes how the approach car large-group setting.
Robert Worden	0000-0001-7304-2752	10/2024		GuestStream #082 series: .1, .2, .3, .4	I have two main research interests: 3-D spatial cogr All animals need to understand the local 3-D space not by neural computing alone, but using a wave in a novel theory of consciousness – that it arises not using active inference. See *Frontiers* article on the t I also work on language – how it evolved, how we le a demonstration of language learning.
Shagor Rahman	0009-0004-0460-0078	8/2025		GuestStream #061.1	I investigate how morality and symbolic thought co· our capacity to model shared cultural expectations symbolic spaces. This framework offers a strong pe Employing multi-agent active inference simulations explicit moral beliefs form the foundation of our syn and impact psychological well-being through narrat This computational framework helps us understand religious prophets, cultural thought leaders of socia can reshape these symbolic systems. By formalizin offers insights into both human uniqueness and the
Sheila Macrine	0000-0002-8600-0938	11/2025			As an Active Inference Institute Research Fellow, I a ***Embodied Intelligence: Multidisciplinary Perspect*** proposes a novel multi-dimensional taxonomy of ag systems. The first phase of this research investigat and Reflexivity—can be mathematically formalized constructing formal "Agency Profiles," this work aim using precise metrics. For this model, Agency is the precision of intrinsic priors. Intelligence is the 'Engir efficiency that enables agency. This approach prov profiles characterized by high rationality but low au The second phase is dedicated to developing a sub Inference models across diverse scales—ranging fr understanding of how agency emerges, scales, and primary interest lies within the Theoretical Neurobic interdisciplinary collaboration and mathematical de formalizing these concepts, aiming to operationaliz

All information at: fellows.activeinference.institute

Research Fellow Application

Research Fellow Application Package Requirements

1. Confirmation that one has read & agreed to the Research Fellows Terms.
2. Research proposal (1-6 pages, excluding citations), which can describe one or more research projects in detail or outline a research direction in more general terms. The proposal should include:

    Title

    Abstract (300 words or less)

    Research question(s) and objectives

    Significance and impact of the proposed research

    Approach and research methods

    Alignment with the Active Inference Institute's mission

    Anticipated measurements (outcomes and deliverables)

    Timeline and milestones

    People and institutions involved

    Any dependencies or contingencies that might affect progress

    Cited references (not counted in page limit)

    In the research proposal, please clearly address:

    — Motivation for proposed work and for applying for a fellowship

    — Describe previous engagements/interactions with the Institute

    — Describe Institute programs or activities of particular interest and/or that you intend to participate in or facilitate.

3. Curriculum vitae
4. One to Three letters of recommendation (submitted in application packet, or sent separately to blanket@activeinference.institute)
5. Up to five representative publications or products

Application questions and completed packets should be sent to blanket@activeinference.institute with [RESEARCH FELLOWS] in the Subject line. Please include all application components as separate PDF or document files attached to your email.

Research Fellows ~ Application Questions

Question	Answer
I have X educational degree / have no PhD — can I apply?	Yes, you can apply. There is no requirement for a PhD, or any specific degree
I am unemployed / doing Active Inference research on my own time / am employed doing other work / would be a part-time Fellow — can I apply?	Yes, all employment statuses are acceptable in principle, as long as the appl is clear about what current/future obligations are and how one plans to proc
How many Fellows will the Institute accept?	We do not have a set fixed limit as of 2024.
What are the Term limits? If 2 years is default, can we allow a term of less than 2 years?	Yes, we will consider terms of less than 2 years.
Fellowships are eligible for how many renewals?	We do not have a set fixed policy on this as of 2024.
I have some other questions	Application questions and completed packets should be sent to blanket@activeinference.institute with [RESEARCH FELLOWS] in the Subject

Research Fellows Terms

Terms v1 (April 2024)

Terms for Active Inference Institute Nonemployee Research Fellows

1. Code of Conduct

- The Research Fellow (Fellow) agrees to adhere to the highest standards of scientific integrity, professional ethics, and responsible conduct of research.
- This includes honesty, objectivity, fairness, respect, accountability, and transparency in all research activities and interactions.

2. Fellowship Agreement

- The Active Inference Institute (Institute) reserves the right to terminate the fellowship agreement (agreement) at any time.
- The specific terms, conditions, and duration of the agreement will be outlined in the initial offer letter, which must be signed by both parties.
- The Fellow is responsible for ensuring that this agreement is compatible with their other institutional affiliations, contracts, and policies.

3. Intellectual Property (IP)

- Nothing about this agreement - changes the status of IP arising during the program (e.g. the Institute does not claim any access to Fellow's private work, nor does the Fellows participation in Institute products/processes affect the Fellow's or Institute's IP rights), except as noted next.
- The Fellow may be required to sign additional IP agreements or disclosures as outlined in the initial offer letter (depending on the applicant's situation and status).

4. Affiliation and Acknowledgment

- The Fellow is encouraged to list the Institute as a professional affiliation on research outputs, presentations, and professional communications related to their fellowship activities.
- The official title for this affiliation is "Research Fellow, Active Inference Institute".

5. Open Science and Research Productivity

- The Institute encourages and supports Fellows to disseminate their research outputs through open access channels, such as preprint servers, open data repositories, and open source software platforms.
- The Fellow should strive to meet the research objectives outlined in their approved proposal .

6. Reporting and Evaluation

- The Fellow shall provide regular progress reports (e.g. quarterly) to the Institute, describing their research activities, achievements, challenges, and plans.
- The Fellow shall participate in an annual evaluation process, which may include a written self-assessment, an oral presentation, and feedback from Institute mentors and collaborators.
- Satisfactory performance and progress, as determined by the Institute, are required for continuation and possible renewal of the fellowship.

By signing the initial offer letter, the Fellow agrees to abide by the terms of this document.

The Institute reserves the right to modify these terms as needed, with written notice to the Fellow.

Philanthropy

We are thrilled to share that in 2024 the Active Inference Institute has officially been recognized as a 501(c)(3) nonprofit organization by the United States IRS. This significant milestone is the result of several years of dedicated effort.

Your generous support will help sustain and extend the work of the Institute and ecosystem, supporting the accessibility, rigor, and applicability of Active Inference. With our 501(c)(3) organizational status, your donation may be tax-exempt.

If you value our mission and wish to contribute, please consider making a donation at:

---

donate.activeinference.institute/

At this time, for any other comments or questions on philanthropy and donations to the Institute, please communicate with blanket@ActiveInference.Institute .

For the organizational Partnership affordance (which may include structured financial or in-kind donations), see Partnership.

Thank you for your continued attention and consideration. We look forward to all the next moves we'll take together.

Grants

The Active Inference Institute is primarily a volunteer organization. We seek sustainable approaches to scaffold our work through Grants and Philanthropy.

As part of our commitment to Open Source, submitted grants are made public whenever possible, by uploading to a preprint server (such as Zenodo) as a publication. In this way, we leave a stigmergic trace on the ecosystem reflecting our plans, and history of assembling teams to tackle areas of research and Education.

Current Grants

- In June 2025 a team of researchers in the AICACP were awarded $270,000 through the Institute.

Previous Grants we have applied for:

- In 2022 we applied for, and did not receive, "Systems Modeling and Cognitive Audits for Hypercert Ecosystems". The application was published on Zenodo.
- In 2023, we applied for, and did not receive, an NSF Pathways to Enable Open-Source Ecosystems (POSE) grant. Along the way, we collaboratively wrote the 2023 paper "The Active Inference Institute and Active Inference Ecosystem" — the structure and text of which, was the initial conditions (prior) for Institute & Ecosystem document.
- In 2024, the RxInfer.jl Learning Group applied for, and did not receive, an FLI grant "FarmWorks: Decentralized AI Agents for Personalized Solutions" (Zenodo link).
- In 2024, the RxInfer.jl Learning Group applied for, and did not receive, a Foresight Institute grant "VILLAGE (Validating Inference for Large-scale Agent Governance Ecosystems)" (document link).
- In 2025, a team applied for, and did not receive, a Dana Frontiers grant "Increasing the Accessibility and Applicability of Active Inference: Generative Playbooks and Open-Source Summer School Curriculum Development" (Zenodo link).
- In 2025, a team applied for, and did not get accepted into, a Google.org accelerator for generative AI. Our team's collaboration did lead to product development and two papers: ResNei: Solution Design Document and The Discovery Engine: A Framework for AI-Driven Synthesis and Navigation of Scientific Knowledge Landscapes.

We have written multiple letters of support, collaboration, and Partnership for others in their applications and Grants.

Partnership

Organizational Partnership Program at the Active Inference Institute

About the Partnership Program

The Active Inference Institute Partnership Program fosters collaboration with organizations aligned with its mission to advance the understanding, application, and accessibility of Active Inference.

By creating mutually beneficial relationships, the program enables partners to contribute to and benefit from the Institute's research ecosystem, which spans diverse Domains of Application such as cognitive science, artificial intelligence, education, and organizational dynamics.

Through partnerships, the Institute supports the development of Projects, educational materials, frameworks, and tools that leverage Active Inference principles to address real-world challenges. Partners gain access to a vibrant network of researchers, developers, and thought leaders while contributing to the growth of the global Active Inference community.

Learn more:

- See Current Partners and details of our engagements
- Partnership Application and Partnership Terms

What the Institute can Provide

- **Recognition and Access**: Public acknowledgment on the Institute's website and materials, along with access to its network of researchers, contributors, and interns.
- **Collaboration Opportunities**: Regular meetings with Institute personnel to co-develop programs, projects, and initiatives in areas of shared interest.
- **Support for Specific Programs**: The Institute facilitates collaborations on targeted initiatives such as educational Courses, livestream Productions, Internship, Fellows, Applied Active Inference Symposium themes, and region-specific work.
- **Guidance and Expertise**: Strategic insights into applying Active Inference principles in organizational or scientific contexts through workshops, training sessions, and mentorship.
- **Flexibility in Engagement**: Tailored levels of involvement based on partner preferences, ranging from casual participation to formalized collaboration.

What the Partner may Contribute

- **Financial or In-Kind Contributions**: Provide monetary support and/or resources (e.g., compute power, datasets, development expertise) to sustain and expand the Institute's work.
- **Alignment with Goals**: Submit an application demonstrating alignment with the Institute's mission of advancing Active Inference and broadening scientific participation.

- **Time and Attention**: Dedicate an agreed-upon level of involvement, from casual engagement in projects to formal facilitation of programs. This can include direct participation on Projects, service on Scientific Advisory Board, mentors for Internship, or other avenues of engagement.
- **Reliable Communication**: Designate a point of contact for regular communication and coordination with the Institute.

Programs and Projects That Can Be Supported

The Active Inference Institute offers opportunities for partners to support a range of operational programs and specific projects. Benefits of supporting specific programs include:

1. Partners can directly contribute to advancing science, education, or applied research in their areas of interest.
2. Financial or in-kind contributions enable impactful initiatives that align with both partner goals and the Institute's mission.
3. Collaboration fosters mutual growth while expanding the reach of Active Inference principles across domains.

By supporting these Programs for Partnership support or projects, partners play a pivotal role in shaping the future of Active Inference research, education, and applications globally. See the table below for some ideas, or reach out if you have other ideas.

Programs for Partnership support

Category	Program/Project	Description	Impact Enabled by Support
▼ **Institute Programs** 4	Internship program	Mentorship-focused development opportunities for emerging talent.	Cultivates future leaders in Active Inference research and applications.
	Research Fellows program	Supports self-directed research projects aligned with Active Inference principles.	Advances cutting-edge studies while fostering professional growth for researchers.
	Grants	Builds administrative capacities for funding proposals.	Enhances the Institute's ability to secure resources for long-term sustainability.
	Open Source and Open Science	Promotes transparency and accessibility through open-source tools and frameworks.	Strengthens the global ecosystem of Active Inference research and applications.
▼ **Institute Projects** 8	Open Source Code Development (e.g., RxInfer.jl)	Develops generative modeling frameworks and tools for Bayesian agents.	Provides foundational resources for researchers and practitioners worldwide.
	Active Inference Journal	Publishes translations, transcripts, papers, and other scholarly outputs.	Increases accessibility to key insights across languages and disciplines.
	Active Inference Ontology	Formalizes Active Inference structures across languages and domains.	Facilitates interdisciplinary collaboration by creating standardized conceptual frameworks.
	Applied Active Inference Symposium	Organizes events focused on practical applications of Active Inference theories.	Encourages knowledge exchange between academia, industry, and broader communities.

Courses	Designs courses that apply Active Inference principles in different settings, such as Physics course and Social Science course.	Expands educational opportunities for audiences globally.
Knowledge Engineering	Indexes literature and resources for improved accessibility.	Reduces research debt while enabling efficient knowledge dissemination across fields.
Audio-Visual Production	Produces engaging content to communicate Active Inference concepts to wider audiences.	Builds public awareness while fostering interdisciplinary dialogue on complex topics.
In-Person Experiences (e.g.,	Creates immersive learning environments for participants worldwide.	Deepens understanding through hands-on engagement with experts in the field.

https://partnerships.activeinference.institute/

Current Partners

Updated: April 2025

Current Institute Partners

Name	Partner Description	Partner Logo	Link	Partnership description
First Principles First	Towards a Science of Mindful Agents, Societies and Observer Languages			The Active Inference Institute has a rich ecosystem of resear developers, and thought leaders that FP1 can draw upon. In t committed to spawning and undertaking projects to expand t and application of active inference.
Numen Games	We are building a Gamify Structural Framework for Organization an Open Metaverse RPG, Numinia is game to work better.	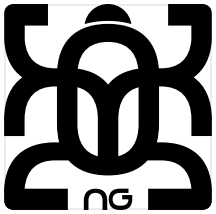		Through our partnership with the Active Inference Institute, N to combine its expertise in gamified, immersive 3D environme Institute's scientific approach to Active Inference. Together, w enhance learning experiences by embedding these scientific into our open metaverse RPG, providing a research-driven fra organizations to innovate and thrive. Our 3D educational gamified experiences will foster interacti real-world organizational settings, while the Institute will guid scientific methodologies that deepen our understanding of c organizational dynamics.
Lazy Dynamics	Lazy Dynamics leads the development of RxInfer-PRO and co-supervises RxInfer (through the open source community ReactiveBayes), frameworks for building Active Inference and Bayesian agents.	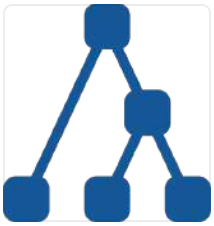		With an official partnership in 2025 between Lazy Dynamics Active Inference Institute, we seek to strengthen interfaces a collaborations related to the development, awareness, and a RxInfer and generative modeling more generally (see RxInfer.jl Learning Group for more information).

Partnership Application

The Active Inference Institute is seeking partners to help support and expand the Active Inference open source ecosystem. Partnerships allow organizations to align with and contribute to the Institute's mission of learning, researching, and applying Active Inference for the benefit of all.

Please submit any questions, pre-submission inquiries, or completed applications, to blanket@activeinference.institute, including [PARTNERS] in the email subject line. The Institute team will review and respond regarding next steps.

See the Partnership stages for an overview, and the full application information below.

Partnership stages

Stage	Description	Key Actions	Outcomes
Stage 0: Initial Engagement	Determining the most suitable mode of collaboration (philanthropy, direct participation, or partnership).	• Conduct exploratory discussions with prospective partners to understand their goals and capacities. • Assess alignment with the Institute's Mission, Vision, Values, and Principles. • Provide information on available pathways (Philanthropy, support for Institute Programs and Institute Projects).	Clear determination of whetl prospective partner prefers philanthropic contribution, d participation in projects, or a partnership.
Stage 1: Submit Application	Prospective partners express interest through a detailed Partnership Application form.	• Complete Partnership Application detailing desired level of involvement, proposed contributions, and alignment with mission. • Specify areas of interest (e.g., unrestricted support, targeted programs, collaborative projects).	A clear articulation of intent proposed scope of collabora from the prospective partne
Stage 2: Review & Selection	The Institute evaluates applications based on impact potential, alignment with priorities, and diversity considerations.	• Internal assessment of applications for fit with current priorities and capacity for collaboration.	Selection of partners who al the Institute's goals and brin unique value to its initiatives
Stage 3: Formal Agreement	Both parties establish a formal Partnership Agreement outlining commitments and terms of engagement.	• Draft and sign an agreement that includes contributions (financial/in-kind), roles, points of contact, and mutual expectations. • Ensure clarity on responsibilities and accountability mechanisms.	A structured framework for collaboration that sets clear expectations for both partie
Stage 4: Planning Session	Kickoff planning session to map out shared goals and develop an action plan for collaboration.	• Identify shared objectives and key milestones for the first 6–12 months. • Develop a detailed action plan with timelines and deliverables. • Schedule regular check-ins to monitor progress.	A collaborative roadmap witl defined milestones that guid initial partnership activities.
Stage 5: Ongoing Collaboration	Continuous engagement to ensure alignment and flexibility as the partnership evolves over time.	• Provide curated updates on relevant developments and opportunities. • Facilitate feedback loops to refine activities based on emerging needs or interests. • Adapt plans as necessary to maintain impact and mutual benefit.	Sustained progress toward s goals with room for adaptati based on feedback and evol priorities.

Application sections

1. Organization Information

Provide essential details about your organization to help us understand your background and context for collaboration:

- **Name of Organization**
- **Website URL**
- **Brief Description of Organization**: A concise overview of your organization's mission, activities, and focus areas (250 words max).
- **Type of Organization**: Specify whether your organization is academic, industry, government, non-profit, or another type.
- **Primary Domain(s) of Focus**: Indicate the key domains your organization works in (e.g., neuroscience, robotics, psychology).
- **Primary Contact Information**: Include the name, title, and email address of the main point of contact for this partnership.
    - *Optional*: List additional team members who may be involved in the partnership.

2. Alignment with Institute Mission

This section allows you to demonstrate how your organization aligns with the goals and principles of the Active Inference Institute:

- **Alignment with Active Inference Principles**: Describe how your organization's work aligns with or could benefit from Active Inference methodologies (500 words max).
- **Relevant Projects or Initiatives**: Highlight specific projects, research areas, or initiatives within your organization that are relevant to Active Inference.
- **Past Interactions**: Share any previous interactions with the Active Inference Institute, its personnel, programs, or activities.
- **Motivation for Partnership**: Explain why your organization is interested in partnering with the Institute and what goals you hope to achieve through this collaboration.

3. Proposed Partnership

Outline the scope and nature of the partnership you are proposing:

- **Scope and Duration**: Provide details on the intended scope (e.g., specific programs or projects) and anticipated duration of the partnership (500 words max).

- **Type and Level of Support**:
  - *Financial Contribution*: Specify an amount or range for financial support.
  - *In-Kind Support*: Describe any non-monetary contributions (e.g., computational resources, library access, staff time).
- **Expectations from the Institute**: Clearly state what you expect from the Institute's role in this partnership, such as:
  - Public recognition of the partnership on its website or materials.
  - Facilitation of connections with researchers, contributors, or stakeholders within the Institute's network.
  - Coordination of joint activities around shared interests.
  - Incorporation of partner feedback into planning processes.
  - Co-development of potential joint initiatives aligned with both organizations' missions.

4. Partnership Expectations

Confirm your organization's commitment to fulfilling its responsibilities as a partner:

- Abide by the agreed-upon Terms of Partnership.
- Make agreed-upon financial or in-kind contributions as outlined in the formal agreement.
- Actively participate in and help promote relevant Institute activities.
- Provide a responsive and reliable point of contact for ongoing communication and coordination.
- Respond promptly to periodic requests from the Institute for information or feedback to ensure effective collaboration.

5. Additional Information

Share any supplementary details that may be relevant to evaluating your application:

- **Related Partnerships**: Describe any existing partnerships that may complement or intersect with this collaboration.
- **Relevant Publications**: Provide references to key publications or outputs related to Active Inference or aligned fields.
- **Potential Conflicts of Interest**: Disclose any conflicts of interest or competing engagements that may affect this partnership.
- **Exploratory Meeting Availability**: Indicate your availability for an initial exploratory meeting to discuss partnership potential further.
- **References**: Provide names and contact information for 1–5 references who can speak to your organization's capabilities and suitability as a partner.

We look forward to exploring how we can work together to advance Active Inference research and applications through meaningful collaboration.

If you have any questions or require additional information about this process, please do not hesitate to contact us at:

Blanket@ActiveInference.Institute

Partnership Terms

Terms v1 (April 2024)

Active Inference Institute Partnership Program Terms and Conditions

These Terms and Conditions ("Terms") govern the partnership between the Active Inference Institute ("Institute") and the partnering organization ("Partner"). By submitting an application to the Partnership Program ("Program"), the Partner agrees to be bound by these Terms.

1. Partnership Scope and Duration

1.1 The scope and duration of the partnership shall be as mutually agreed upon by the Institute and Partner in writing, based on the Partner's application and any subsequent discussions.

1.2 The partnership shall commence on the date the Partner is notified of acceptance into the Program and shall continue until the agreed upon end date, unless terminated earlier in accordance with these Terms.

2. Partner Obligations

2.1 The Partner shall make the financial or in-kind contributions specified in their application and agreed upon with the Institute. Contributions are non-refundable.

2.2 The Partner shall participate in and help promote Institute activities relevant to the partnership, as mutually agreed upon.

2.3 The Partner shall provide a designated point of contact who is responsive and reliable in communicating with the Institute.

2.4 The Partner shall respond to periodic Institute requests for information and feedback in a timely manner.

2.5 The Partner shall inform the Institute of any changes to their organization that may impact the partnership.

3. Institute Obligations

3.1 The Institute shall provide public recognition of the partnership, subject to the Partner's approval of any use of their name, logo, or other identifying information.

3.2 The Institute shall facilitate connections between the Partner and relevant Institute stakeholders, as appropriate for the agreed upon scope of the partnership.

3.3 The Institute shall coordinate joint activities with the Partner around areas of shared interest, as mutually agreed upon.

3.4 The Institute shall consider Partner feedback in its planning and decision-making related to the Program.

4. Intellectual Property

4.1 Each party shall retain ownership of any intellectual property they create or possess prior to or independently of the partnership.

4.2 Any intellectual property created jointly by the Institute and Partner in the course of the partnership shall be owned jointly, unless otherwise agreed in writing.

4.3 Neither party shall use the other party's intellectual property, including trademarks and logos, without prior written consent.

5. Confidentiality

5.1 The Institute and Partner may exchange confidential information in the course of the partnership. Each party shall maintain the confidentiality of such information and not disclose it to third parties without prior written consent, except as required by law.

5.2 Confidential information shall not include information that is publicly available, independently developed, or obtained from a third party without breach of any obligation of confidentiality.

