Active Inference: A Neuro-Inspired Blueprint for General Intelligence

From the Free Energy Principle to Scalable Intelligent Agents

Dr Alexander D. Shaw. University of Exeter

SiC25. 17-18 Sep 2025

cpnslab.com

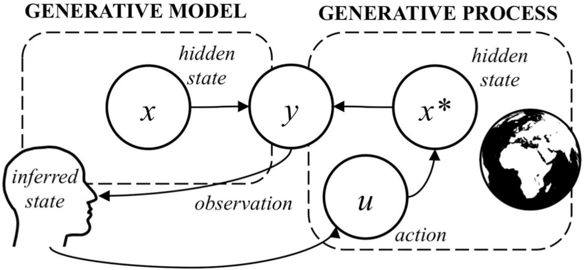

Active Inference: *A Neuro-Inspired Blueprint for AGI*

Why this matters for AGI

Any system that *persists* must avoid disorganising ("surprising") states. Biological agents do this by building models of their world, *inferring* hidden causes of sensory data, and *acting* to make future data match their predictions.

Active Inference in one sentence

Agents **update beliefs to explain data (perception)** and **select policies to make data predictable (action)**, jointly *minimising expected free energy*.

Core decision objective: Expected Free Energy (EFE)

$$G(\pi) = \underbrace{\mathbb{E}_{q(o,s|\pi)}[-\ln p(o)]}_{\textbf{pragmatic: reach preferred outcomes}} - \underbrace{\mathbb{E}_{q(o,s|\pi)}\Big[D_{\mathrm{KL}}(q(s\,|\,o,\pi)\,\|\,q(s\,|\,\pi))\Big]}_{\textbf{epistemic: reduce uncertainty (learn)}}.$$

Why this matters for Superintelligence

- **Unified objective:** One principle drives perception, learning, exploration, and control.
- **Built-in curiosity:** Exploration emerges intrinsically.
- **Model-based & embodied:** Learns a world model and uses it for inference, planning, and action.
- **Scalable architecture:** Generalises to hierarchical, modular, self-modelled agents.

key:
$G(\pi)$: expected free energy
$\mathbb{E}[\cdot]$: expected value (average)
$q(o, s \mid \pi)$: predicted outcomes and states
$p(o)$: preferred outcomes
$q(s \mid o, \pi)$ vs. $q(s \mid \pi)$: posterior vs prior beliefs

Neuro-Inspired Active Inference Agent (Pong)

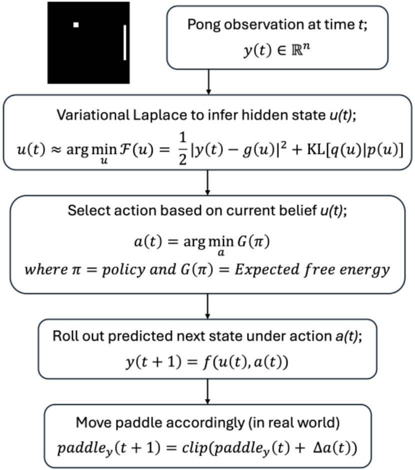

Schematic: per-timestep inference–action loop.

Overview

Predictive-coding / variational message-passing agent that *updates beliefs* about hidden states and *selects actions* to minimise free energy.

Generative model

$$x_{t+1} = f(x_t, a_t) + \omega_t, \qquad y_t = g(x_t) + \epsilon_t$$

x_t: hidden state, a_t: action, y_t: observation, ω_t, ϵ_t: noise.

Free-energy objective (perception)

$$\mathcal{F}(q