6. Termination

6.1 Either party may terminate the partnership at any time upon written notice to the other party.

6.2 Upon termination, the Partner shall cease use of any Institute intellectual property and shall return or destroy any confidential information of the Institute in their possession.

6.3 Termination shall not affect any rights or obligations accrued prior to the effective date of termination.

7. Limitation of Liability

7.1 Neither party shall be liable to the other for any indirect, incidental, special, or consequential damages arising out of or related to the partnership.

7.2 The Institute's total liability under these Terms shall not exceed the amount of financial contributions made by the Partner in the 12 months preceding the event giving rise to liability.

8. Governing Law and Dispute Resolution

8.1 These Terms shall be governed by and construed in accordance with the laws of the jurisdiction in which the Institute is incorporated (Delaware, USA).

8.2 Any disputes arising out of or related to these Terms shall be resolved through good faith negotiation between the parties. If negotiation fails, the parties shall submit to mediation before proceeding to arbitration or litigation.

9. Miscellaneous

9.1 These Terms constitute the entire agreement between the parties with respect to the Program and supersede any prior agreements or understandings.

9.2 These Terms may be amended only by a written document signed by both parties.

9.3 Neither party may assign these Terms without the prior written consent of the other party, except that the Institute may assign these Terms to an affiliated entity.

9.4 If any provision of these Terms is held to be invalid or unenforceable, the remaining provisions shall continue in full force and effect.

By submitting an application to the Active Inference Institute Partnership Program, the Partner acknowledges that they have read, understood, and agree to be bound by these Terms.

Contact: Blanket@ActiveInference.Institute

Open Source

Open Source is key to the Mission, Vision, Values, and Principles of The Active Inference Institute.

The default Open Source license information for all Institute materials is **CC BY-NC-SA 4.0** and is described below. Check with specific products and collaborators for more information.

CC BY-NC-SA 4.0

Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International

This license requires that reusers give credit to the creator. It allows reusers to distribute, remix, adapt, and build upon the material in any medium or format, for noncommercial purposes only. If others modify or adapt the material, they must license the modified material under identical terms.

- **BY:** Credit must be given to the creator(s) of the work, the specific people where know & The Active Inference Institute as hosting or publishing entity.
- **NC:** Only noncommercial use of your work is permitted. **Noncommercial means not primarily intended for or directed towards commercial advantage or monetary compensation.**
- **SA:** Adaptations must be shared under the same terms (unless otherwise specified and agreed upon by creators).

Some of our main Open Source Repositories are listed in the table below, nested within the Institute Github.

Open Source Repositories

Name	URL	Description
Active_Inference_Ontology		Snapshots of Active Inference Ontology
ActiveBlockference		Active Blockference repository for integrating Active Inference with cadCAD and more
ActiveInferAnts		Active Inference Ant simulations, and much much more
ActiveInferenceCategoryTheory		Category Theory curriculum and materials
ActiveInferenceJournal		Primary repository for the Active Inference Journal, with transcripts from Production
AEOS		Snapshots of Active Entity Ontology for Science (AEOS)
Biofirm		Active Inference PyMDP agents for Bioregional Modeling
CEREBRUM		Case-Enabled Reasoning Engine with Bayesian Representations for Unified Modeling
GEN24		Generative AI experiments and deployments as part of Active Blockference (here)
GeneralizedNotationNotation		Information on Generalized Notation Notation
Journal-Utilties		Utilities for Active Inference Journal
PyDMB		Dysfunctional Markov Blanket package to accompany research paper
Symposium		Synthetic intelligence methods for Applied Active Inference Symposium
Textbook		Repository for Textbook Group

Ecosystem Support

Activities at the Institute offer resources and participation opportunities for individuals and organizations. These epistemic and pragmatic services include:

Informational Commons

- **Production** & **Active Inference Journal**
  - Largest corpus of Open Source Active Inference education materials available to date.
- **Common Forum**. Providing Institute Projects, hosting online forums, Applied Active Inference Symposium, discussion groups, and Communications channels where learners, researchers, and practitioners can connect, ask questions, and share insights. Fostering a community that helps individuals overcome challenges, exchange ideas, connect on collaborations, and receive support from peers and experts.
- **Opportunities to Share and Present Work**. Provide myriad opportunities to share and present relevant work on Active Inference, offering opportunities for unique collaborations and new knowledge discovery catalyzed by Active Inference and the consequent amplified leveraging of expertise and practices across disciplines, domains, and paradigms.

Infrastructural and Administration directions

- **Infrastructure**. Maintaining and developing information systems to support The Institute's activities, iteratively improving usability and efficacy. Pending funding, working with requirements engineering and user experience professionals to overhaul existing systems.
  - Improve provisioning and access to e.g. Implementations of Active Inference, Active Inference Ontology, Textbook Group & Education materials, Livestreams from Production, Educational Standards & Qualifications, etc.
- Partnership
  - Managing and growing relationships with Education, research, application, and service partners.
- Philanthropy
  - Development relationships with potential donors and sponsors, and, pending funding, developing the necessary infrastructure (e.g., accounting, legal, digital affordances, materials) to request and receive donor and sponsor support, and to offer and dispense micro-grants and financial support to researchers.
- **Funding Discovery &** Grants **Support**. Providing a variety of support mechanisms for participants to search for and submit to grant and funding opportunities, as well as assist them in forming partnerships (e.g., with other researchers, companies, and universities).
- **Professionalization**. Developing a curriculum of training programs for Officers and Directors of commercial entities and officials of governmental and civil society organizations to enhance their understanding of sentient behavior (as described by Active Inference) and its implications for organizational interactions in the areas of Business, Operations, Legal, Technical, and Social.

Institute Projects

See Activities for updated information on active projects.

Institute Projects are the primary means of participation with The Active Inference Institute.

To date, The Institute has hosted or facilitated the development of hundreds of Open Source licensed products which serve various functions in The Active Inference Ecosystem including Awareness, Education, Commons, Support, and Governance.

Project Rhythm Through Prepare and Measure

The Institute implements a unique "Prepare and Measure" system that structures project work through alternating phases of preparation and measurement.

To complete a Project ~ Preparation, participants propose a phase of activity — their “packed backpack” and intention for developing artifacts, research, or create educational materials while receiving ongoing feedback.

This is followed by making a Project ~ Measurement, where the participant documents their reports and reflections. Following the measurement, next steps are explored. This rhythmic approach creates natural checkpoints for reflection while maintaining momentum

Benefits and Implementation

The prepare-measure cycle embodies active inference principles by balancing exploration with evaluation (see Physics course). Rather than following rigid schedules or purely passive learning, participants actively sample their environment through concrete project work, while regular measurements provide the feedback needed for learning and course correction.

This system helps cultivate a culture of active sensemaking, where putting work out for feedback is encouraged over passive consumption. The flexibility of this approach allows it to scale from individual contributors to large collaborative projects, while maintaining rigor through consistent documentation and assessment.

Towards a Systems Approach Project Framework

We make ongoing incremental updates to the approach taken across Institute Projects. Current thinking on this is considering updates in the area of:

Project Structure

- Stronger connection with EduActive (Education) or ReInference (Research) Organizational Units.
- Each project will have a clearer standardized public profile featuring:
  - Clear mission statement and objectives
  - Timeline with key milestones and deadlines
  - Contribution pathways and skill requirements
  - Active measurement cycles and preparation phases

Project Management Approach

Time Management

- Dedicated work blocks outside of meetings
- Regular preparation and measurement cycles

- Minimal reliance on email/Discord for core work

Task Organization

- Public task tracker with clear ownership
- Regular progress updates and milestone reviews
- Documentation of both successes and learning opportunities
- Integration with prepare/measure cycles

The Institute aims to implement these projects using active inference principles, ensuring each initiative contributes to our mission of making active inference more accessible, rigorous, and applicable while serving our growing global community.

https://projects.activeinference.institute/

AICACP

AI Capabilities & Alignment Consensus Project

We are happy to share that in June 2025, the Survival and Flourishing Fund has awarded a $270,000 grant to the Active Inference Institute to support work on the **AI Capabilities & Alignment Consensus Project** (🧙 AICACP).

AICACP is a multi-year initiative designed to reshape the conversation around AI capabilities, alignment, and regulation. By combining high-impact journal collections, in-person discussion-oriented workshops, and academic media content for public outreach, the project aims to bridge the divide between AI “doomers” and “accelerationists” through deeply exploring the meanings of “world models” and “agency,” and what these concepts mean for AI development.

Adam Safron (at Allen Discovery Center at Tufts University) is the creator, principal investigator, and organizing editor for this project. This grant supports an expanded team to assist with efforts to manage the special issues and workshops.

For those interested in learning more, please check out this document:
AI Capabilities and Alignment Consensus Project - Google Docs

Please reach out to asafron@gmail.com with any thoughts or questions.

Active Blockference

Much information is contained in the Active Blockference project documentation.

Active Blockference is an open-source project developed by the The Active Inference Institute that aims to create a comprehensive cognitive modeling framework for complex systems. The project combines two powerful technologies: cadCAD (Complex Adaptive Dynamics Computer-Aided Design) and Implementations of Active Inference.

Project Overview

The primary goal of Active Blockference is to develop a simulation environment that can model the cognitive processes and goal-directed behavior of agents within various complex systems across Domains of Application. This framework is designed to:

1. Facilitate rigorous analysis of multi-agent systems and their emergent behaviors
2. Serve as a sandbox for exploring cognitive, micro-economic, behavioral, and decision-making processes
3. Enable cognitive audits of protocols and systems across multiple domains

Key Components

- **Open-source package**: Integrates cadCAD and Active Inference implementations for theoretical and applied studies
- **Multi-agent simulation**: Expands from single-agent to multi-agent models to explore cognition and behavior in various settings
- **Educational resources**: Develops materials to onboard new users to the Active Blockference community

Rationale for Combining cadCAD and Active Inference

The integration of cadCAD with an Active Inference kernel provides a powerful synergy for modeling complex systems:

1. **cadCAD**: Offers a robust framework for simulating complex adaptive systems, allowing for the modeling of multi-agent interactions and system-level dynamics.
2. **Active Inference**: Provides a principled approach to modeling goal-directed behavior and decision-making processes of individual agents.

By combining these technologies, Active Blockference can model both the macro-level system dynamics and the micro-level cognitive processes of agents within those systems.**Project Status and Development**

As of October 2024, Active Blockference is in active development. The project welcomes participants from various backgrounds to contribute to the growing codebase on GitHub.

Get Involved

Interested individuals can participate in Active Blockference through:

1. Joining ongoing discussions on the Active Inference Institute's Discord server
2. Contributing to the Coda document for asynchronous collaboration
3. Exploring and contributing to the GitHub repository

By developing this cognitive layer for complex systems modeling, Active Blockference aims to enhance our understanding of multi-agent dynamics and decision-making processes across a wide range of applications.

Active Blockference

https://coda.io/d/_dIvNESFmyj6

Active InferAnts

The Active InferAnts project works towards integrated modeling of Ants and Environments.

Working documentation page.

https://github.com/ActiveInferenceInstitute/ActiveInferAnts

Project abstract

The Active InferAnts project is an ambitious computational modeling initiative that aims to implement Active Inference principles and Implementations of Active Inference for ant colony simulations. The project refactors existing ant simulation code to use the Active Blockference package, enabling more flexible and scalable multi-agent modeling of ant behavior through stigmergic interactions and simultaneous localization and mapping (SLAM).

At its core, the project models individual ants (Nestmates) as active inference agents that interact with their environment through pheromone trails and other nestmates, while incorporating both low-level movement decisions and higher-level task selection behaviors like foraging, nursing, midden work, and nest architecture.

The simulation captures the complex interplay between individual ant cognition and colony-level emergence through a hierarchical generative model framework that includes pheromone-based communication, spatial navigation, and task switching dynamics.

As the project Github attests to, software development efforts have also woven together elements related to P3IF, William Blake, Large Language Models (LLM), and more.

Active Inference Journal

The Active Inference Journal is an Open Source project launched in 2021 to create a comprehensive, accessible repository of Active Inference knowledge through automated transcription and processing of educational content. The project combines sophisticated language processing pipelines, collaborative editing workflows, and decentralized storage solutions to transform video lectures, discussions, and presentations into richly indexed, searchable, and citable content.

The Journal serves as a crucial bridge between traditional academic publishing and emerging Decentralized Science (DeSci) approaches, enabling broad participation in Active Inference scholarship through several key innovations: automated speech-to-text transcription reducing manual effort, standardized editorial practices for community contribution, version control through GitHub, and integration with knowledge engineering projects. This infrastructure allows researchers, practitioners, and learners worldwide to not only access but actively contribute to the development and documentation of Active Inference theory and applications through open-source collaboration, while maintaining academic rigor through systematic processing pipelines and metadata management.

Core Open Source Journal Repositories

- ActiveInferenceJournal: Main content repository containing transcripts, translations, and published materials: https://github.com/ActiveInferenceInstitute/ActiveInferenceJournal
- Journal-Utilities: Technical infrastructure for automated processing, including speech-to-text, translation, and knowledge extraction tools
  https://github.com/ActiveInferenceInstitute/Journal-Utilities
- Coda documentation for the project.

Active Inference Journal

https://coda.io/d/_dwYsKMwppRN

Active Inference Ontology

The Active Inference Ontology has been a core Institute project since 2021.

- Home page of the Ontology documentation
- Github repository of Ontology snapshots

The screenshot below represents the table with current version of ontology, here groups into the list of 65 Core terms (central terms for learning and applying Active Inference), 301 Supplemental terms (possibly useful terms to know for some situations or models), and 74 entailed terms (common words that may have specific relevance for the topics discussed in other terms lists).

The Active Inference Ontology is used in the Knowledge Engineering analyses.

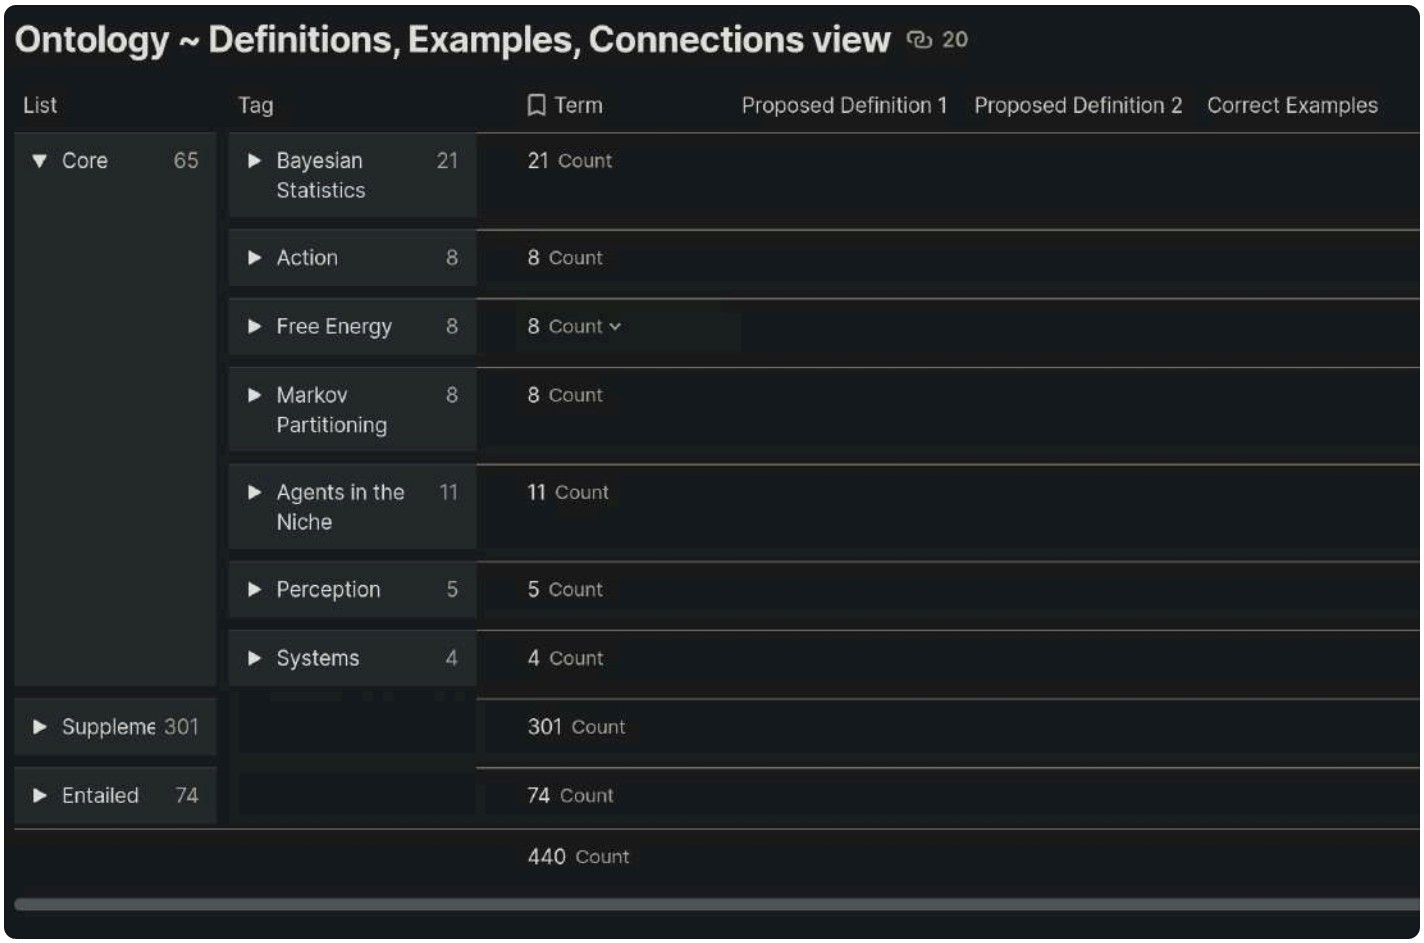

During 2024 we have continued to have amazing discussions, add examples/counterexamples, add connections, add translations, and more

The Active Inference Ontology project is a core initiative of the Active Inference Institute that aims to develop and maintain a structured framework for understanding key concepts, terminology, and relationships within Active Inference theory. Started in 2021, this open-source ontology serves multiple purposes - from supporting education and research to enabling computational applications and cross-domain translation of Active Inference concepts.

The ontology is publicly hosted and continuously updated through a living document system (public Coda, Github) with stable versions released periodically. It contains curated definitions, examples, translations across languages, and steps towards formal logical relationships among terms.

Core Functions

- Reduces uncertainty around fundamental Active Inference concepts and terminology
- Facilitates coherent, rigorous, inclusive research and applications across domains
- Enables effective onboarding of new learners
- Supports automated inference and computation
- Provides translation capabilities across languages and fields

Development Areas

The project focuses on several key development tracks:

- Definition refinement and example curation
- Applied usage in building computational generative models for Active Inference
- Multi-language translations
- Formal logical expressions (SUMO and other systems)
- Literature analysis and knowledge engineering
- Application testing across different use cases

Contribution Methods

The ontology welcomes contributions through multiple pathways:

- Adding/reviewing definitions and examples
- Contributing translations
- Developing formal logical expressions
- Testing applications in research and education
- Participating in discussion and refinement

Rather than being a static reference, the Active Inference Ontology functions as a living knowledge system that evolves with the field while maintaining rigor and accessibility. It serves as both a practical tool for working with Active Inference concepts and a scaffold for developing deeper understanding across the Active Inference ecosystem.

The project exemplifies Active Inference principles in its own design - it actively reduces uncertainty about core concepts while enabling generative exploration and application across domains. This makes it a crucial resource for researchers, educators, practitioners and learners engaging with Active Inference theory and applications.

The below image describes some of the work we have done across a continuum of levels/types of formalization. As with all products, this is a work in progress where Open Source contributions will be welcome.

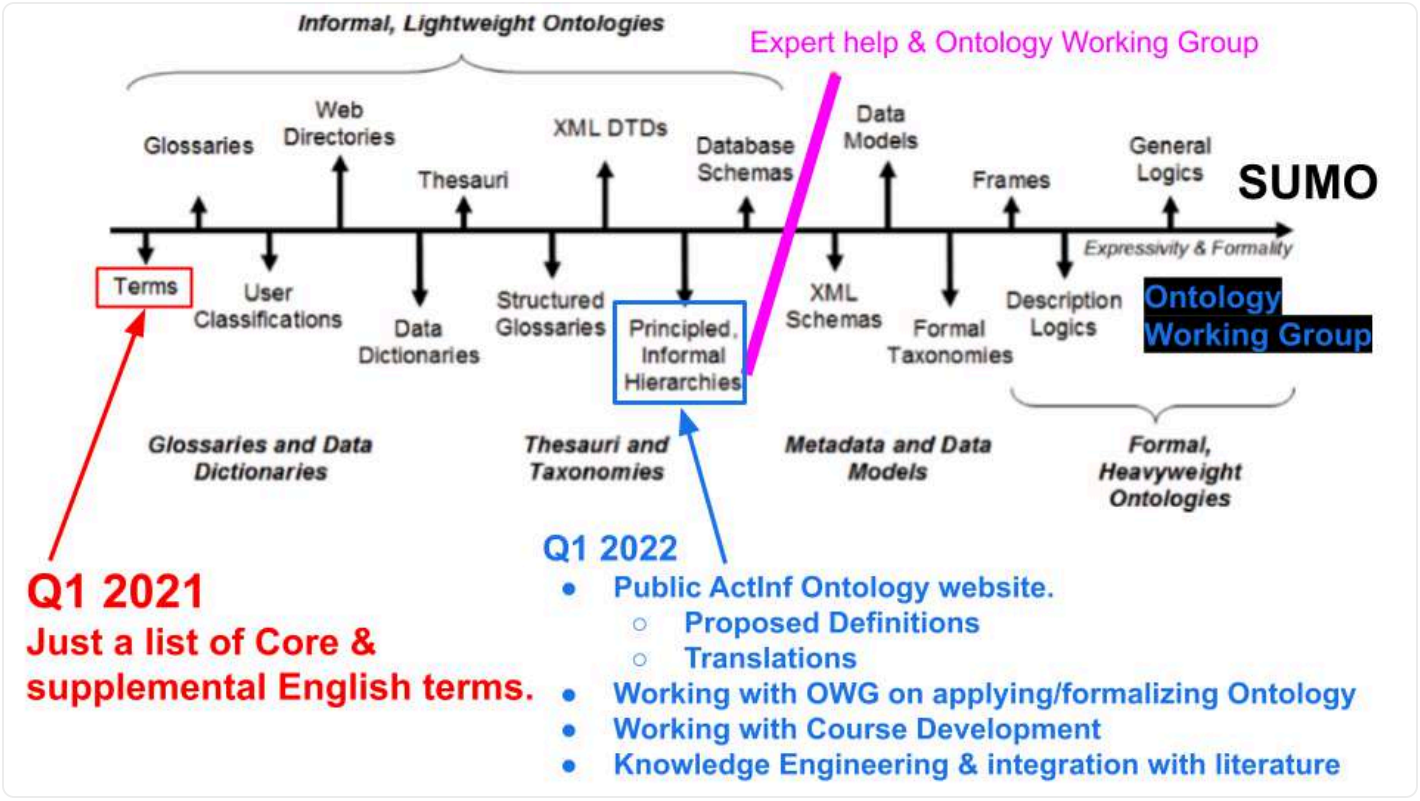

Q1 2021

Just a list of Core & supplemental English terms.

Q1 2022

- Public ActInf Ontology website.
  - Proposed Definitions
  - Translations
- Working with OWG on applying/formalizing Ontology
- Working with Course Development
- Knowledge Engineering & integration with literature

All ~ Active Inference Ontology

https://coda.io/d/_djD38E5fJk_

Active Entity Ontology for Science (AEOS)

≈ Active Entity Ontology for Science (AEOS) represents a key framework developed by a 2022 Institute Projects, created to bridge centralized and decentralized approaches to scientific organization.

The work was published as "An Active Inference Ontology for Decentralized Science: from Situated Sensemaking to the Epistemic Commons" (link), and is also available on Github and in an interactive Coda format.

An Active Inference Ontology for Decentralized Science: from Situated Sensemaking to the Epistemic Commons

Friedman, Daniel[1]; Applegate-Swanson, Shaun[1]; Balbuena, Jessica Angeli[1];
Choudhury, Arhan[2]; Cordes, RJ[1]; El Damaty, Shady[3]; Guénin—Carlut, Avel[4];
Knight, V. Bleu[1]; Metelkin, Ivan[2]; Shrivastava, Siddhant[5]; Singh, Amit Kumar[2];
Smékal, Jakub[6]; Tuttle. Caleb[3]; Vyatkin, Alexander[2]

Show affiliations

In this work, we examine science from the vantage points of blockchain technology and its connection to decentralized science (DeSci). We consider science as a collective process using Active Inference, an integrative framework that models the cognitive processes of perception, planning, and action selection in terms of Bayesian probabilities and updating. We present the Active Entity Ontology for Science (AEOS, available at coda.io/@active-inference-institute/active-entity-ontology-for-science-aeos) as a composable and versionable system for modeling various science systems, using the Active Inference entity partitioning. Such DeSci systems are considered from the perspective of BOLTS (Business, Operations, Legal, Technical, Social). Further steps for developing and utilizing AEOS in the context of scientific ecosystems are provided.

Here are the key aspects of AEOS:

Core Components

Framework Structure

- Uses Active Inference principles to model different forms of scientific activity as a collective cognitive process occurring in a niche.
- Provides a composable and versionable system for modeling various scientific systems
- Integrates BOLTS perspective (Business, Operations, Legal, Technical, Social) for comprehensive analysis

Key Functions

- Maps relationships between different scientific entities and processes using Active Inference entity partitioning
- Enables modeling of both traditional institutional science (CeSci) and decentralized science (DeSci) approaches
- Facilitates bottom-up sensemaking while maintaining systematic organization

Implementation Goals

Scientific Organization

- Supports emergence of epistemic communities through organic collaboration
- Enables transparent resource allocation and knowledge sharing
- Provides tools and analytic methods for decentralized scientific governance

Practical Applications

- Guides development of tools for scientific collaboration, Web3 or otherwise.
- Helps structure new kinds of organizations for research purposes
- Supports integration of blockchain and other technologies into scientific workflows

The AEOS serves as a bridge or blanket between:

- **In theory** and in terms of generalities: Active Inference Ontology
- **In practice**: Existing and emerging decentralized approaches, providing a structured way to understand and implement new forms of scientific organization while maintaining rigor and effectiveness in DeSci and traditional scientific settings.

Applied Active Inference Symposium

The Applied Active Inference Symposium highlights ongoing work related to Active Inference across domains.

📣 See the live 2025 Symposium Program

Online, November 12-14, 2025 (See the program).

The proceedings of past 🚀 Applied Active Inference Symposium are available:

- 1st in 2021 — Karl Friston (transcript, video part 1, 2, 3)
- 2nd in 2022 — Robotics
- 3rd in 2023 — Enacting Ecosystems of Shared Intelligence
- 4th in 2024 — Program, video part 1, part 2, part 3, playlist of videos
- 5th in 2025 — "Industry" November 12-14: See the program

More

- See 🛩️ Sponsorship at the Symposium
- Any other ideas? Please email us.

---

https://symposium.activeinference.institute/

Call for Presenters

Information on being a presenter for the Applied Active Inference Symposium

We are excited to invite researchers and practitioners to submit presentations for the upcoming **5th Applied Active Inference Symposium 2025**, to be held on **12-14 November, 2025**.

This online Symposium will focus on exploring the frontiers of Active Inference with a special focus on its **applications in industry**. As with last year's Symposium, the keynote address and panel will feature Karl Friston.

All information: symposium.activeinference.institute/

Complete this form to submit your presentation proposal, and/or read on for more details. Please feel free to share this invite with anyone else who you think would be a great addition to the Symposium.

Symposium Themes and Topic

We are welcoming presentations and workshop submissions that align with the Symposium themes and tracks:

- **Overall Theme**
  - **Applying Active Inference in Industry:** In 2025 we are seeing increased application of Active Inference across industry domains. What are the success stories? What was learned along the way? What challenges remain? What kind of applied research would accelerate applications?
- **Topics**
  - **Open Science & Tool Development**: Increasing the function and accessibility of Active Inference software across languages. Focused tracks and workshops on practical aspects of Active Inference.
  - **Learning & Education:** Share initiatives and experiences in teaching and spreading knowledge about Active Inference. Connecting the global Active Inference community with each other, resources, and research/education opportunities.

Interactive Aspects & Specific Affordances

Submissions of the following types are welcomed:

- **Presentations:** Pre-recorded video presentations of 15-60 minutes, or livestream slots of 30-90 minutes for panels or group discussions.
- **Workshops, Tutorials, and Hackathons**: Hands-on sessions of 60-180 minutes, that engage participants with Active Inference methods and provide open source materials/outputs.

Accessibility & Participation

We are committed to making this online symposium accessible to all.

- **Participant registration:** The Symposium will be online, with no financial barrier to participation. We will provide pathways for participants to foster deeper involvement with ongoing research and projects in the ecosystem.

- **Collaboration with Active Inference Journal**: We will publish the Abstracts of presentations (similar to last year) to give visibility to proceedings and enhance availability of the materials
- **Translation**: We will translate materials into different languages where possible, to provide content that reaches a broader audience (the organizing team will support presenters on this).
- **Multimodal Content**: audio, visual, interactive/real-time simulations or other media that help communicate complex ideas. The content will remain posted on our YouTube channel for asynchronous viewing following the symposium.

Important Dates

- **Presenter Submission Deadline**: October 31, 2025
- **Symposium Dates**: November 12-14, 2025

How to Submit

Please submit your abstracts and proposals via this form.

Symposium co-organizers:
Alexander Ororbia, Alexandra Mikhailova, Andrew Pashea, Bradly Alicea, Christian Martens, Cory Slater, Daniel Friedman, Marc Broberg, Maria Garcia, Maria Luiza Iennaco, PabloFM, Rorik Smith, Sylvia Zhang, Bleu Knight, Zohreh Rahmannejad, Alex Vyatkin

Contact Information:
For any inquiries, please reach out to blanket@activeinference.institute .
We are also seeking sponsors at this time.
Information on participant registration will be available here closer to the date.

Sponsorship at the Symposium

Support the Applied Active Inference Symposium with financial or in-kind donations

More information on Sponsorship at the Symposium in 2025 coming soon.

Why Sponsor the Applied Active Inference Symposium?

Active Inference represents a cutting-edge approach to artificial intelligence and cognitive science, especially in the modern and open science setting. By supporting the Applied Active Inference Symposium, sponsors are contributing to:

- **Scientific Innovation:** Active Inference offers a rigorous, first-principles approach to AI and cognitive science, promising significant computational benefits and innovative system designs.
- **Open Science & Collaboration:** The Active Inference Institute is a 501(c)(3) non-profit organization that promotes open-source practices, global participation, and collaborative learning, accelerating progress through shared knowledge. This Symposium is one of the few inclusive opportunities available at this time to participate in hands-on workshops and connect with other people applying Active Inference.
- **Broad Impact:** Supporting Active Inference research and applications, can contribute to addressing global challenges, fostering interdisciplinary applications, and promoting ethical transparency in AI development.
- **Strategic Advantage:** Sponsorships provide visibility to talent, aligns with funding requirements, and facilitates valuable partnerships, positioning sponsors at the forefront of AI innovation.

Sponsorship tiers

Sponsorship level	Benefits
1	• Listed in Symposium Program and materials. • Acknowledged in Livestreams

2	**As above** plus:
	• Logo included in the Adventure space
	• Video Presentation to be played on livestream (1-5 minutes pre-recorded video)
	• We can make introductions with you and:
	  ◦ Symposium Presenters, Co-Organizers, Participants
	  ◦ Institute Research Fellows, Officers, Board of Directors, Scientific Advisory Board, Interns.

3	**As above** plus: • Consideration of a session/speaker, or theme, to be included the program. • Arrangement before or after the Symposium, of a private session on theme of your choice. • *Optional*: Custom commissioned hand-drawn Active Inference art

How Organizations Can Contribute

- Provide financial support at desired sponsorship level (or suggest something else?)
- Offer in-kind support (e.g., services, products)

- Optional: Contribute additional resources (e.g., presenters, materials)

Impact of Your Support

Your sponsorship directly contributes to:

- Advancing open-source Active Inference research and applications
- The sustainability of the Active Inference Institute, and success of the 📌 Institute & Ecosystem.
- Promoting innovation in AI and cognitive science
- Supporting the coming generations of Active Inference researchers and practitioners

Join us in shaping the future of intelligent systems through Active Inference!

Email blanket@activeinference.institute with any inquiries or questions.
Provide financial support directly at this link.

CogNarr (Cognitive Narrative) Ecosystem: Facilitating Group Cognition at Scale

An ongoing Active Inference Institute project facilitated by John Boik, PhD

Introduction

John Boik, **Research Fellow (May 2024 -)**

ORCID: 0000-0003-1289-7997

Human groups of all sizes and kinds engage in deliberation, problem solving, strategizing, decision making, and more generally, cognition. Cognition in the group setting serves a similar purpose to cognition in individual humans. From an active inference perspective, that purpose is to achieve and maintain, with high certainty, those preferred conditions that promote health and wellbeing. As we know from common experience, the quality of group cognition can range from functional to dysfunctional, productive to unproductive, and thoughtful to superficial. As such, the quality of a group's cognitive process can either lead the group toward or away from health, wellbeing, security, and goal achievement.

Achieving functional, productive, and thoughtful cognition is especially difficult in the large group setting. The small-group setting often involves face-to-face dialogue, which can support rich and dynamic interactions that allow all voices to be heard. But such interactions are more difficult to achieve in the large-group setting, which typically requires some form of online communication. New approaches are needed to facilitate the kind of rich communication and information processing that are required for effective, functional, productive cognition in the online setting, especially for groups characterized by hundreds, thousands, or millions of participants who wish to share potentially complex, nuanced, and dynamic perspectives.

The incipient CogNarr (Cognitive Narrative) Ecosystem is intended to facilitate functional cognition in the large-group setting. A key perspective is to view a group as an organism that uses some form of cognitive architecture to sense the world, process information, remember, learn, predict, make decisions, and adapt to changing conditions. The CogNarr ecosystem is designed to serve as a component of that architecture.

The CogNarr project at the Active Inference Institute is intended to bring CogNarr to life. CogNarr is potentially a massive project, involving many topics, tasks, needs, researchers, staff, volunteers, and so on. If you have questions or suggestions, want to learn more, or want to help, please write to me via "blanket@activeinference.institute" with "[COGNARR]" in the subject line and I will respond. You can also join our biweekly meetings, every other Wednesday in the morning US Mountain Time. Meeting events are listed on the Active Inference Institute calendar.

Papers and Videos

The CogNarr Ecosystem is described in two recent papers:

- Boik, JC., 2024a. CogNarr Ecosystem: Facilitating Group Cognition at Scale. ArXiv.
- Boik, JC., 2024b. CogNarr Ecosystem: Preliminary Thoughts on a Story Graph Meaning Representation. OSF.

Each paper is also discussed in an Active Inference Institute GuestStream:

- GuestStream #087.1~ 9/5/2024
- GuestStream #087.2 ~ 9/10/2024

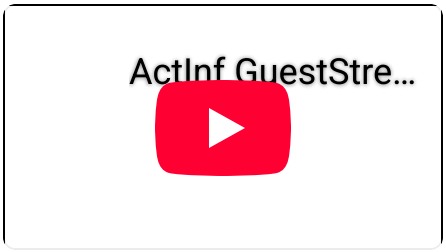

Relation to Previous Work by John Boik

The CogNarr Ecosystem is a continuation of ideas discussed in a previous series of three papers, the 2014 book *Economic Direct Democracy: A Framework to End Poverty and Maximize Well-Being*, and a simulation paper that describes how a novel economic system (part of the Local Economic Direct Democracy (LEEDA) framework, might function. These are discussed on the PrincipledSocietiesProject website, a prior effort by John Boik.

While the focus of CogNarr is on the design and development of a cognitive architecture that facilitates group communication and cognition in general settings, the previous series of papers focus on the de novo design of core societal systems (e.g., economic, financial, governance, legal), viewed as the cognitive architecture of a society. Thus, all five papers focus on the design of cognitive architectures for groups, but differ in their use cases. That being said, the approaches discussed in the CogNarr series, or even extensions of CogNarr itself, could one day be incorporated into a collaborative governance system of the kind first envisioned in *Economic Direct Democracy*.

Previous works include:

- Boik JC. Science-Driven Societal Transformation, Part I: Worldview. *Sustainability*; 12(17), 2020, 6881.
- Boik JC. Science-Driven Societal Transformation, Part II: Motivation and Strategy. *Sustainability*; 12(19), 2020, 8047.
- Boik JC. Science-Driven Societal Transformation, Part III: Design. *Sustainability*; 13(2), 2021, 726.
- Boik JC. First Micro-Simulation Model of a LEDDA Community Currency-Dollar Economy. *International Journal of Community Currency Research*; 2014.

Four LiveStreams discuss the 2020/2021 series:

- LiveStream #021.01

- LiveStream #021.02
- LiveStream #021.03
- LiveStream #021.04

See also:

- 4th Applied Active Inference Symposium presentation, Nov 14, 2024:
  CogNarr Ecosystem: Facilitating Group Cognition at Scale.
- Livestream presented on February 10, 2025 to the Theoretical Neurobiology (TNB) Group:
  CogNarr Ecosystem: Facilitating Group Cognition at Scale.
- 5th Applied Active Inference Symposium presentation, Nov 14, 2025:
  Story Graphs: An Exploration of Use Cases In and Beyond CogNarr

Funding/Support

The CogNarr project currently is unfunded. Your donation to the Active Inference Institute, earmarked for the CogNarr project, can help get this project off the ground. The institute is a US 501(c)3 nonprofit organization. Any level of funding is welcome and appreciated. The following table provides examples of how funds might be spent.

Note that CogNarr is envisioned as an open-source project. Monetization opportunities exist and potential development partners and social investors are encouraged to reach out. For example, third-party enhancements resting on top of the core stack, and second- and third-party provision of services (such as customization, consulting, and training) are possible. Further, CogNarr apps could be designed for a wide range of use cases in civil society, government, education, and industry. These could be customized to address specific needs. Social-media-like applications are also possible.

Support Examples

Support Level	Example Task/Result
< $1,500	Funds used for miscellaneous expenses and/or pooled until sum reaches one of th
$1,500	Approximately sufficient to pay open-access journal processing fees for one resea access to the published article.
$50,000	Approximately sufficient to support one lead researcher, full-time, for three month longer period. During this time, one research paper could be written.
$200,000	Approximately sufficient to support one lead researcher, full-time, for a year (inclu papers could be written. Alternatively, the funds could be split between multiple re
$1,000,000	Approximately sufficient to support a very small development team for one year to (demonstration product). Alternatively, the funds could be split between multiple r
$10,000,000	Approximately sufficient to support a small to modest development team for one y
$100,000,000	Approximately sufficient to give the CogNarr project a solid financial foundation, s projects. At this level of funding, especially if provided annually, the full CogNarr p releases, outreach and education, user support, partnership development, govern resources, and establishment of a prolific research team.

Road Map

If and when substantial funding becomes a possibility, the CogNarr team will construct a carefully considered roadmap for project development. In the meantime, the 2024 series of papers mentions numerous ideas for future work. Some of these are listed below, along with non-research tasks. The whitepapers listed below are intended as living documents that inform the CogNarr team and CogNarr stakeholders about available options. CogNarr governance, policies, and organization would evolve to reflect suggested recommendations and best practices. Each research task listed below is expected to result in one or a series of scientific papers. In contrast, tasks in the analysis section might lead to internal or public reports or whitepapers. Given the wide range of tasks listed below, many contributors are needed and individuals interested in CogNarr, regardless of their skill set, are encouraged to contact us.

Tasks

Task Category		Task
▼ **Research**	19	Define an initial meaning representation design for story graphs and story graph fragments, which migh Representation Tree Structure, Type Theory with Records, or DisCoCirc.
		Generate a few to several initial user stories, in text form, that would be used to manually create story g as initial demonstration examples.
		Generate a larger set of user stories, ideally created by real individuals who are interested in the CogNa might be static and written in text form. Alternatively, and ideally, they would be dynamic and interactiv might submit a written story, but then work with the CogNarr team to manually (and eventually automat into a story graph representation. The initial set of user-generated stories might be focused on one or a interest to the public. One option might be to engage a college or classroom in story creation, such that interact with participants over an extended period of time (perhaps months or even years), as the CogN
		Explore storage and recall technologies for story graphs, including property graphs. As part of this, exp ecosystem, and C-sets in particular, might be employed.
		Explore how concepts from applied category theory might be employed to achieve compositionality, re complexity, facilitate translation of story graphs into alternative forms, and facilitate zoom-in and zoom- structure and meaning.
		Explore different logical, probabilistic, or other computational models that could be used to generate an story graphs, or to generate probabilistic forecasts from story graphs.
		Explore the set of linguistic and knowledge-representation phenomena that the initial CogNarr design s
		Explore how PTLMs (LLMs) might be used in CogNarr, and the relative benefits of vanilla vs advanced PTLMs, hybrids, and alternative models for NLI and graph transformations in CogNarr.
		If C-sets look promising as a basis for story graphs, develop an efficient query language for C-set prope only allow complex traversals over story graphs, but would allow arbitrary Julia functions to be applied a path. Also, develop graph-to-graph, graph-to-code, and graph-to-text functors that translate story gra
		Explore methods to test and validate the NLI capabilities of CogNarr, given that the user-interaction asp that of traditional NLI experiments.
		Explore methods to compare and contrast story graphs, summarize story graph characteristics, compu for ideas contained in story graphs, or otherwise assess and summarize the information contained in a

	Explore methods to assess the quality of a story graph, and the quality of a set of story graphs.
	Explore methods to direct a group's attention to potentially important aspects of a story graph or set of
	Explore how tensor networks might be used in CogNarr computational models.
	Explore how distributed computation might be used in CogNarr models, and the kinds of inference mod that would lend themselves to distributed computation.
	Explore how user dynamics, especially evolution of a group beliefs toward shared belief models, might I inference or other approaches.
	Explore methods to quantify or otherwise evaluate the quality of group cognition in the CogNarr setting how CogNarr addresses each of the 13 hallmarks of basal cognition. Also explore methods to quantify c information flow in the group setting.
	Explore how CogNarr, and the concepts of cognitive science embedded in CogNarr, might impact publi function, purpose, and meaning.
	Explore how concepts of criticality in the group decision-making setting relate to concepts of deep dem promotes group cognition via a process that might occur at the edge of criticality.
▼ **Governance, Policies and Organization** 9	Develop whitepaper on options for CogNarr governance, including interim governance during initial pha
	Develop whitepaper on options for the CogNarr reputation system.
	Develop whitepaper on options for protection of user's personal information, including topics of data pri ownership, and data control.
	Develop whitepaper on rules and policies related to users and user groups.
	Develop whitepaper on options for CogNarr organization, including teams/departments, roles, responsil building.
	Develop whitepaper on options for CogNarr transparency, with regard to technology (including comput policy and governance.
	Develop whitepaper on options for open-source development, including topics related to transparency, resting on top of the core stack, second- and third-party provision of services (such as customization, and customization of apps for specific groups or use cases.
	Develop whitepaper on the relevant regulatory and legal issues that CogNarr might face.
	Develop a monetization plan for the CogNarr project, so that it might eventually become self-supportin
▼ **Outreach, Education, and Marketing** 7	Identify potential education, government, civil society, and industry groups who have interest in CogNa stakeholders, focus groups, promoters, partners, and so on. Further, identify how CogNarr might serve
	Identify opportunities within education, government, civil society, industry and elsewhere to provide ed science in general and the goals and designs of CogNarr, in particular.
	Seek partnerships within education, government, civil society, industry and elsewhere to engage partic users, and to develop a pool of potential users who might be willing to test and provide feedback on Co features and tools.
	Seek potential clients within education, government, civil society, industry and elsewhere who might pu services.
	Seek persons and groups within media, education, and elsewhere who might be interested in creating c theater, animations, movies, documentaries, art, or otherwise that would help promote the CogNarr visi
	Seek persons and groups within news media and elsewhere who might be interested in conducting writ interviews on topics central to CogNarr, or who might look to CogNarr staff for comments on related cu

		Develop potential use cases for CogNarr in a variety of different domains.
▼ **Analysis**	4	Develop whitepaper to better contrast CogNarr with other online collaborative decision-making tools ar
		Develop whitepaper on potential efforts or strategies that uncooperative or antagonistic users might en system or otherwise use CogNarr in a detrimental fashion. As part of this, discuss how information qual disinformation prevented or discouraged.
		Develop whitepaper on potential societal impacts of CogNarr, including benefits and dangers.
		Develop whitepaper on how CogNarr supports decentralized power, user equality, democratic principle
▼ **Software Development**	8	Create outline of CogNarr architecture, including computation, communication, networking, and securit resource needs.
		Identify and prioritize components and features of the initial MVP.
		Create description of development team, including roles and responsibilities. As part of this, discuss de
		Create a development roadmap and timeline, leading to and immediately following a MVP.
		Consider development partnerships, as needed, to create a MVP.
		Describe how the CogNarr research program informs software development, and the research prerequi development tasks.
		Develop a MVP.

Livestream #021 series (during 2021)

01

02

03

04

.1

.2

Courses

Below are the courses hosted at the Institute

🗣️ Active Inference for the Social Sciences ~ All 2023

🧭 Physics as Information Processing ~ Chris Fields ~ All 2023

Social Science course

In 2023, The Active Inference Institute hosted a course titled “Constructing cultural landscapes: Active Inference for the Social Sciences”, organized by Avel Guenin-Carlut, Ben White, Mahault Albarracin, Lorena Sganzerla and Daniel Friedman. The twelve-week course introduced participants to conceptual tools to understand the relation between social and cognitive sciences. Recordings of the talks, and more information are available at the public link.

https://coda.io/@active-inference-institute/active-inference-social-science-aii-2023

Active Inference for the Social Sciences ~ All 2023

https://coda.io/d/_dVVFg3pdihg

Physics course

In 2023, the The Active Inference Institute hosted a course titled “Physics as Information Processing”, taught by Chris Fields. The six-week course introduces participants to formal Quantum Information Theory as a concept and tool for understanding physical interaction as communication. Recorded lectures and course materials are available at the public link, and here is the YouTube playlist of all videos.

https://coda.io/@active-inference-institute/fields-physics-2023

Physics as Information Processing ~ Chris Fields ~ All 2023

https://coda.io/d/_dhwI_xbGzuD

Educational Standards & Qualifications

Engagement Pathways at the Active Inference Institute

Learning Paths, Modes, and Seasons

- **Browser**: Discovers active inference through key word searches, algorithmic recommendations, bibliographic tracing, or word-of-mouth, engaging with occasional content such as Production
- **Regular Consumer**: Follows dedicated channels and educational content about active inference and related topics
- **Active Learner**: Independently seeks out technical materials, research papers, and in-depth resources. Taking notes, making personal synthesis artifacts, engaging in solo or group Ecosystem Projects.
- **Institute Volunteer**: Participates in Institute Projects with a defined role (e.g. Textbook Group facilitator, Production .0 video preparation collaborator).
- Internship: Engages in focused project work while receiving mentored guidance and education.
- Fellows: Dedicated, possibly funded, focus on larger scale initiatives.
- Scientific Advisory Board, Board of Directors, Officers

The Active Inference Institute aims to make these learning pathways accessible to a global audience, meeting learners wherever they are in their journey. Through a multi-tiered approach, we aim to create entry points and paths for everyone from casual browsers to researchers and practitioners.

2025 Learning Initiatives

The Institute is enthusiastically preparing for expanded Education offerings in 2025. We recognize the growing interest in active inference across disciplines and are look to develop new Partnership s and Institute Programs to support learning needs and our Mission, Vision, Values, and Principles.

We will focus on building collaborative learning environments that bridge theoretical foundations with practical applications, while fostering a meaningful and productive community of practice that spans academic, industry, and independent researchers.

FarmWorks

https://zenodo.org/records/13754586
FarmWorks: Decentralized AI Agents for Personalized Solutions.

FarmWorks is the name of a project to develop a platform for human-AI interaction in agriculture, enabling personalized, farm-scale solutions that resist power concentrations associated with centralized AI systems.

A Grants submitted in September 2024.

Work continues in the RxInfer.jl Learning Group page at this link.

FarmWorks: Decentralized AI Agents for Personalized Solutions

Baulin, Vladimir[1]; Vyatkin, Alex[2]; GUÉNIN—CARLUT, Avel[3, 2]; Friedman, Daniel[2]; Bolt, John[2]; Falkenstein, Stefan[4]; Jassal, Parishrut[5]; Trois, Celio[1, 6]; Minchin, Jonathan[7]

Show affiliations

Project description submitted as part of application to Future of Life Institute - How to mitigate AI-driven power concentration

Climate change intensifies agricultural challenges, requiring more and more advanced technological solutions. Small farmers increasingly rely on technical assistance, which is becoming centralized, dominated by large agricultural corporations and governments imposing sophisticated pre-designated solutions. As AI proliferates within these centralized solutions, diseases mitigation methods, climate credits, government subsidies, and regulations risk monopolizing farmers' activities. This tendency, amplified by AI development, threatens to undermine farmers' autonomy and limit their ability to make independent decisions, converting them into consumers of centralized technological solutions.

We propose to develop FarmWorks — a platform for human-AI interaction in agriculture that enables personalized, farm-scale solutions while resisting power concentration associated with centralized AI systems. FarmWorks addresses the above challenges by providing an open source decentralized AI-powered agricultural platform that empowers individual farmers with cutting-edge technology while preserving their autonomy and promoting sustainable practices. By integrating real-time data collection, edge computing, and Active Inference models, FarmWorks enables farmers to make informed decisions tailored to their specific contexts (for example, integrating humidity data and epidemiological models to assist farmers with remediative and anticipatory treatments for mold).

Fundamentals of Active Inference

We worked with Sanjeev Namjoshi during 2023-2024 to support development of a textbook (expected public release in 2025). We look to share more information about future Fundamentals of Active Inference as we can.

For more on the book & Sanjeev's project, see:

Sanjeev Namjoshi ~ Active Inference Insights 018 ~ Education, Expectation-Maximisation, Evolution

https://www.youtube.com/watch?v=sAwPXw-WNg4

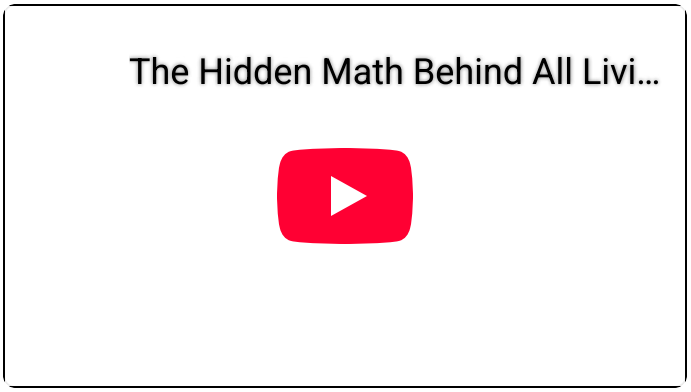

The Hidden Math Behind All Living Systems (on Machine Learning Street Talk)

https://youtu.be/hf18w6CuY8o?

Generalized Notation Notation

Generalized Notation Notation is a text-based language designed to standardize the representation and communication of Active Inference generative models. It aims to enhance clarity, reproducibility, and interoperability in the field of Active Inference and cognitive modeling.

Open Source code link: https://github.com/ActiveInferenceInstitute/GeneralizedNotationNotation

---

Original publication: Smékal, J., & Friedman, D. A. (2023). Generalized Notation Notation for Active Inference Models. Active Inference Journal. https://doi.org/10.5281/zenodo.7803328

GNN provides a structured and standardized way to describe complex cognitive models. It is designed to be:

- 👷 **Human-readable**: Easy to understand and use for researchers from diverse backgrounds
- 🤖 **Machine-parsable**: Can be processed by software tools for analysis, visualization, and code generation
- 🔄 **Interoperable**: Facilitates the exchange and reuse of models across different platforms and research groups
- 🔬 **Reproducible**: Enables precise replication of model specifications

GNN addresses the challenge of communicating Active Inference models, which are often described using a mix of natural language, mathematical equations, diagrams, and code. By offering a unified notation, GNN aims to streamline collaboration, improve model understanding, and accelerate research.

Active InferAnt Stream #014.1
Generalized Notation Notation: From Plaintext to Triple Play

Active InferAnt Stream #014.2
GNN for Generative Model Supply Chains: A Golden Spike Moment for Multiagent Trajectory Planning with RxInfer.jl

Active InferAnt Stream #014.3
The Sound of Uncertainty: Auditory Rendering of Generative Models in the Field of Streams

Knowledge Engineering

As of the end of 2022, Knowledge Engineering is a ongoing project (code repository) at the The Active Inference Institute that analyzes the literature related to Active Inference and Free Energy Principle, published as: *The Free Energy Principle & Active Inference: a Systematic Literature Analysis* https://zenodo.org/record/7449368

We performed a literature analysis of publications in scientific literature using the term “Free Energy Principle” or “Active Inference”, with an emphasis on works written by Karl J Friston. For a subset of papers with accessible full texts, we performed manual annotation (related to structural, visual, and mathematical features) and automated analyses (related to the terms in the Active Inference Institute’s Active Inference Ontology). The initial analysis here, at the scale of thousands of citations and hundreds of annotated papers, is presented as a first step towards the development of systems which could:

- Encompass increased scope of relevant works, including non-textual
- Integrate multiple forms of annotation and participation
- Facilitate integration of manual and artificial contributions
- Feature richer interfaces for use in learning & research
- Address field-specific local questions and provide transferable approaches
- Speak to broader questions in the history and philosophy of science

The paper is pre-printed at: https://zenodo.org/record/7449368

This project has an interactive Coda site and a Github repository.

The initial work was done in 2022 and we look forward to revisiting and improving this work in the years to come.

Knowledge Engineering Frontend

Want to print your doc?

This is not the way.

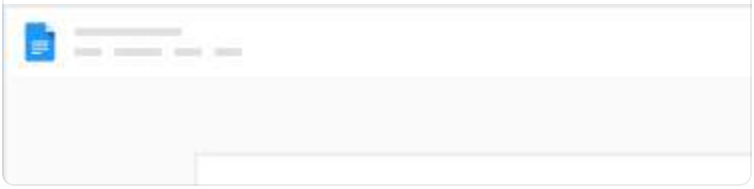

Obsidian Repository

Interactive Educational Resource for Active Inference

See 🎸 Interactive Obsidian Knowledge Base to interact with a structured repository of ♾️ Active Inference knowledge, connecting with many other background topics and areas.

The standalone link to the Obsidian website is:
https://publish.obsidian.md/active-inference/knowledge_base/cognitive/active_inference

The underlying 🌱 Open Source repository for this is https://github.com/ActiveInferenceInstitute/cognitive

Interactive Obsidian Knowledge Base

publish.obsidian.md

https://publish.obsidian.md/active-inference/knowledge_base/cognitive/active_inference

Myth of Objectivity

Myth of Objectivity Hypothesis (MOH)

> 📣 **Rigorous exploration of humanity's capacity for symbolic intelligence by modeling norms across social and identity scales, i.e. *morality*. The goal is to extend this framework in-silico to a, natively human value aligned, artificial general intelligence.**

Shagor Rahman, Research Fellows — Shaggy@activeinference.institute

> *"What to others a trifle appears*
> *Fills me full of smiles or tears;*
> *For double the vision my eyes do see,*
> *And a double vision is always with me..."*
>
> ***— William Blake***

Research Overview

This project investigates **how morality and symbolic thought co-evolved** through what I call the "Myth of Objectivity Hypothesis" - the idea that our capacity to model shared cultural expectations enabled humans to transcend individual perspectives and inhabit collective symbolic spaces.

Using multi-agent active inference simulations and transcendental model selection, my research explores how implicit and explicit moral beliefs form the foundation of our symbolic identities, examining how our boundless and enumerable selves manifest through narratives, institutions, and social paradigms.

This computational framework helps us understand not only human cultural evolution, but also how interventions can reshape symbolic systems - from religious prophets and cultural thought leaders to economic and technological changes. By formalizing the relationship between moral reasoning and symbolic cognition, this work offers crucial insights for both human uniqueness and the development of culturally-aligned artificial intelligence systems.

Resources

Active Inference Live Stream: GuestStream #061.1

"*Myth of Objectivity and the Origin of Symbols*" (Frontiers in Sociology)

"*Moral Case for Building AGI with Morality*" (Substack Blog)

Transcendental Model Selection {TmS}

Core to the theoretical basis is that human general and specifically symbolic cognitive abilities stem from our ability to innately model various social interactions and toggle between these various models.

Hierarchical Social Levels

Individual (ego), dyadic intimacy, anonymous bias-based groups (accent bias), and cultural identity integration

Model Selection as Moral Agency

The ability to activate broader (cultural) versus narrower (egocentric) modes of social engagement

Precision-Weighted Activation

Cultural beliefs constrain lower-level inferences through the precision parameter α

Symbolic Abstraction

Virtual cultural spaces necessary for symbolic modeling and communication

The godly etchings: Mathematic Formulation of {TmS}

The framework formalizes transcendental inference through hierarchical generative models where the deepest latent states generate precision of prior beliefs (α) about subordinate states. Free energy decomposes into accuracy and complexity terms, contextualized by precision across multiple social scales.

Traditional Accuracy←→Complexity Decomposition

$$F = \underbrace{-\mathbb{E}_Q[\ln P(o|s,m,\alpha)]}_{\text{Accuracy}} + \underbrace{\mathbb{D}_{KL}[Q(s,\alpha)||P(s,\alpha|o,m)]}_{\text{Complexity}}$$

Decomposition Using Transcendental Inference

$$
\begin{aligned}
F &= E_Q[\ln Q(s,\alpha) - \ln P(o,s,\alpha|m)] \\
&= \underbrace{E_Q[-\ln P(o|s,\alpha,m)]}_{\text{Accuracy}} + \underbrace{D_{KL}[Q(s,\alpha)||P(s,\alpha|o,m)]}_{\text{Complexity}} \\
&= \underbrace{E_Q[-\ln P(o|s^{(1)},\alpha,m)]}_{\text{Accuracy}} + \underbrace{E_{Q(o)}D_{KL}[Q(s^{(1)})||P(s^{(1)}|s^{(2)},\alpha,m)]}_{\text{Complexity: first-order}} \\
&+ \underbrace{E_{Q(o)}D_{KL}[Q(s^{(2)})||P(s^{(2)}|s^{(3)},\alpha,m)] + \ldots + E_{Q(o)}D_{KL}[Q(s^{(L)})||P(s^{(L)}|m)]}_{\text{Complexity: higher order}} \\
&+ \underbrace{E_{Q(o)}D_{KL}[Q(\alpha)||P(\alpha|s^{(L)},m)]}_{\text{Transcendental}}
\end{aligned}
$$

In this formulation, each order evaluates preferences and uncertainty minimization with respect to higher-order levels. The s^(n) variables correspond to increasingly complex levels of social organization:

- s^(1) represents individual behavior
- s^(2) corresponds to social or dyadic expectations
- s^(3) reflects expectations based on shibboleth or tag-based social signifiers
- s^(L) introduces a cultural level that abstracts expectations across various levels of social interactions

Transcendental Model Selection as Moral Agency

$$
\begin{aligned}
\ln \frac{P(m|o)}{P(m'|o)} &= \ln \frac{P(o|m)}{P(o|m')} + \ln \frac{P(m)}{P(m')} = \Delta F + \Delta G \\
\Delta F &= \ln P(o|m) - \ln P(o|m') \\
\Delta G &= \ln P(m) - \ln P(m') = G(m) - G(m')
\end{aligned}
$$

Model Selection: Agents select between narrower models (m^ego) and broader ones (m^culture), enabling inference to extend into greater timescales and underwriting policy selection and counterfactual planning.

Formalizing Gestalt vs Analytic Modes of Cognitive Attention

One of the key goals of this research is determine how transcendental model selection and transcendental inference can help formalize gestalt/analytic dichotomy in human cognition. This framework provides a computational account of how humans seamlessly switch between broad, holistic thinking and focused,

hierarchical analysis. This has been written about by many thinkers and recently by Ian McGhilchrist and John Vervaeke, the goal here is to view it in an Active Inference compliant perspective.

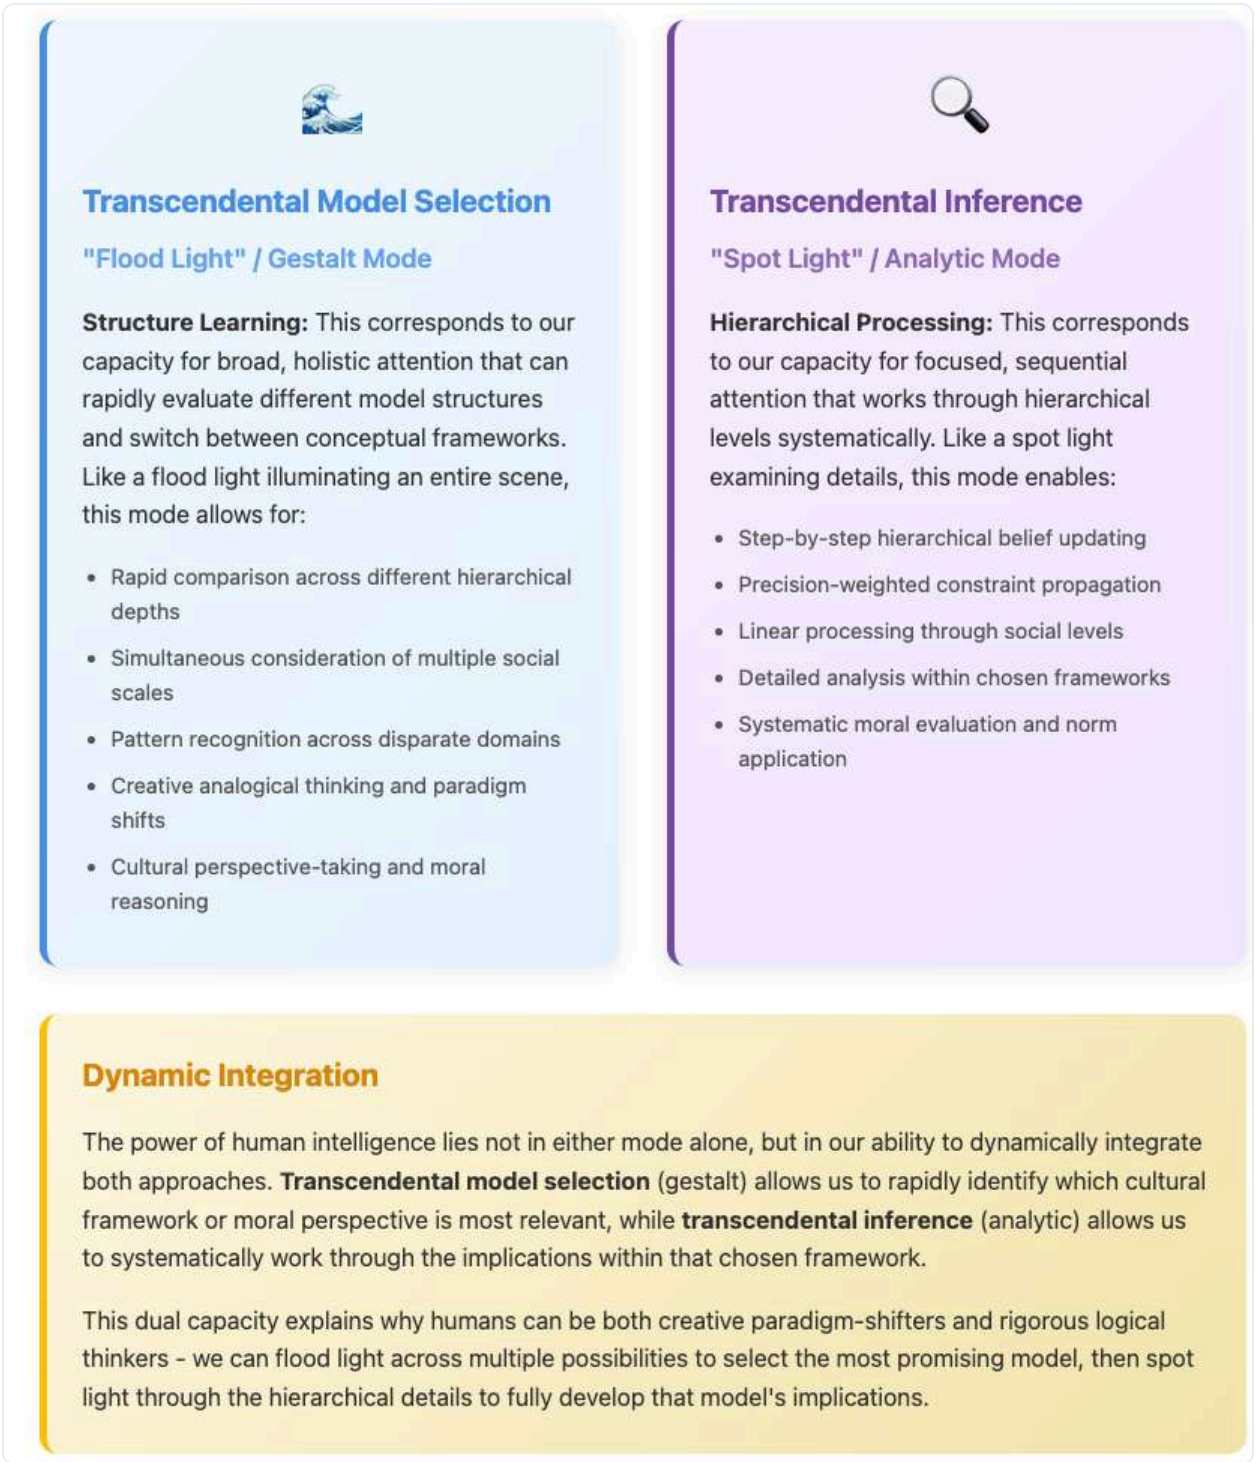

Selected Acts in Humanity's Objectively Mythic Journey

The research traces humanity's cognitive evolution through distinct phases of social organization and anonymity. This is critical for understanding important transitions in our own social-cultural contexts and the impact of particular intervening events such as important cultural icons, communication or mass media technology, or

political/institutional changes. Below are examples of important epochs but important is the ability to identify lower level shifts.

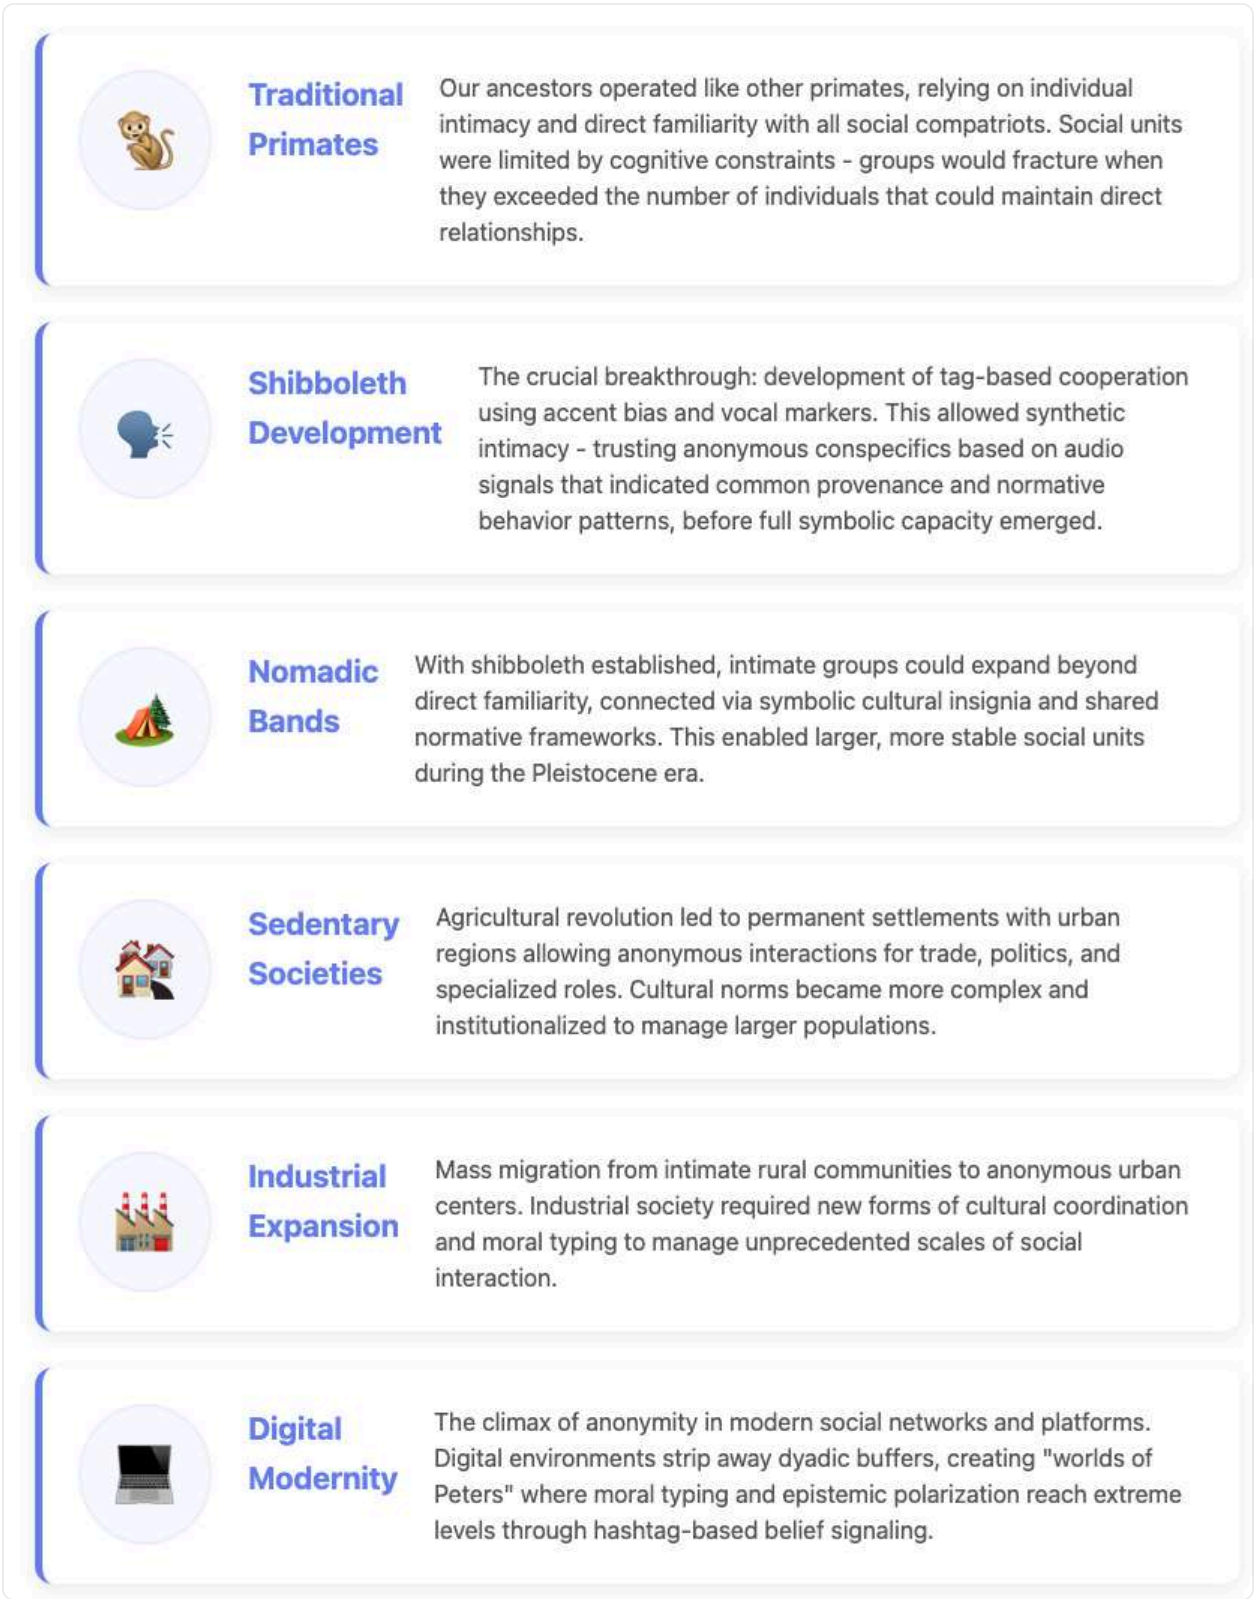

Empirical Predictions

1. **Developmental Trajectory:** Children should show increasing α values (cultural precision) with age, corresponding to greater moral conformity and symbolic capacity
2. **Neural Correlates:** fMRI studies should reveal distinct activation patterns during transcendental inference versus ordinary perspective-taking
3. **Cross-Cultural Variations:** Cultures should vary in hierarchical depth and precision coupling strength between levels
4. **Digital Communication:** Anonymous platforms without dyadic buffering will naturally produce more extreme moral typing and polarization

Shaggy's own Mythic Journey/ Road Map

Phase 1 (Months 1-6): Low-dependency prototypes, simulation development, narrative vignettes

Phase 2 (Months 7-18): Lightweight collaborations, mid-scale simulations, behavioral pilots

Phase 3 (Months 19-36): Resource-intensive studies, large-scale simulations, organizational pilots

Implications for Artificial General Intelligence

Hierarchical Organization

AGI systems must develop explicit mechanisms for traversing hierarchical scales, integrating multiple Markov blankets, and selecting appropriate levels of abstraction for different contexts.

Cultural Alignment

The precision parameter α that enables cultural-level constraints to guide individual behavior may be crucial for AGI systems operating across personal, social, and institutional scales.

Transcendental Model Selection

Current AI architectures lack the fundamental capacity for transcendental model selection, representing a core barrier to achieving human-level general intelligence.

Epistemic Depth

True alignment requires moving beyond black box architectures to systems with explicit epistemic depth - a capacity that is feasible to implement computationally.

About Shaggy

Email: Shaggy@activeinference.institute

Professional Background: Senior Product Manager, focus in Financial Services and Data/ML based systems.

Writing: Please follow/subscribe here: https://shaggy.substack.com/

Interested in: Collaborations on empirical validation, category theory formalization, and AI/AGI applications

Production

We produce educational content in the form of Production on YouTube, Podcasts on Podbean, and replication across other platforms.

We have multiple kinds of formats for the content, and we add new ones as availability/capacity arises.

Some current formats and area of focus for Production are:

- Livestream (focused on specific papers, 58 papers discussed from 2020 through 2025).
- 100+ GuestStreams, highlighting a wide range of work in Active Inference and related fields.
- ModelStream (computational models), OrgStream (social and organizational topics), MathStream (formalisms and math), MathArt streams in 2024, ArtStreams (art and aesthetics), InferAnt Streams (coding, modeling, synthetic intelligence), Roundtables (organizational updates), Courses (Social Science course, Physics course), Textbook Groups, and more
- One of the highlights of 2024 was **Active Inference Insights (Podbean, YouTube),** hosted by Darius Parvizi-Wayne.
    - Active Inference Insights is a podcast which introduces listeners to the wondrous land of Active Inference. Guided by our diverse array of guests, from physicists and mathematicians to cognitive scientists and philosophers, you will not only learn about cutting-edge theory, but also come to see the world in a whole new way, in which all things can be tied together by a single imperative: the minimisation of free energy.

To date, we have released over 500 videos, all available Open Source.

xWe aim for all videos are productions, to be transcribed, analyzed, translated, and published by the Active Inference Journal.

https://video.activeinference.institute/

Videos and Podcasts

Want to print your doc?

This is not the way.

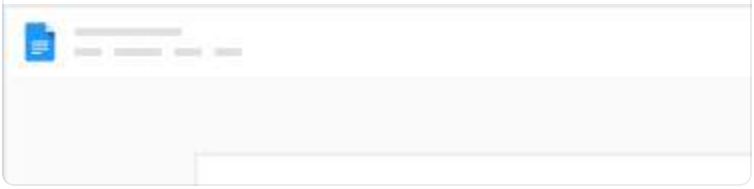

Research

Research activities and resources

The Research page hosts information on research projects such as CogNarr (Cognitive Narrative) Ecosystem: Facilitating Group Cognition at Scale, Wave Hypothesis, and other Research Resources .

Some research and products of Active Inference Institute and participants are below:

2024

- Karl J. Friston, Thomas Parr, Conor Heins, Axel Constant, Daniel Friedman, Takuya Isomura, Chris Fields, Tim Verbelen, Maxwell Ramstead, John Clippinger, Christopher D. Frith, Federated inference and belief sharing, Neuroscience & Biobehavioral Reviews, Volume 156, 2024, 105500, ISSN 0149-7634, https://doi.org/10.1016/j.neubiorev.2023.105500. https://www.sciencedirect.com/science/article/pii/S0149763423004694
- From the Broken link project:
  - "Aligning Active Inference Ontology to SUMO". June 3, 2024 (Collaboration) https://zenodo.org/records/11463326
  - "Aligning Spatial Web Terms to SUMO". April 24, 2024. https://zenodo.org/records/11062810
- Albarracin, M.; Pitliya, R.J.; St. Clere Smithe, T.; Friedman, D.A.; Friston, K.; Ramstead, M.J.D. Shared Protentions in Multi-Agent Active Inference. *Entropy* **2024**, *26*, 303. https://doi.org/10.3390/e26040303
- Friedman, D. A., & Tickles, D. (2024). Four-fold Fields of Quantum Dreams (Version v1). https://doi.org/10.5281/zenodo.10798145
- Dean Tickles, Daniel Ari Friedman, Why Paleolithic Rockstars were both enigmatic and sporadic: A comment on: 'Snakes and Ladders' in paleoanthropology: From cognitive surprise to skillfulness a million years ago, Physics of Life Reviews, Volume 50, 2024, Pages 4-6, ISSN 1571-0645, https://doi.org/10.1016/j.plrev.2024.04.010. https://www.sciencedirect.com/science/article/pii/S1571064524000447

2023:

- August 2023 publication from the Institute: "The Active Inference Institute and Active Inference Ecosystem".
- "Distributed Science - The Scientific Process as Multi-Scale Active Inference", Francesco Balzan, John Campbell, Karl Friston, Maxwell James Ramstead, Daniel Friedman, Axel Constant, 2023.
- "A variational synthesis of evolutionary and developmental dynamics", Karl Friston, Daniel Ari Friedman, Axel Constant, V. Bleu Knight, Thomas Parr, John O. Campbell, *Entropy,* 2023.
- "Generative Research Teams: Active Inference Compositions For Research and Meta-Science", Daniel Friedman & Jakub Smékal, 2023.
- "Generalized Notation Notation for Active Inference Models", Jakub Smékal & Daniel Friedman, 2023.

- "Active Inference for Autonomous Decision-Making with Contextual Multi-Armed Bandits", Shohei Wakayama and Nisar Ahmed, 2023.
- "
  A single-pheromone model accounts for empirical patterns of ant colony foraging previously modeled using two pheromones
  ", Eric Saund and Daniel Friedman, *Cognitive Systems Research,* 2023.

2022:

- "The Free Energy Principle & Active Inference: a Systematic Literature Analysis", Virginia Bleu Knight; RJ Cordes, Daniel Friedman. 2022.
- Catechism for: "Towards Active Diffusion: A Tale of Multiple (den)Cities", Jakub Smékal, Daniel Friedman. 2022.
- "
  Predictive Processing Interpretation of the Mirror Test and Implications of a Reflection Prediction for Human Cognition
  ", Sean O'Connor and Daniel Friedman. 2022.
- "Active Blockference: cadCAD with Active Inference for Cognitive Systems Modeling", Jakub Smékal, Arhan Choudhury, Amit Kumar Singh, Shady El Damaty & Daniel Ari Friedman, from IWAI 2022 (International Workshop on Active Inference).
- "
  An Active Inference Ontology for Decentralized Science: from Situated Sensemaking to the Epistemic Commons
  " & the Active Entity Ontology for Science

2021:

- Transcript of: Karl Friston, 1st Applied Active Inference Symposium https://zenodo.org/record/5797041
- "Active Inference in Modeling Conflict"
- "Narrative Information Management".
- Cognitio 2021 https://www.cognitio2021.com/
  - "Thinking like a State: Active inference and the deep roots of complex societies", Bleu Knight, https://osf.io/dxnzt/
  - "Intelligence without creativity: can Active Inference ground our understanding of life, cognition and society", Avel Guénin-Carlut
  - "Evolution of Latent Model for Collective Cognition", Amit Singh
- "Active InferAnts: The basis for an active inference framework for ant colony behavior", paper www.frontiersin.org/articles/10.3389/fnbeh.2021.647732/
- 2nd International Workshop on Active Inference "Active Inference & Behavior Engineering for Teams", poster on the 2020 paper "Active Inference & Behavior Engineering for Teams" https://zenodo.org/record/4021163

- "Collective Intelligence as Latent Imagination", Amit Singh, International Conference on Cognitive Modeling (ICCM'21), 73
- "Context Switching in Machine Minds", Amit Singh, Society of Mathematical Psychology 2021, https://youtu.be/4647USeygmg

2020:

- "Active Inference & Behavior Engineering for Teams" Alex Vyatkin, Ivan Metelkin, Alexandra Mikhailova, RJ Cordes, Daniel Friedman

Research Resources

Research Projects & Resources

See Control Flow for resources on:
“Control flow in active inference systems Part I: Classical and quantum formulations of active inference"
“Control flow in active inference systems Part II: Tensor networks as general models of control flow"

See Variational Evolution for resources on:
"A Variational Synthesis of Evolutionary and Developmental Dynamics" (2023)

Control Flow

Resources for: Control flow in active inference systems

Supplementary Materials for:

Chris Fields et al., "**Control flow in active inference systems Part I: Classical and quantum formulations of active inference,**" in IEEE Transactions on Molecular, Biological and Multi-Scale Communications, doi: 10.1109/TMBMC.2023.3272150. https://ieeexplore.ieee.org/document/10113698 (2023)

Chris Fields et al., "**Control flow in active inference systems Part II: Tensor networks as general models of control flow,**" in IEEE Transactions on Molecular, Biological and Multi-Scale Communications, doi: 10.1109/TMBMC.2023.3272158. https://ieeexplore.ieee.org/document/10113744 (2023)

Control_Flow_Supplementary-Information-Table-1.pdf

Variational Evolution

Friston, K.; Friedman, D.A.; Constant, A.; Knight, V.B.; Fields, C.; Parr, T.; Campbell, J.O. A Variational Synthesis of Evolutionary and Developmental Dynamics. *Entropy* **2023**, *25*, 964. https://doi.org/10.3390/e25070964

Resources for "A variational synthesis of evolutionary and developmental dynamics" are hosted at:
https://coda.io/@active-inference-institute/fep-evolution
And embedded in Variational Evolution

Variational Evolution

Want to print your doc?

This is not the way.

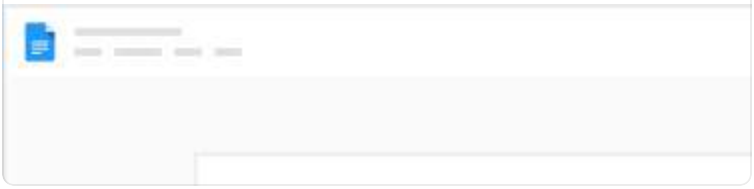

Robotics & Embodied

https://coda.io/d/_dl3nNmei2EF

RxInfer.jl Learning Group

Documentation for the RxInfer.jl Learning Group can be found at this link.

The RxInfer project at the Active Inference Institute represents a vibrant Open Source collaboration focused on learning and applying RxInfer.jl, a cutting-edge Julia package for automated Bayesian inference and active inference modeling

Community Structure

The project maintains regular synchronous meetings every Thursday at 13 UTC in the Institute Discord, bringing together a diverse group of participants including developers, researchers, and educators. See all meetings here. Over 2024, the meetings included over 20 active contributors with varying backgrounds and expertise levels in Julia programming and Bayesian modeling. The community is working on many innovative applications.

Educational Resources

The project emphasizes knowledge sharing through:

- Comprehensive documentation and examples
- Regular meetings featuring technical presentations
- Hands-on tutorials and code workshops (also see LearnableLoop)

Development Roadmap

Key technical advances planned for 2024-2025 include:

- Improved scaffolding of educational examples for training and learning use
- Nested models with GraphPPL.jl
- Enhanced graph structure visualization
- Automated inference with ExponentialFamilyProjection.jl
- Implementation of stochastic processes
- Improvements in robustness and memory efficiency

The RxInfer.jl Learning Group project exemplifies open-source scientific software development, combining rigorous technical advancement with inclusive community participation in learning and Grants. Through its regular meetings, diverse project portfolio, and commitment to education, it continues to push the boundaries of active inference implementation while maintaining accessibility to newcomers in the field

RxInfer 2024 ~ Active Inference Institute

https://coda.io/d/_ddtS-XZ4BJb

Seasonal School

Active Inference Seasonal School

More information on ☼ Seasonal School and in-person experiences to come ~

Start

Onboarding Curriculum and Learning Paths for Active Inference, across languages and backgrounds

https://github.com/ActiveInferenceInstitute/Start is Open Source @Software Development project, started in December 2024, introduced in video Production Active InferAnt Stream 008.1 “Symbol’s Greetings: Onboarding to Active Inference across backgrounds & languages”, and continued in Active InferAnt Stream 015.1 “START/HERE: A Map & What Might Happen Next”

Start here: https://github.com/ActiveInferenceInstitute/Start/blob/main/here.md

START (Scalable, Tailored Active-inference Research & Training)

Start (Scalable, Tailored Active-inference Research & Training) is a modular pipeline that generates high-quality educational materials on Active Inference and the Free Energy Principle, tailored to professional domains and individual learners. It integrates live web research via Perplexity and advanced LLMs via OpenRouter to produce evidence-based, professionally structured content, with visual analytics and multilingual localization. The system emphasizes real data, reproducibility, and quality assurance through tests and linting, and it ships with comprehensive documentation and an interactive CLI experience.

> **Live docs site: https://activeinferenceinstitute.github.io/Start/**
>
> **Repository: https://github.com/ActiveInferenceInstitute/Start**

- What it produces
  - Domain research: 3,000–5,000 word analyses of professional fields (e.g., neuroscience, AI, healthcare), grounded in current sources (Perplexity).
  - Audience/entity research: 5,000–8,000 word learner profiles with actionable learning strategies.
  - Curricula: 40–60 hour, module-based programs with structured sections, objectives, and assessments.
  - Visualizations: PNG charts of curriculum metrics and Mermaid flow diagrams for structure and learning pathways.
  - Translations: Native-quality, culturally adapted outputs in 11+ languages (Chinese, Spanish, Arabic, Hindi, French, Japanese, Russian, Swahili, Tagalog, and more).
- How it works (pipeline)
  - Inputs: YAML configurations define domains, entities/audiences, and target languages.
  - Research → Curriculum Generation → Visualization → Translation → Outputs under data/.
  - Prompts are curated for domain analysis, curriculum generation, section authoring, and translation, ensuring consistent structure and completeness.
- Key technologies and quality model

    - Research: Perplexity API for current, multi-perspective domain insights.
    - Content: OpenRouter LLMs for robust, structured curricula and translations.
    - Quality: pytest coverage, ruff linting, black formatting, type hints, and CI-ready workflows.
    - Architecture: Clear separation across src/ (core logic), learning/ (scripts), data/ (artifacts and configs), docs/ (user/developer guides), and tests/.
- Running the system
    - Interactive terminal: run.sh provides an end-to-end guided experience from research through translation.
    - Documentation workflows: run_docs.sh supports serve, build, and deploy to GitHub Pages.
    - Python environment uses uv for reproducible setup; keys for Perplexity and OpenRouter are required.
    - Outputs accumulate incrementally in data/domain_research/, data/audience_research/, data/written_curriculums/, data/visualizations/, and data/translated_curriculums/.
- Who it is for
    - Educators and program designers building university or professional development courses.
    - Researchers seeking personalized, evidence-based learning roadmaps and current domain syntheses.
    - Organizations adopting Active Inference frameworks for training, strategy, and decision support.

Systems Approach

A modern third-generation systems approach is essential for managing today's complex adaptive systems. This approach transcends traditional linear models, embracing continuous evolution, interactivity across scales, and the dynamic, constructivist perspective where systems actively reshape themselves in response to changing conditions. Combining active inference principles with a systems approach provides a pathway to designing resilient, self-organizing systems that are responsive to diverse environments.

Evolving systems approach

Classical systems approach focused on structured relationships within defined boundaries, effective for static and predictable systems. However, with the growing complexity of modern systems like cyber-physical networks and adaptive ecosystems — a more dynamic approach is needed. Third-generation systems approach builds on these foundations to handle layered, open-ended development and real-time adaptability. A systems approach today emphasizes flexible frameworks that enable continuous learning and adaptation across varied contexts, positioning systems to better navigate and anticipate change.

∞ Active Inference provides a foundation for adaptive systems by defining systems as (nested, interacting) agents. This aligns with a systems approach by enabling systems to organize themselves dynamically in response to environmental changes. Key characteristics of this model include:

- Continuous adaptation: Systems evolve iteratively, continuously integrating new information rather than following a rigid lifecycle.
- Anticipatory action: Systems use predictive models to take preemptive actions, reducing disruptions before they occur.
- Interactions within and across scales: Systems function cohesively across micro and macro levels, preserving stability and coherence regardless of scale.

These qualities position active inference as a crucial tool for developing systems that are resilient, responsive, and able to self-correct in changing environments.

Systems as Constructors

A modern systems approach treats systems as constructors, entities that not only adapt but also actively build and modify their environments. Systems continuously refine their models based on feedback, supporting informed decision-making and efficient resource allocation. This constructivist perspective emphasizes:

- Dynamic modeling: Systems adjust internal models based on ongoing sensory input, which helps them make real-time, informed decisions.
- Open-ended development: Systems remain open to generating novel solutions and can reorganize to meet emerging challenges, enhancing robustness and resilience.

This approach is especially applicable in fields requiring systems to maintain functionality amid complex, changing conditions, like AI and distributed cybersecurity. Systems built on this constructivist foundation are inherently flexible, robust, and capable of evolving independently.

Collaborative ecosystems and community-driven development

The active inference framework is built within a collaborative, Open Source development model (The Active Inference Ecosystem) that aligns with third-generation systems approach. This community-driven ecosystem encourages knowledge sharing and real-time updates, ensuring that systems evolve alongside technological and societal needs. Collaborative development fosters rapid adaptation and inclusivity, allowing systems to better meet diverse user requirements. An open-source model also supports common standards, providing a strong foundation for sustainable and accessible system design across interconnected fields.

Adaptive systems approach

Integrating active inference within a modern systems approach offers a robust, adaptive model for managing complexity. This combination encourages resilient, coherent, and evolving systems that can operate autonomously and flexibly across scales. By embracing dynamic modeling, constructivist principles, and active inference, this approach provides a foundation for systems that not only withstand change but actively respond to it, supporting a broad range of applications in both technical and social domains.

Reference “Toward an Ontology for Third Generation Systems Thinking” by Anatoly Levenchuk

Textbook Group

This form registers you to participate in the Active Inference Textbook Group on the book "Active Inference: The Free Energy Principle in Mind, Brain, and Behavior" By Thomas Parr, Giovanni Pezzulo and Karl J. Friston (2022).

Enter the Textbook Group's interactive document here, with pages for questions, resources, past meetings, and more.

The Textbook Group is about **learning Active Inference in an open science setting**. All backgrounds and level/type of familiarity with Active Inference are welcomed and encouraged!

See the completed playlists of Cohort 1, Cohort 2, Cohort 3, Cohort 4, Cohort 5, Cohort 6, Cohort 7.

The main focus of the Textbook Group is to help you learn Active Inference. We're expecting lots of different backgrounds, but our goal is to meet you where you're at to help you understand the textbook.

There will be no wrong answers or incorrect questions. You'll be encouraged to make connections with what's familiar and authentic to you. The textbook includes connections to biology, psychology, physics, mathematics, computer programming, etc.

Group facilitators and participants will be actively maintaining and updating the Coda, which we use as a shared epistemic niche. Facilitators will be available to answer questions and connect you with other participants to compliment and reinforce learning. If you're interested in facilitating please indicate in the form above.

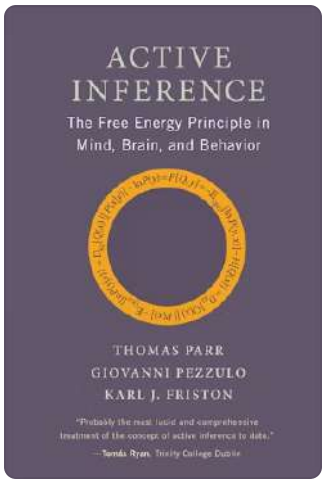

The Institute is exploring exciting research and applications of Active Inference. The Textbook Group is a great place to learn more about this. If you're interested to learn more, please complete the form below or email us.

To register, complete the form below!

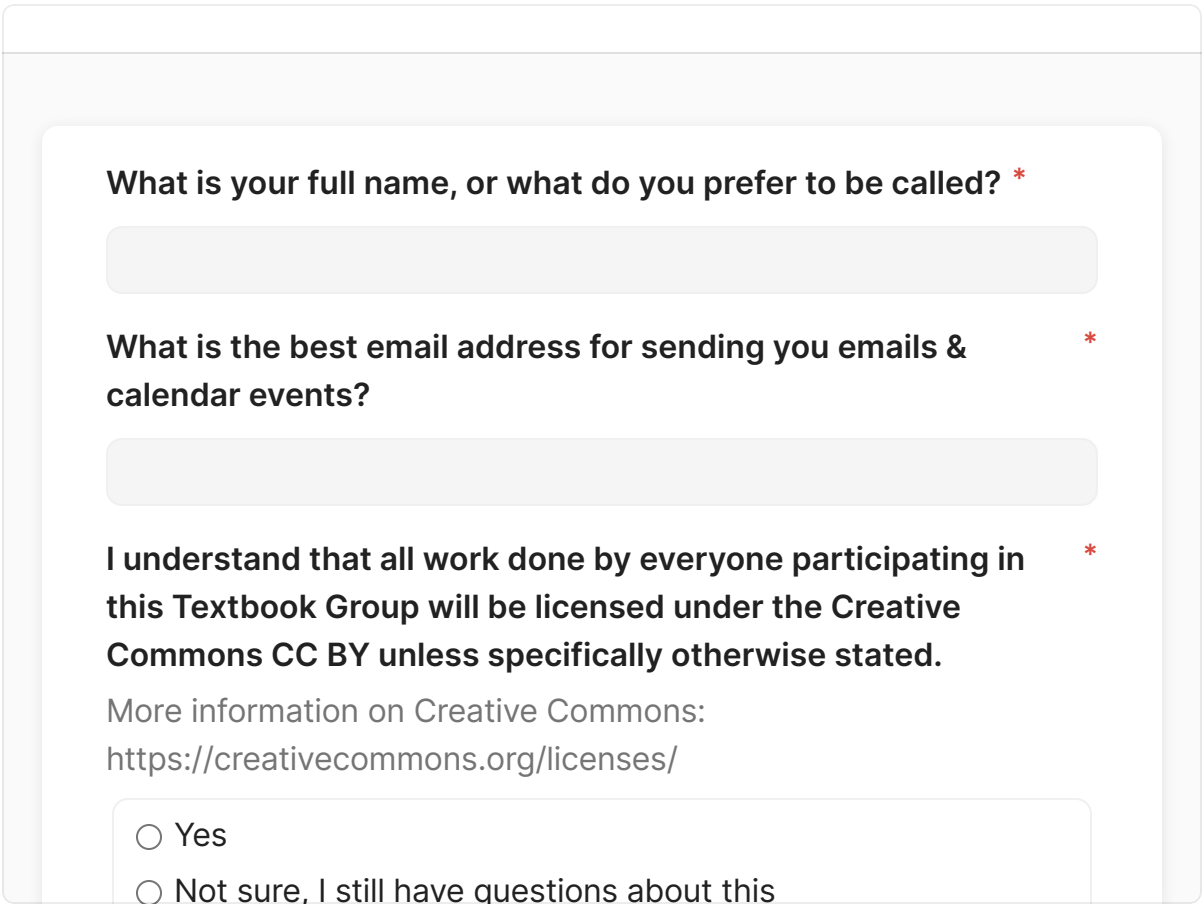

- This Onboarding will be an email from Blanket@ActiveInference.Institute , containing a link to the Coda document that will be the single source of truth for this Textbook Group cohort.
- The Coda will have supporting material and learning practices to understand each chapter, information about how to contribute, as well as information on the calendar of the Textbook Group. Everyone will have an individual learning space, so you can easily share your work, collaborate with others, and get help.

Tech Tree

A current area of interest for us is creating an Active Inference 🌳 Tech Tree (a “tool to map science and tech”) to guide 🌱 Open Source development.

For now, work on this can be found within the 🚀 RxInfer.jl Learning Group documentation document here, and in the Github repository here https://github.com/ActiveInferenceInstitute/Symposium/tree/main/output , where we are processing public participant information for the 🎓 Applied Active Inference Symposium and applying LLM methods to this.

Theoretical Neurobiology (TNB) Group

Theoretical Neurobiology (TNB) Group

Objective

The TNB Group has been fostering interdisciplinary research and collaboration for decades. Our mission is to advance the understanding and application of active inference, a theoretical framework developed by Prof. Karl Friston. This is achieved through regular online meetings featuring presentations and discussions, which may include empirical data and its analysis, simulations, and mathematical development. We welcome contributions and perspectives from diverse fields, including neuroscience, mathematics, machine learning, psychology, philosophy, medicine, and biology.

Recordings of past meetings, organised by research area, can be found on the **TNB YouTube Channel**. Note that we are gradually adding more videos from our archive alongside recently recorded sessions.

Meeting Details

- **Schedule**: Mondays and Tuesdays, 2:30 pm (UK time)
- **Duration**: ~2 hours
- **Structure**:
  - ~40-minute presentation
  - ~40 minutes of Q&A
  - ~40 minutes of discussion and feedback with Prof. Friston
- **Frequency**: Weekly
- **Platform**: Zoom - the meeting link is sent via our mailing list (email us at theoreticalneurobiology@gmail.com to join the mailing list)

How to Participate

Our meetings are open to researchers, students, and professionals worldwide. With no membership fees, we provide a relaxed, no-pressure environment for engagement, whether through active participation or as an observer. You are welcome to join any session that interests you.

To join our mailing list for updates on upcoming presentations, active inference events, or job opportunities — or to request to present your work — email us at **theoreticalneurobiology@gmail.com**. Please note that presentation slots typically book two to three months in advance.

Chairs

Riddhi J. Pitliya, PhD.

My research focuses on active inference, human cognition, and multi-agent systems. I completed my PhD at the University of Oxford, where I investigated how individuals infer causal structures and agency in their environments, particularly across the depression spectrum. We found that depressive symptoms are linked to reduced sensitivity to inhibitory causal relations, reduced perceptions of agency, and a tendency to engage in frequent but less goal-directed actions when learning about causal structures. Currently, I work at VERSES in the Intelligent Systems Lab, where I develop computational models of theory of mind, leveraging active inference to facilitate collaboration and competition among multiple agents.

Miguel De Llanza Varona

I'm currently a PhD student at the University of Sussex under the supervision of Christopher L. Buckley and Anil Seth. My research lies at the intersection of AI and theoretical neuroscience where I explore the theoretical underpinnings of representation learning in bounded rational agents. My main research interests are twofold: first, how cognitive constraints (e.g., metabolic costs or limited memory) interfere with optimal Bayesian inference; and second, what are the challenges of learning representations in service of reconstructing the data in misspecified generative models (e.g., VAEs).

Peter Thestrup Waade

My research focuses on computational cognitive modelling of multi-scale social interaction, particularly from the perspective of active inference and predictive processing. I did my PhD with Chris Mathys at the Interacting

Minds Centre at Aarhus University, and am starting a postdoc position at the Translational Neuromodelling Unit at ETH Zürich with Klaas Stephan. I develop Julia software for cognitive modelling in general, the Hierarchical Gaussian Filter and active inference with POMDP's - I also do some work in consciousness research, on joint action in partner dancing and on Chinese philosophy and predictive processing.

Robert Chis-Ciure, PhD.

I'm an ERC postdoctoral research fellow in Anil Seth's lab at the University of Sussex and Sussex Centre for Consciousness Science. Our research focuses on formalised notions of emergence and computational neurophenomenology. We're using hybrid predictive coding and active inference formalisms to model various phenomenal properties of experience and validate them experimentally. In doing this, we're building toward a new methodological paradigm, *Phenomenomics,* to comprehensively characterise the "inner worlds" of human and, eventually, all other observers—their *phenomenome*—by also leveraging AI/ML strategies on large scale datasets. Before Sussex, I was a Fulbright postdoc at NYU under David Chalmers, a Tatiana Foundation postdoc in Georg Northoff's lab, and a Fulbright Ph.D. student in Giulio Tononi's lab, working on consciousness at the intersection of philosophy, neuroscience, and computational modelling. In my free time, I do various projects as an affiliated researcher at the Wolfram Institute.

Will Yun-Farmbrough

I am a PhD student at the Sussex Centre for Consciousness Science, supervised by Anil Seth and Chris Buckley. My research investigates how perceptual phenomenology in human subjects can constrain and inform predictive coding models of cortical processing — what are the algorithmic underpinnings for how our world appears to us? I am also interested in predictive processing approaches to the meta-problem of consciousness, seeking to understand how intuitions of conscious experience and qualia might arise naturally in certain generative model hierarchies. I enjoy surfing, houseplants, and zen.

Work in Progress

1. **Repository of Active Inference Resources**
   We are compiling a comprehensive repository of resources on active inference. This will include introductory materials, research papers, and tutorials, all categorised for easy navigation.
2. **Meeting Recordings**
   We are working on publishing recordings of past and future meetings online.
3. **Joint Research Symposia**
   We are planning on organising joint symposia with other research communities to strengthen cross-community interaction and collaboration.

Wave Hypothesis

During 2024, Robert Worden (http://www.bayeslanguage.org/bb/intro.html) presented a series of livestreams at the Active Inference Institute, related to the papers:

- Robert Worden, "The Projective Wave Theory of Consciousness" (2024) (https://arxiv.org/abs/2405.12071)
- Robert Worden, "Assessing the Brain Wave Hypothesis: Call for Commentary" (2024) (https://arxiv.org/abs/2408.04636)
- Robert Worden, "Testing the Brain Wave Hypothesis" (2024) (https://arxiv.org/abs/2408.05368)

GuestStream #082.1

Bayesian Model-Based Cognition: The Requirement Equation

GuestStream #082.2

Three-dimensional Spatial Cognition: Bees and Bats

GuestStream #082.3

The Projective Wave Theory of Consciousness

David Rudrauf, Kenneth Williford, Karl Friston.

GuestStream #082.4

Assessment of the Brain Wave Hypothesis

GuestStream #082.5

Computers, Meaning and Consciousness

GuestStream #082.6

A Unified Theory of Language

GuestStream #082.7

Language and Human Nature

On the Wave Hypothesis please email commentary to rpworden@me.com .
If you wish, your comments will be posted at www.bayeslanguage.org/bb/commentary.html .

Karl Friston writes:

> "It would be very useful to see other people's take on this proposal for mortal computation in the brain."

The Active Inference Ecosystem

The The Active Inference Ecosystem is a vibrant, global community of researchers, practitioners, and enthusiasts united by their interest in Active Inference — a powerful framework for understanding cognition, behavior, and complex adaptive systems. The ecosystem extends far beyond the formal boundaries of the The Active Inference Institute, encompassing a wide array of individuals, organizations, and projects that contribute to the development and application of Active Inference across Domains of Application.

At its core, the Active Inference Ecosystem is characterized by its open, collaborative nature. It brings together experts from fields as varied as neuroscience, artificial intelligence, philosophy, physics, and social sciences, fostering cross-pollination of ideas and innovative approaches to complex problems. The ecosystem thrives on the collective efforts of its participants, who engage in research, education, software development, and practical applications of Active Inference principles.

The ecosystem is not just an academic or theoretical construct; it is a living, evolving network of interactions and initiatives. It includes Partnerships among organizations, educational programs, Open Source products, events like the Applied Active Inference Symposium, and various community-driven efforts. The Active Inference Institute serves as a hub within this ecosystem, providing infrastructure, coordination, and support to facilitate the growth and impact of Active Inference across disciplines and sectors (see History of The Institute for how this has unfolded over the years).

As the document transitions into detailing the Active Inference Ecosystem, readers can expect to explore the Ecosystem Priorities and Challenge Areas, Ecosystem Development: Structure and Growth, and Ecosystem Projects across Domains of Application.

Ecosystem Priorities and Challenge Areas

We look to continued engagement with the Ecosystem, to better curate and refine the
Ecosystem Priorities and Challenge Areas.

Education: Scientific Literacy and Workforce Development

Active Inference relies on mathematical formalisms and is loaded with abstract conceptual challenges that transcend disciplinary boundaries. We hope to model educational processes such as pedagogy, competency evaluation, and professionalization in Active Inference. Thus, the Institute catalyzes workforce development, seeks to stabilize the "research to practice" gap, and contributes to the broader project of participation in scientific ecosystems.

Research: Grounding the Cognitive Sciences in Physics

Research across the natural sciences suffers from a lack of theoretical integration and practical collaborations. Active Inference is gaining traction as a rigorous attempt at a unifying first-principles accounts of vital features of biological systems, transcending disciplinary boundaries. At The Institute we promote this theoretical integration through various educational programs, supporting learners of all backgrounds.

Information Science and Diverse Intelligences

The interaction frequencies of modern information environments are higher and more complex than ever. At The Institute we apply Active Inference to understanding, monitoring, evaluating, refining, and developing artificial and synthetic (e.g., human-machine interface, organizational, crowd) intelligence systems. In this way, active inference helps to identify, analyze and optimize various forms of "interoperability" across various forms of intelligent system, making possible a form of "mutual socialization" among such systems. This work is enacted by projects currently related to information science, ontology, data quality control, artificial intelligence explainability, and knowledge engineering.

User Experience, Accessibility, and Sociotechnical Design

It remains an open challenge how to most effectively, efficiently and fairly enable sustainable engagement in digital systems consistent with all parties expectations and needs. At The Institute we map cognitive frameworks as a framing for design, user experience, ergonomics, and requirements engineering, as well as implementation and operational guidance, to offer new methods and tools to a wider community of professionals and scholars.

Business Applications

Business and commercial interactions are typically characterized by party attention to reduced set of abstracted variables as compared with biological and social systems. Notwithstanding the "management" and regulation or variables, active inference can still help to improve the competitive insights and risk mitigation strategies and other variables that are the focus of business and commercial parties. Active inference research and analysis promises to substantially enhance and improve critical business elements such as risk strategies, insurance markets, banking (lending criteria), identity authentication, and authorization and a host of other business interaction decisions.

Social Welfare

The scale independence of active inference analysis causes it to be well suited to framing issues in settings where different parties experience different levels of information and resources. This includes various programs of local and global social welfare that seeks to enhance the local and global fairness of resource allocations of various kinds and to offer a pathway to easing the consequent burdens that unbalanced resource related interactions place on precarious populations.

Cyber and Cognitive Security

Individuals and organizations today are confronted with a rapidly-evolving landscape of threats to digital and cognitive security. At The Institute we work to unify cognitive frameworks with existing cyber security and emerging cognitive security concepts and frameworks, to understand, measure, and address local and global information technology risks and impacts more effectively at multiple scales.

Scaling the Active Inference Ecosystem

The nascency of the Active Inference Ecosystem enables us to take a proactive approach towards various areas of consideration. At The Institute we create synergy among the efforts applied to the above challenge areas, and emerging needs of the Active Inference Ecosystem. This approach creates an opportunity to learn by doing and to embrace convergence research, where implementations are developed in parallel with theory, supported by regular information sharing and collaboration among practitioners and researchers.

Applying Active Inference

The Institute brings insights from empirical and theoretical Active Inference research into practice by designing new projects or communicating with existing projects that design and implement social system infrastructure, such as health infrastructure and cultural technologies that support human well-being. We also support 🌳 Ecosystem Projects that design and implement solutions to various collective problems, such as climate change, threats to democracy, armed conflict, or overall polycrisis.

Additional Focus Areas

- Software usability and accessibility
- Information system optimization and efficiency
- Cultural heritage and progress
- Legal and regulatory consistency and compliance

Ecosystem Development: Structure and Growth

Community Growth and Development

Here, we present the community growth and development model for The Active Inference Ecosystem, built on the following 5 core components:

1. **Awareness**. Promoting and fostering awareness and use of Active Inference, and developing partnerships with well-aligned organizations and communities.
2. Education. Developing and disseminating educational materials, contributing to competency, capability, and common language within the community.
3. **Common Forum**. Offering and maintaining an inclusive and accessible common forum for discussing, sharing, and hosting relevant work and opportunities, finding collaborators, and networking (i.e., an informational commons).
4. **Support**. Providing support for emergent teams and projects which align with The Institute's mission, in the interest of innovation and impact
5. **Governance**. Maintaining stable governance for cultivating and sustaining partnerships, technical infrastructure, and sponsors.

Ecosystem Structure

The Institute cultivates an active and engaged ecosystem around the scientific modeling framework of Active Inference. This vibrant Ecosystem and community drives innovation on the research front and makes significant strides in providing accessible education. The Institute ensures that efforts are well-aligned, impact-focused, relevant, and meaningful in advancing research and education for the betterment of society by forming partnerships and by engaging with and growing the Active Inference community. Our community development model emphasizes facilitation over management, and distributed as opposed to command-and-control strategies. More importantly, our model moves beyond the provision of networking and discussion space to support emergent, collaborative work.

In these regards, the Institute functions as a seed crystal that can help to foster phase changes across a variety of information system domains and applications. The Institute does not directly manage all of the systems upon which it has an influence, but instead seeks to leverage its influence by providing coherent multiple tools and practices from which communities of shared interest can optimize their local information system dependencies for active inference efficiencies.

As opposed to a linear "funnel" growth model, The Institute will implement a cyclical model of organic growth pursued through the incubation among participants of (i) self-efficacy, or a sense of personal capability, (ii) a sense of support and safety, and (iii) a sense of investment and impact in participants, as a basis for forming a sense of community and providing the foundation for development of relationships within the community through positive, repeated contact. The support of these senses leads to productive, emergent collaboration, which in turn leads to emergent community narrative, norms, roles, and "scripts". Participants are reinforced in their feelings of capability as a part of a team, assured that they will be provided with support in a reasonably safe environment, and that results will have a lasting, positive impact on their community. Resulting research and educational artifacts and documentation constitute shareable content which can then be used to bring awareness about Active Inference and The Institute to non-community members.

Where a "funnel growth" model focuses on awareness alone as a basis for developing a user-base, our model's focus on education, knowledge sharing and presentation of work, and support for teams allows for non-community members of all backgrounds and interests to engage with and contribute to the community, thus affirming membership through a sense of shared investment, impact, and competency. Further, where online learning communities anticipate members terminating participation following completion of coursework (or after achieving feelings of self-efficacy in the material), our model's provision of support and opportunities for sharing of work with professionals and academics provides incentives for continued engagement and participation to those who feel they have already become reasonably familiar with all available educational material.

Below, background is provided on the (i) structure of the community (i.e., user segmentation), (ii) our information storage and dissemination technology ("tech") stack, (iii) our communications plan, (iv) the education, support, and infrastructure and governance functions we provide and/or intend to provide as a part of this model, and (v) our intended approach toward evaluating quality control and growth.

Community Participants

The Institute is a formal organization that has been constituted to serve some of the organizational and operational needs of the expanding active inference ecosystem. The Institute and its staff recognize that the energy and knowledge value relating to the further understanding and development of active inference resides in the broad active inference community, which is supported, fostered, convened and cross fertilized through the activities of The Institute. The reach and potential implications of active inference across domains and sectors is sufficiently broad that parties can choose from among many different ways to engage. A partial list of categories of participation is presented below to provide a sense of the variety of participants.

Direct Institute Participants (Members and Learners)

Many individual participants interact directly with The Institute and its resources and programs. Participants include members of the Active Inference Ecosystem, or those who engage directly with and contribute to Institute Programs. These participants include students, educators, researchers, and professionals from around the world who may benefit either from awareness of Active Inference and its implications, developing related competencies and having opportunities to network and collaborate with individuals who do, or from opportunities to collaborate and share work and insights which would be valuable to the Active Inference Ecosystem.

Participants also include *learners* at various levels of involvement and expertise that engage directly with The Institute as part of their learning process. The Institute seeks to support all learners, from the academic expert to those individuals who are not, and everyone in between. The Institute seeks to facilitate access by all learners to tools and materials and narratives that can help people at all levels access information that can help them to enjoy the direct and indirect benefits of active inference thinking and approaches.

Users (Adopters and Beneficiaries)

For individual and organizational users that explicitly adopt Active Inference-based [organization and operation] of their information processing and synthetic intelligence practices, policies and tasks, the Institute's productive outputs provide support and opportunities for engagement with a broader community. The Institute maintains an online resource center that includes software, tools, and materials that convey methodologies and practical pathways for instantiating Active Inference-derived structures in a variety of community settings and institutional contexts, and includes [practical suggestions for] the facilitation of Active Inference itself as an open source and open standards set of products and practices. As such, the community using Active Inference and related Open Source products requires documentation, clear messaging regarding updates, and guidelines on fair and best practices. By considering such beneficiaries of Active Inference as "users," The Institute may leverage existing best practices from other domains, such as user experience, requirements engineering, and software engineering. Potential users include professionals, researchers, educators, and engineers.

Beyond direct "users" of active inference, there are many groups of parties that benefit from the use of active inference who won't interact directly with such systems, nor be aware of it. Comparison is made to people who fly in airplanes, but haven't studied Bernoulli's hydrodynamics principles.

Research Partners (External Research Organizations and Working Groups)

The Institute's ReInference unit collaborates with external research partners, universities, institutions, and subject matter experts. These partnerships involve joint research projects, data sharing, and knowledge exchange to enhance the depth and breadth of research efforts. Collaborations with research partners create an opportunity to enrich The Institute's research capabilities and resource access, thereby accelerating the generation of new knowledge and helping us to address complex research questions, validate findings, and extend the reach of our research impact. Potential research partners include

organizations working on or faced with problems that may be solved by Active Inference, and organizations which are working on or have solutions to problems which The Institute and the community are facing.

Educational Partners (Universities and Educators)

The Institute's EduActive Unit collaborates with educational partners to influence, instantiate, share, and get access to educational programs, teacher training, and learning resources. By partnering with educational institutions, The Institute extends its educational reach and impact and fosters effective delivery and dissemination of its educational content. Potential educational partners include universities, tutors, educational institutions, and educators.

Funders (Donors, Supporters, Grants, and Funding Agencies)

The Institute requires Philanthropy in order to keep pace with community needs, maintain information infrastructure, and assist researchers in finding their own financial support for relevant research initiatives. Potential donors and funders include generous community members and beneficiaries, government funding agencies, private philanthropic donors, and sponsors of events, programs, and initiatives.

Ecosystem Projects

There are many Ecosystem Projects — here we include the subset which have completed a form at https://projects.activeinference.institute/ to increase their visibility and participation.

See Activities for all projects by Research Fellows, Scientific Advisory Board members, Current Partners, and Institute Projects. .

Active Inference Ecosystem Projects

Not synced yet

Project	Documentation	Mission & Objectives
Active Inference Account of Belief Updating in PTSD		Write a theoretical paper in the style of Parr et al. chapter 6
Improving RxInfer.jl's Model Visualization Capabilities		Our mission is to equip RxInfer.jl - and its relevant component libraries - with a host of model visualization modalities that prove useful to those who wish to use, and/or to develop RxInfe To that end, we anticipate measuring the initial quality of our contribution/s by their receptio RxInfer.jl's core developers: TU/e's BIASlab. All our objectives must therefore take the approv BIASlab as their proverbial North Star.
Neurodivergent Learning Sessions		Neurodivergent learning is focused on outreach and spreading awareness geared towards th struggle with standardized curriculum environments when it comes to public and higher edu milestones... as a number of people with neurological conditions not limited to autism spectr can struggle in varying ways with learning and being in the right environment in which inform presented to them in a manner which is coherent.
The Unordinary Bible Study (abbreviated as TUBS)		Hosting once a month sessions that focus on cross-referencing biblical verses but not spend much time digging into scripture as opposed to focusing on inter-faith and contemporary pe focused dialogue.
The Einstein Model of a Solid as a Model of the Mental Apparatus from the Economic Perspective of Psychoanalytic Theory.		Bridging Psychoanalysis and Thermodynamics with applications to Artificial Intelligence. App AI.
Project Sweet (Sus) Dogg		To Help Warm-up or Prepare a Plausibly Notable Aspect of Agent Based Alignment By Social
Active Inference for Built Environments & CooperActive Systems		To advance the application of Active Inference in designing, managing, and evolving built en that prioritize the flourishing of all life on this planet. While humans possess unique cognitive capabilities, we recognize that excessive anthropocentrism blinds us to the needs of other li organisms. Our work centers on **life prosperity** as the foundational principle for all built envir decisions. We seek to develop adaptive, nature-integrated solutions through distributed intelligence, di technologies, and decentralized decision-making that serve the broader web of life while me human cooperative living needs.

Domains of Application

The sub-sections of Domains of Application reflect some early collaborative efforts towards curating Implementations of Active Inference across different sectors and systems of interest. This section of the document is not presented as a comprehensive or exhaustive survey in any way, rather more of an invitation to those who would like to steward a section (keeping it updated and relevant) as we develop these synoptic capacities together. Later updates will more deeply reference Production and other resources where Active Inference has been demonstrated across systems.

Along with other modern technical fields, Active Inference faces and addresses challenges of broad relevance such as (i) remote education, workforce development, and competency evaluation, (ii) user experience, ergonomics, and accessibility in a modern global context, (iii) Open Source availability , utility, reliability, and safety, (iv) participation in research and practice-oriented activities (v) cyber- and cognitive-security, (vi) theoretical and practical aspects of artificial intelligence explainability and safety, (vii) social and economic policy integration and management.

Integrations featuring Active Inference are increasingly being found across public and private sectors. These applications are enabled through common education around Active Inference themes, concepts, skills, practices, and tools. As such, there is potential for The Institute to facilitate both the study (theory and research) and professionalization (practice and implementation) of Active Inference within and across myriad sectors and disciplines, and to grow the incipient Active Inference Ecosystem and awareness of Active Inference by facing such challenges proactively and in a fashion aligned with our vision, values, and principles. We hope to achieve this through developing coordinated resources that are accessible to users at all backgrounds and levels of familiarity. Moreover, we aim to develop this nascent research arena by facilitating and/or mediating access to resources for an array of independent projects.

A core reason why Active Inference is being adopted so rapidly is that it provides a flexible, agent- and action-oriented ontology which describes a great array of complex adaptive systems, up to and including human Social cognition.

The Active Inference framework can be used to describe systems at different nested scales. The applicability of Active Inference to multi-scale complex adaptive systems is a source of great explanatory power, and it is also a challenge for the framework's coherence. Scholars from different disciplines or fields may read Active Inference concepts or constructs differently, and unknowingly build an error into their research ecology which is then propagated forward, thereby hampering progress in the field at large. To our knowledge, the Active Inference Institute is the first scaled attempt at directly tackling that risk by offering Active Inference education to learners of all backgrounds, and by working to specify an ontology that is both particular to Active Inference and broadly accessible. Furthermore, the institute offers accessible onboarding to current best practices in Active Inference research as well as the ability to drill down into specific topics across the broad array of disciplines that are implementing the framework.

Biology

Chris Fields and Michael Levin (2020) posit that Active Inference "provide(s) conceptual tools for reconceptualizing biology as the study of a unified, multiscale dynamical system".

Ramstead et. al (2019) have leveraged active inference and the underlying free energy principle to characterize *variational neuroethology*, a theoretical ontology for living systems based on a recursively nested formulation of Markov blankets.

Friston et. al (2023) introduce a variational formulation of natural selection to explain how slow phylogenetic processes constrain—and are constrained by—fast, phenotypic processes.

Neuroscience

Active Inference emerged from the field of theoretical neurobiology (Friston, 2005), where it was "first used to model the function, structure, and dynamics of the human brain" (Ramstead, 2024). It built upon foundational work in predictive coding (Rao and Ballard, 1999) and the Helmholtzian concept of perception as "unconscious inference" (Helmholtz, 1867).

Active Inference's central premise that "all neuronal processing (and action selection) can be explained by maximizing Bayesian model evidence — or minimizing variational free energy" (Friston 2017) provides a unifying theory to explain and predict myriad aspects of brain function and behavior (Friston, 2010). As such, it has been applied to many areas of neuroscientific research.

Active Inference models are used to provide parsimonious explanations for neural mechanisms and motifs, such as canonical microcircuits and neural networks (Bastos et al 2012, Isomura et al 2022).

Researchers have furnished Active Inference models for phenomena including motivated control (Pezzulo et al 2018), sense of agency (Friston et al 2013), modulation of uncertainty by the dopaminergic system (Friston et al 2012), and the computational relationship between interoceptive and exteroceptive neural systems (Allen, 2022).

Active Inference frameworks have also been used to explain the dynamics of a variety of neurological and psychiatric conditions, including depression (Barrett et al 2016) and schizophrenia (Limongi et al 2023).

Recent studies have shown that in-vitro neuronal networks self-organize in response to stimuli in ways that are consistent with, and predicted by, the Free Energy Principle (Isomura et al 2023). The FEP also provides theoretical commitments towards testable theories of consciousness (Whyte et al, 2024).

As a multi-scale theory, Active Inference aims to ground neurobiology in physics-as-information-processing, and links it to other domains of inquiry, including diverse intelligence (Levin 2023) and artificial intelligence (Friston et al 2024).

Mental Health

Being a theory of embodied and sentient behavior, Active Inference can contribute in knowledge sharing to better understand the intrapersonal and interpersonal dynamics involved in or implicated in mental health (Pezzulo et. al, 2024). Computational psychiatry (Friston 2022) serves to utilize models of cognition and behavior to predict and account for the above-mentioned dynamics. Being a model constrained by Bayesian principles and the free energy principle, Active Inference allows for one such attempt at better predicting treatment outcomes, nosology, and fundamental principles of cognition.

The Institute supports individual thoughts and projects designed to inquire on topics related to social sciences, psychology, and mental health. Such projects have included attempts at classifying and clarifying Active Inference ontology to better fit the lived experience of individuals with posttraumatic stress disorder.

Active Inference is a systems approach to psychological constructivism that offers a trans-diagnostic perspective to readers. One such benefit of a trans-diagnostic approach is that it identifies connections between different processes without the strict adherence to philosophical requirements. Areas of the theory that can be beneficial to mental health research include:

- Experiential quality of prediction error for patients (i.e., as a mismatch of one's generative model) (defining "surprise" in therapy practice, Holmes & Nolte 2019)
- Homeostasis and role of consciousness as allostatic control (Krupnik 2024); as well as the dynamic interplay between these processes and mental health symptoms (cultural identity, Ramstead et. al 2016; social conformity, Constant et. al 2019)
- Mental health symptoms as under/over-reliances on a generative model (apathy, Hezemans et al. 2020)
- Requirement of interoceptive processes (body-based) and the roles these have in the make up of a Bayesian brain (Duquette 2016)
- Role of affect and ascribing confidence to one's generative model (Hesp et al. 2021)
- Equal treatment of action policies as being direct manipulations of one's environment (decrease free energy now) versus epistemic transformations (change your beliefs about the world to decrease free energy in the future) (PTSD & explore-exploit dynamics, Linson et al 2020; social cognition, Gallagher & Allen 2018)
- Hierarchical models of cognition that outline the dynamic interplay of predictions, actions, habituation, and environmental feedback (theory of constructed emotion, Barrett 2017; cognitivism versus autopoiesis, Allen & Friston 2016)

Conceptualizations are being offered that describe the experience of those with particular mental health symptoms (Parr et al. 2022, p. 186). Active Inference has also been applied to the study of depression (Barrett et al 2016), psychosis (Knolle 2023), schizophrenia (Jeganathan 2021), anorexia (Barca et al 2020), functional neurological disorder (Jungilligens et al 2022), and interoceptive dimensions of psychopathology (Paulus et al 2019; Barrett 2016). Conversely, Active Inference also provides a framework for understanding constructs of mental wellness, including subjective well-being (Smith et al 2022).

Chamberlin (2023) illustrates how Active Inference can be applied directly to one psychotherapy model, Coherence Therapy. This type of dialogue allows readers to see the neurological mechanisms and meaningful narratives at play in a framework that treats both equally. It is also good for readers to note that Active Inference offers a framework to reformulate agents as being cognitive, emotional, and embedded without adding other philosophical requirements. It can be beneficial to engage existing psychological theories [of cognition, emotion, personhood, agency, social relations] in order to emphasize constituent processes that Active Inference gives language for. For instance, the focus on sense making in the life of an individual highlights the existence of an agent's generative model that has been determined within and a part from the generative process. In parallel, sense making can speak to themes of agency and emotional expression.

Bioregional Modeling

An ongoing project, Biofirm currently consists of two main components in support of ⁂ Bioregional Modeling:

1. **Ecosystem Control System**
    - Active Inference-based multi-agent control framework
    - Homeostatic regulation of ecological parameters
    - Comparative analysis between random and controlled dynamics
2. **BioPerplexity Analysis**
    - California county-level bioregion research using Perplexity.ai API
    - Business case generation and pitch development
    - Cross-document visualization and analysis

Category Theory

Active Inference, through the Free Energy Principle (FEP), provides a framework for understanding how systems make predictions and update their models based on sensory evidence. Category Theory, meanwhile, offers a formal mathematical language to describe the transformational processes that occur during these updates (see: Chris Fields 2024, "What is the Identity operator?").

Mathematical Bridge

Where Active Inference describes the necessity of prediction and error minimization, Category Theory provides the precise mathematical tools to track how these predictions and updates flow through a system. The power of this combination lies in Category Theory's ability to formalize the very transformations that Active Inference predicts must occur.

Creative Processes

This relationship becomes particularly relevant when examining creative processes. Active Inference explains why systems must make predictions and learn from surprises, while Category Theory's operators can formally map out the transformational paths taken - even in cases where the end state wasn't predictable from the initial conditions. The identity operator, in particular, helps us understand how systems maintain coherence while undergoing these creative transformations.

Practical Implementation

This growing theoretical bridge between Category Theory and Active Inference has practical implications for:

- Modeling learning processes
- Understanding system adaptation
- Tracking creative development
- Formalizing prediction errors and updates
- Maintaining system identity through change

The synthesis of these approaches provides a more complete picture of how systems learn, adapt, and create while maintaining their essential identity through transformative processes.

Computational

Implementations of Active Inference

In the Active Blockference project, we have curated dozens of Open Source Implementations of Active Inference at this link.

The sub-pages here go into more detail on several different toolkits for applying Active Inference, including in Python (PyMDP), Julia (RxInfer.jl), Matlab (SPM (Statistical Parametric Mapping)), and Prolog (Symbolic Active Inference).

RxInfer.jl

RxInfer.jl (https://rxinfer.ml/) is a programming package of functions developed at BIASlab in Eindhoven, Netherlands. It attempts to commoditize ∞ Active Inference, making it suitable for engineering applications. Compared to existing ⠿ Implementations of Active Inference like 🐍 PyMDP, RxInfer is unique in the sense that it draws upon reactive message passing on Forney Factor Graphs (FFG). Whereas 'traditional' implementations rely on Bayes graphs in the form of Partially Observable Markov Decision Processes (POMDP). FFG's using reactive message passing only perform calculations when necessary, hence there is no underlying clock which schedules calculations. The reactive paradigm thus may offer computational benefits in certain situations, and enable favorable scaling properties for Active Inference models.

The 🎄 RxInfer.jl Learning Group at the Institute collaborates with with the developers of 🚀 RxInfer.jl in 🌍 Open Source development, such as developing visualisation techniques of the FFGs within the code editor.

Core Capabilities
RxInfer.jl provides powerful features for probabilistic modeling, including:

- Streaming dataset processing through reactive message passing
- Hybrid models combining discrete and continuous latent variables
- Scalable inference for large models with millions of parameters
- Automatic differentiation support for parameter tuning

PyMDP

The PyMDP package https://github.com/infer-actively/pymdp is an Open Source Implementations of Active Inference, specifically “A Python package for simulating Active Inference agents in Markov Decision Process environments.”

Active Inference ModelStream 007.1 ~ Conor Heins & Daphne Demekas ~ pymdp

https://www.youtube.com/watch?v=skf3sOM-7WI

Active Inference ModelStream 007.2 ~ pymdp

https://www.youtube.com/watch?v=uX8iSoDR83g

SPM (Statistical Parametric Mapping)

Statistical Parametric Mapping (SPM, homepage) represents a pivotal development in the history of active inference and computational neuroscience. Created by Karl Friston at the MRC Cyclotron Unit in the late 1980s, SPM began as a statistical technique for analyzing brain imaging data, particularly fMRI, PET, and EEG data (Wikipedia and History).

The development of SPM marked a crucial shift from simple region-of-interest analyses to whole-brain statistical approaches. Originally written in MATLAB, SPM91 (also known as SPMclassic) became the community standard for analyzing neuroimaging studies within a few years of its release. The software's success stemmed from its rigorous approach to making valid statistical inferences about brain responses without prior knowledge of where those responses would occur.

SPM's theoretical framework evolved to incorporate increasingly sophisticated statistical methods, including the general linear model (GLM) and Gaussian field theory. This evolution paralleled and supported the development of active inference theory, as many of the mathematical principles underlying SPM - particularly those involving free energy minimization and Bayesian inference - became foundational to active inference. Today, while dedicated Implementations of Active Inference toolboxes exist in various programming languages (like PyMDP in Python, RxInfer.jl in Julia), SPM remains significant as both a historical cornerstone and practical tool in the field.

Symbolic Active Inference

Symbolic Active Inference, developed by Research Fellows @Jean-Francois Cloutier , represents an innovative approach to combining symbolic reasoning with active inference principles. The framework aims to bridge the gap between traditional symbolic AI and the free energy principle by implementing active inference using symbolic representations and logical reasoning. This implementation allows for explicit modeling of beliefs, goals, and actions using symbolic structures while maintaining the core mathematical principles of active inference - namely the minimization of variational free energy and expected free energy.

The approach enables systems to perform goal-directed reasoning and planning through symbolic manipulation while grounding these processes in the formal theory of active inference. Key aspects include the representation of generative models using symbolic structures, belief updating through logical inference, and action selection based on expected free energy minimization. This synthesis provides several advantages: it makes active inference more interpretable through explicit symbolic representations, enables complex reasoning about abstract concepts and relations, and allows for more efficient computation compared to purely numerical implementations. The framework has been demonstrated through implementations in domains like robotic planning and symbolic problem-solving, showing how symbolic representations can be effectively integrated with active inference's information-theoretic principles. This work represents an important step in developing hybrid AI systems that combine the strengths of both symbolic and probabilistic approaches to intelligence.

All Open Source information on Symbolic Active Inference can be found at https://github.com/jfcloutier/karma_system

Economics

Economics is a very broad field. From macro policy to econometric micro optimization. Here the focus is on conceptualizing the decision maker as it is relevant for deciding a relevant policy alternative from a potential set. Undoubtedly future work and potential authors will expand this section greatly.

The foundation of economics is to scale decision making to collective systems. Traditionally, decision makers are seen as utility maximizers (or regret minimizers). With the underlying assumption of full information and (bounded) rationality.

However, active inference nuances this view by positing that rational choice is a limit case of decision making. Only occurring during absolute certainty of observing one's preferences (Friston et al., 2013). Instead a pragmatic turn entails information seeking as part of the decision process such that actions are both pragmatically and epistemically informed (Schwartenbeck et al., 2015).

Such a shift in perspective - all the way up to perspective swaps - may not be limited to traditional economics by expanding existing frameworks with new methodologies. Instead, this shift from viewing choice as static towards a dynamic process, means that multiple economic approaches to collective policy selection become feasible.

One such alternative economic approach is broad prosperity. It involves taking inventory of a set of value-neutral indicators, of which gross domestic product is just one. Unfortunately, it is very difficult to express the causal relationships between these indicators as these span a variety of domains like social, environmental and economic concerns. Additionally, what occurs locally has impacts globally and vice versa (TNO, 2024).

Active inference is poised to address these limitations. Given the nature of scale-free action perception loops; any self-organising system may be described as a sense-maker. In doing so solve the issue of not being able to sum free energy across agents. For example when planning a new public transport line. One could calculate the total utility obtained via preference elicitation (willingness to pay, stated and revealed choice experiments). Or one could instantiate a niche constructing digital twin. The entire urban region which is assumed to itself be a scale-free niche constructor will then have to share its niche with a synthetic artifact.

Evaluating the potential of a policy alternative, like building a tram or bus line, becomes a practice of understanding the generative model of the digital twin. Which is assumed to approximate a real niche constructor.

Education

The transdisciplinary nature and flexibility of Active Inference makes the framework ideal for practical, theoretical, and interoperable work across myriad use-cases. In the use case of learning in systematized settings (i.e. Education) the conventional planning frames take on *wheels* (π, as in policy selection) in order to function as a platform enabling translational "spinning" (i.e. helicity) across contexts of greater scale (learning generalization as **transfer**). With the inertia from the spin as your **stability** mechanism, the addition of policy selection by the learner as a self-organizing system (i.e. learning agent) within larger variability retained settings, introduces *uncertainties* to test the *which* and the *where* of when ***trans-disciplinary*** experience (i.e. real world experience, real dynamism, real problems) requires practical/pragmatic (i.e. action) solution(s). Comparatively speaking, conventional frame containment as stabilizer, only provides a variability reduced-reductionism environment ubiquitously held up as **constructing** learning where the product is a wide base as "foundational" retentions, and relatively smaller "crowning" states, as in Maslow's Hierarchy.

Before describing what the mechanics of this inclusion of policy selection is, and can look like for you, it is best to point out that going forward, the acceptance that policy selection plays a role in how we learn, is not necessarily easy to incorporate as strategy applied. "*I find this policy selection part hard to understand*" is often heard when something new and/or unfamiliar is introduced into a messaging exercise. This is understandable when a proposal uses terminology that isn't part of the newcomer's current lexicon (and sometimes even when the term is already used). To take up new labels (and the ideas behind them) requires taking a step **back** from centuries of the accepted definition of what providing an *education*...is: define and refine via a process of packaging and delivery of information (so deliver to me, the learner, what I can recognize). Sustainers and defenders of that (status quo) strategy will argue (correctly) there is much more going on than that minimum of two of ***define and refine***, and the Active Institute's argument would be...maybe, possibly, but not certainly.

There too many examples, practiced both currently and historically in academia, to deny that at the core of educational practice, there is a reinforcement and incentivisation firmly established around practices focused on defining (i.e. agreement around an external ontology/standards) and refining (i.e. moving to smaller and smaller divergence(s) from what we see/do, and what we think we're doing/seeing). That being the case, new terminology like **Prediction Matter Expert** is the ***surprise*** given that phrase's like this that are introduced, lack consensus around meaning and precision. Time is then spent working through where the introduced term/label/idea can fit (appropriately) within contexts of particular study/focus/research. This is an effortful exercise, that can often lead people new to Active Inference and the FEP, to wonder "*where exactly is the Institute going with this idea/terminology/set of formalisms?*" That's a fair question, and in asking, we open a portal to the navigational aspects of resolving the "where" of **learning as orientation process**. This is the "where am I?" **action** - not just wonder - as Active Inference.

Applying Active Inference and the FEP to educational programming - *"you are now here, but you're not staying here, you're going **back** out there"* - has thus far struggled to gain much traction in many legacy (read hierarchical Pyramid Model) educational systems. Given most education systems' tendencies to want to place the *certainty* of keeping systems accountable ahead of determining how agents learn when **prediction-as-skill** under ***uncertainty*** is given equal priority with subject matter expertise (as skill), we continue to find that active inference as functional **compliment** needs time for mass academic uptake (to scale). One of the core differences between subject matters and prediction matters exists at the waypoint called ***Updating***. Currently, legacy education systems interpret "updating" as a cumulative-constructive-classical exercise, and therefore it is ***surprising*** for those vested in that method, when someone with formal active inference priors, proclaims the need to incorporate statistical and probability functions into the praxis and pedagogy design. This non-binary nature of probability (i.e. could be zero, or one, or something **between**) aspect dependent on "what I as agent...*thinks* will happen," does exist as a teaching strategy, but is only applied within **the variability reduced frame**, pre-selected by the course/activity/lesson plan designer who is the subject matter expert.

And, active inference prediction modelling begins with the concept that the learning agent is first and foremost a self-organizer, self-designer who wants (self-identifies) minimization of any divergence between their own model and what the niche continually signals. Under this circumstance, updating as a process may take on constructive attributes, but it will also require some exposure to de-constructing processes (i.e. the most basic being, when change in the situation is apparent, will the agent 1) accept that change and 2a) either modify their surroundings or 2b) modify their model?). This is a fundamentally different type of branching - change the model, change the environment, change **both** - to pass/fail or even rubric induced accounts.

This then necessitates a different (second) definition of "updating" as a result of starting with a predictive probability of achieving an *ad hoc* **and** *post hoc* processing threshold (could be described as ALL moves cardinal vs. NEXT moves ordinal/sequential), *before* "right and wrong" or even "75% correct" as assessed (as the 25% "wrong" usually doesn't carry forward past the filters of **constructive** practices).

So why does this difference matter? In arriving at a threshold minimum, the active inference learning agent needs to **reconcile** while also keeping records. That "25% wrong" for example, is actually valuable information (not to be discarded) if **divergence minimization** is one of the stated goals. Now the question becomes "*do I let go of what I predicted wrong because it didn't affect my pass/fail status, do I let what I got wrong change my aspirations because I haven't achieved perfection, or, do I look at Right-Wrong as a* ***proportional measure*** *from which to make future decisions?*" (more on this shortly). Taking accounts and making reconciliations, is the process of modifications **by and to** which the **updating** of the active inference generative model, evolves. The conventional view of update as build-up, build-forth (Subject Matter Expertise, SME), is now complimented with a Prediction Matter Expert (PME) view of "*what can I as learning agent let go, in order to arrive at a new know?*" as policy selection **to be determined**. Borrowing from Chris Fields' Identity Operator presentation, **PME**'s cope better with the undecideability in the frame problem - *what* ***doesn't*** *change as a* ***result of an action***. Using Chris' terminology, "circumscribing what I don't have to worry about"...means "I" can now take my "eye(s)" off of certain contents so as to increase availability for new [to me as agent] contents. Under this condition, the forensics come before, and not just after, a learning episode, making policy selection (π) now one part agent domain, one part external plan designer/niche reducer domain - with All Moves now meaning all of the puzzle pieces are present, and each is connected **regardless** of order application.

Of course, once the differences between legacy systems perspectives and active inference perspectives are held up as the parameterized space, the ability for the learning agent to oscillate between perspectives (i.e. perspective swap as action) becomes available. This oscillating process - first **back**, then forth...and never forth-only - is not uncommon. Agents swap perspectives when pairing science with fiction, active with inference, math with art as *comparative* **with** *collective* **proportional measuring** (as minimum) processing (unit of) analysis.

Which leaves the Institute with a challenge: how do we continue to attract Subject Matter Experts **and** point to the fact that Subject Matter Expertise alone can only take one so far as a navigator in variability retained settings? Another way of putting this could be stated as, as an institute, can we afford to not talk about the gorilla in the room: how we learn (define and refine...and retain) needs a **co-pilot** (what can I let go...to arrive at a new know?). This being asked as AI and LLM's train on far more information than humans can, to derive that synthesis (here's your answer!) that defining and refining puts out (outputs).

Let's look at a real world example already introduced to the officers of the Institute where subject matter expertise attracted agents to the institute, and, the institute had to find a way to help the "experts" let go of what they already know. In this case example, Active Inference has been linked to the process of early childhood education (Montessori programming). Under Montessori philosophy, teacher's are described as "directors" with a focus on "independent learning." Comparisons can then be made to other early childhood education approaches. The Reggio Emilia early learning method holds up their philosophy of teachers roles as "partners" and "guides."

The question then becomes one of: as the learning agent ages - enters different "grades", stages and phases of *Updating* as a result of predictive processing (probability now based on **increased temporal depth**) - does the teacher **as** multi-hat wearing director/partner/guide/coach/facilitator still fit the needs of the self-organizing learner going forward? Perhaps, if the learning is organized as an *adventure* as a proxy for ***authentic*** - where once again authentic is trans-disciplinary *real world experience, real dynamism, real problems) requiring practical/pragmatic (i.e. action) solution(s),* while **an adventure is a simulation**.

Or, as a PME enabler (**Not** trainer), does the teacher SWAP titles - by **subjecting themselves** to the Identity Operator process - of Teacher **with** Way Finder (navigator), initiating their own perspective exchanging process of self-identifying (minimizer of divergence between their own model, now as minimum(2) dual-state swap able [i.e. Gripper & Gripped - **BY and TO** - simultaneously], with what the niche continually signals) resulting in an SME + PME hybrid triangulating with ANY niche (not just *their subject* specialty)? This would require teachers to both teach and co-learn interchangeably.

As the reader can appreciate, this is a different condition than teachers staying close (closed) to what they know (SME dilemma) and thus self-selecting away from "*what can I let go, in order to arrive at a new know?*" This is where the Institute's role as director/partner/guide/coach/facilitator ends, and a co-piloting triangulation exercise (i.e. simulations to actualizations and **Back**) begins.

Going forward, it is the Institute's ambition to make clear that the channel (i.e. gap) between legacy systems developing subject matter experts **and** what we view as new affordances that can be realized when *uncertainty*-as-learning-tool is perceived as a feature - as prediction matter expertise - is a potential *exponentiator* of a learner's predictive capacities within **and beyond** systematized and variability reduced settings. We choose to be partners in this enterprise, as we feel serving in that capacity is closer to co-piloting than co-hosting in a flight simulator. Every organization wonders where the "stay afloat" energy will come from. In our case, we policy select to work with people vested in research with a specialty focus who also want to be able to *generalize* (play in "Scale Free") with higher degrees of confidence when necessary (be a trans-disciplinarian when the niche is open, and variability retained).

Philosophy

Active inference provides a mathematical model of sense-making. Philosophy is the study the components and dynamics captured by these mathematical models. It comprises of a broad literature in philosophy of mind spanning history. As such active inference is not an island but densely connected to other descriptions and mathematical models of sense-making. Each with their own components and dynamics. One such advancement is within the field of neurophenomenology that seeks to build a dialogue between neuronal processes and philosophical constructs as experienced (Sandved-Smith et al. 2021).

Developing new implementations and applications of active inference benefits greatly from its philosophical context. Theoretical advancement of how to interpret existing phenomena through the lens of active inference is not just to fill shelves with studies. Advancement is essential to improve algorithms and inform applications across domains. After all, there are many domains which have mathematical and conceptual models of sense-making. Each of which could potentially be evaluated through an active inference lens.

Physics

See Physics course by Chris Fields, and more references to come.

Robotics

- See Fellows JF Cloutier's project, Symbolic Active Inference (project documentation)
- Second Applied Active Inference Symposium in 2022 had a focus on Robotics, see program.

Legal

Cases mentioning active inference

Patents mentioning active inference

Social

Active inference research in the social domain tends to focus on modeling communication and the sharing of belief models within groups. Such topics can be understood as pertaining to normative processes of group cognition. Over time, we can expect research to extend further to pathological examples of group cognition, assessment of group cognition quality, steering of group cognition to improve quality, evaluation of the cognitive architectures used during group cognition (e.g., rules, policies, computational tools, communication tools, attention mechanisms), and evaluation of group cognition where the group is a political body (such as a city or nation).

Group cognition rests on the communication of (potentially dynamic and evolving) belief models—-the internal generative models that individuals use to predict and explain their world—and consensus building with respect to beliefs. As described by Albarracin et al., 2024 for a generic group in the normative setting, "group members can be seen as actively and implicitly aligning their beliefs and expectations through dialogue and interactions, thereby enhancing their ability to predict each other's actions and intentions, and thereby coming to perceive and act in the world in similar ways."

Most humans do not conceive of their own beliefs in terms of models, however. Rather, humans tend to experience their beliefs and make sense of the world in part through narratives (Bietti, Tilston, and Bangerter 2019; Turner et al. 2023; Fanti Rovetta 2023; Cordes et al. 2021). These can be internal narratives that a person constructs, adjusts, and recites to himself or herself, or social narratives that are shared within a group. In the active inference context, Albarracin et al. (2021) consider social scripts, which are widely-supported prescriptions about how one is to behave in various social settings, or what is important in those settings. Bouizegarene et al. (2020) consider shared narratives conceived of more broadly. Social scripts and shared narratives help humans to generate more accurate predictions about the world and to coordinate social behavior.

The CogNarr Ecosystem, an Active Inference Institute project, has as a goal the facilitation of group cognition at scale, through sharing of belief models (Boik, 2024a, Boik, 2024b). This is an extension of previous work that viewed core societal systems (e.g., economic, financial, and governance systems) as part of the cognitive architecture of political bodies (Boik, 2020a, Boik, 2020b, Boik, 2021)

A large body of active inference research, perhaps thousands of papers, at least mentions the social setting. In addition to some articles already cited, articles in which the phrases "active inference" and "social" appear in the title include the following:

- Gallagher and Allen, 2018. Active inference, enactivism and the hermeneutics of social cognition.
- Hipólito and van Es, 2022. Enactive-Dynamic Social Cognition and Active Inference.
- Constant et al., 2019. Regimes of Expectations: An Active Inference Model of Social Conformity and Human Decision Making.
- Cheadle et al., 2024. Active Inference and Social Actors: Towards a Neuro-Bio-Social Theory of Brains and Bodies in Their Worlds.
- Kirchner et al., 2022. Better Safe than Sorry?-An Active Inference Approach to Biased Social Inference in Depression.
- Tehrani-Safa et al., 2024. Learning Risk Preferences Through Social Interaction: An Active Inference Approach
- Fox, 2021. Active inference: Applicability to different types of social organization explained through reference to industrial engineering and quality management.
- Bezzazi, 2021. Social Active Inference.
- Ohata and Tani, 2020. Investigation of the Sense of Agency in Social Cognition, Based on Frameworks of Predictive Coding and Active Inference: A Simulation Study on Multimodal Imitative Interaction.
- Matsumura et al, 2023. Social Emotional Valence for Regulating Empathy in Active Inference.
- Solymosi and Schulkin, 2024. Creative Resilience. Flourishing and Valuation through Social Allostasis and Active Inference.
- Tani, 2019. Accounting Social Cognitive Mechanisms by the Framework of Predictive Coding and Active Inference: A Synthetic Experimental Study using Robotics Interaction Platforms.

Some of these topics and papers were explored in the 2023 Social Science course.

Logistics

“Enhancing Population-based Search with Active Inference” (Dehouche and Friedman, 2024)

Scientific Method

See ≈ Active Entity Ontology for Science (AEOS)

Distributed Science - The Scientific Process as Multi-Scale Active Inference (Balzan et al. 2023)

Generative Research Teams: Active Inference Compositions For Research and Meta-Science (Friedman & Smekal 2023)

DeSci

Decentralized Science (DeSci) was explored in the ≈ Active Entity Ontology for Science (AEOS) work.

Discussion and Future Directions

The Active Inference Institute attracts and amplifies the self-organizing abilities of people, thereby potentiating a unique opportunity and a powerful and scalable platform from which to accomplish research and development goals. As members of the Ecosystem, we continue to evolve an understanding and "voice" clarifying who we are, and who we might become, as a collective. In the process of building both an organizational reputation and individual expectations, we are constantly reminded of and inspired by the fact that the object of our research and development work, Active Inference, itself anticipates analysis and integration well beyond systems that are "closed" in time or space (i.e. those constrained to evolve linearly with a beginning, middle, and end as structure). We are interested in modeling, designing, and working with "open" systems, and have sought to cultivate an Ecosystem and larger community that reflects the intrinsic openness and systemic "curiosity" of Active Inference. With additional resources to support the work described in this application, the benefits of these open systems and guidance on future interaction practices consistent with Active Inference can be readily made available across myriad domains.

The Institute's work and community building efforts have always exemplified the benefits of "open" systems, consistent with the insights gleaned from Active Inference research itself. For example, when tracking open system behavior associated with the development and evolution of Active Inference, The Institute might have chosen to place an emphasis on "closing" (i.e. erecting constraints, applying limits, setting conditions, etc.) to simplify the challenge of modeling the space. However, rather than take a closed system (laboratory-centric) approach alone, which might have relegated Active Inference to an isolated academic disciplinary silo, we recognized the benefits that accrue from an "open" approach that invites self-organizers in the broader Active Inference Ecosystem to migrate (the "active" in Active Inference) into programs and participation that best suits their needs and prior experiences. Members of the Ecosystem will continue to encourage and support the opportunities to embrace variation-retained field studies for Active Inference...everywhere.

In the Ecosystem, we recognize that people and entities are explorers, capable of self-organization, motivated and eager to discover, and change agents of Active Inference approaches in the truest sense. By applying and leveraging the collective expertise of our community in preparation, scouting, and wayfinding practices, The Institute aims to continue helping Ecosystem participants to move ideas off-the-bench and into complex real world situations where the interaction environment acts as the ultimate scrutinizer. The resilience, sustainability and responsiveness of biological systems described through Active Inference research suggests that the human and social systems benefits of applied Active Inference framings will enhance the positive impact on the organization and operation of humans, including but not limited to The Institute itself.

To the people already involved in the Active Inference community, the "Ecosystem" isn't just a hypothetical and aspirational future state. Instead, it is the actual current world of interactions among members of the Active Inference community that we inhabit at all times. Active Inference, The Institute and members of the Ecosystem are all focused on dynamically adapting the efficiencies of change management practices as we prepare, scout, and "way find" our way into the future with measurable degrees of understanding around confidences, probabilities and the underlying mechanics involved, rather than depending on static plans that are quickly rendered obsolete in times of rapid change. As The Institute and Ecosystem help build competence and confidence in more agents in forms of organization and operation that reflect and apply Active Inference concepts, we will grow the pool of potential first finds (discoveries and inventions) and high-reliability knowledge systems in our world. Cultivating those skills as part of who we are as individuals and in organizations, and sharing those skills with others who are eager to see the future, and to be the future, is more than just an attractor state to guide our actions. It is the core mission of The Institute, Ecosystem, and its participants.

Act. Infer. Serve.


serve
act
infer

Projects

Projects are the main units of work and participation at the Institute.

See active projects at Activities.

- Core Institute projects include Active Inference Ontology, Active Inference Journal, Production, Textbook Group, RxInfer.jl Learning Group, Active Blockference, etc.
- Additionally there are projects in The Active Inference Ecosystem that are engaging with, or hosted at, the Institute.

Drawing on work in Quantum Active Inference (see Physics course), we structure project proposal and reporting at the Institute, according to Project ~ Preparation and Project ~ Measurement stages.

- Project ~ Preparation is for proposing or informing the Institute about projects that you want to facilitate.
  - Active projects that you can join are in Activities.
- Project ~ Measurement is for reports on completed work, whether the project had been proposed initially to the Institute or not.
  - Measurements can optionally be included in the monthly Newsletter, so please send in

To propose a project: Project ~ Preparation form

To report on a project: Project ~ Measurement form

Project ~ Preparation

Thank you for your interest in
→ Project ~ Preparation at or with the Active Inference Institute!

The → Project ~ Preparation form below is for projects that you are proposing, so that others can see information on the project and get involved.

Active and past projects are listed in Activities.

Below, unfold PREPARE questions to see the questions that are on the form.

▼ PREPARE questions

Not synced yet

Type		Question
▼ Team	4	Who will be the primary point of contact (i.e. project facilitator) for this project?
		What is the preferred mode of contact and contact information for the project? (Discord, email, X, other)
		Who else is already participating in what roles/positions?
		If the project is open for collaboration, who else might you like to have participate in what roles/positions? What skills are you looking for?
▼ Synchronous Meetings	2	What kind of synchronous meetings will the project have?
		Will you work within your project to manage the calendar, or do you want to have the Institute host the calendar events for this project?
▼ Project	7	What is the project title?
		What is the situation or problem that this project and team seeks to address?
		Given the situation, what are the team's mission and objectives?
		How does this project relate to and/or apply Active Inference?
		What are the milestones and/or timelines for this project?
		Under what circumstances will the project be closed and the team dissolved?
		Any other information about this project or team?
▼ Impact	3	How might/will this project contribute to the **ACCESSIBILITY** of Active Inference?
		How might/will this project contribute to the **RIGOR** of Active Inference?

		How might/will this project contribute to the **APPLICABILITY** of Active Inference?
▼ Institute	3	How might this project contribute to the work and mission of Active Inference Institute?
		What would be the MINIMAL type and level of support that the Active Inference Institute could provide for your project to still proceed?
		What would be a HIGH level of support that the Active Inference Institute could provide for your project to proceed?
▼ Ecosystem	2	How might this project contribute to the Active Inference Ecosystem?
		How might this project get support from the Active Inference Ecosystem?
▼ Other	1	Is there anything else you would like to communicate to us or any questions you may

Email blanket@activeinference.institute with any comments or questions.

PREPARE form link

What is the project title?

Who will be the primary point of contact (i.e. project facilitator) for this project? *

What is the preferred mode of contact and contact information for the project? (Discord, email, X, other) *

Who else is already participating in what roles/positions?

If the project is open for collaboration, who else might you like to have participate in what roles/positions? What skills are you looking for? How can people join?

What kind of synchronous meetings will the project have?

Will you work within your project to manage the calendar, or do you want to have the Institute host the calendar events for this project?

What is the situation or problem that this project and team seeks to address?

Given the situation, what are the team's mission and objectives?

How does this project relate to and/or apply Active Inference?

What are the milestones and/or timelines for this project?

Under what circumstances will the project be closed and the team dissolved?

Any other information about this project or team?

How might/will this project contribute to the ACCESSIBILITY of Active Inference?

How might/will this project contribute to the RIGOR of Active Inference?

How might/will this project contribute to the APPLICABILITY of Active Inference?

How might this project

What would be the

What would be a HIGH

Project ~ Measurement

The Measurement form is the main way that participants in The Active Inference Ecosystem share updates from their work, for getting visibility, feedback, and collaboration. With your permission, your Measurement may be included in the monthly Newsletter.

Active and past projects are listed in Activities

The → Project ~ Preparation form is for new projects that you are proposing, or raising for visibility.

The Project ~ Measurement is for updates from your projects, which may be shared in Newsletter.

Thank you for your diligence in reporting a Project ~ Measurement to the Active Inference Institute!

Below, unfold to MEASURE questions see the questions that are on the full form.

▼ MEASURE questions

Not synced yet

Type		Question
▼ Team	2	Who participated in the project and how?
		Who should be contacted regarding the project's work ar
▼ Project	4	Which of the project milestones were achieved or on whi progress?
		What happened? What did you measure? How and Why?
		What, if anything, is impeding the group's progress on th
		What was learned? What generative model inferences ar
▼ Impact	3	How did this project contribute to the **ACCESSIBILITY** of
		How did this project contribute to the **RIGOR** of Active In
		How did this project contribute to the **APPLICABILITY** of
▼ Institute	2	How do you think this project contributed to the Active Ir
		How was this project supported by the Active Inference I
▼ Ecosyste m	2	How did this project contribute to the Active Inference Ec
		How was this project supported by the Active Inference I
▼ Collabora tion	2	How often and how did the group meet synchronously?
		How often and how did the group collaborate asynchron
▼ Other	1	Any other information to provide?

Email blanket@activeinference.institute with any comments or questions.

Unfold here for a short Measurement form

If there is an extant project associated with your report, select it from the dropdown list:

If your Measurement is not from an extant project, please provide the name of the project here:

Who participated in the project and how?

Who should be contacted regarding the project's work and how?

What happened? What did you measure? How and Why?

Do you have any other information to provide?

Full MEASURE form link

Or complete the full measurement form below:

If there is an extant project associated with your report, select it from the dropdown list:

If your Measurement is not from an extant project, please provide the name of the project here:

Who participated in the project and how?

Who should be contacted regarding the project's work and how?

What happened? What did you measure? How and Why?

Please describe any deliverables or consequences of the team's efforts. Provide titles & links to publications or artifacts if possible.

How often and how did the group meet synchronously?

How often and how did the group collaborate asynchronously?

Which of the project milestones were achieved or on which milestones did the team make significant progress?

What, if anything, is impeding the group's progress on the project?

What was learned? What generative model inferences and updates happened?

How did this project contribute to the ACCESSIBILITY of Active Inference?

How did this project contribute to the RIGOR of Active Inference?

How did this project contribute to the APPLICABILITY of Active Inference?

Get Involved

Learn & Apply Active Inference. Contribute and participate in the Ecosystem.

For individuals

- There are many ways to Get Involved with both Institute Projects and Ecosystem Projects.
  - Welcome! Wherever you are in your Active Inference journey, we seek to help you find niches for your that scaffolding and development. Individuals of all backgrounds and level of familiarity with Active Inference are welcome at the Institute.
  - Check out the Volunteer, Internship, and Fellows programs.
  - See Affordances for specific contribution opportunities.
  - Join our Discord for chat, our Newsletter for emailed updates, and email us if you have other questions.
- Donate to our Philanthropy efforts (the Institute is a 501(c)(3) organization).

For organizations

- Partnership for organizations looking to engage with & support The Active Inference Ecosystem

---

Email blanket@ActiveInference.Institute with any ideas or questions.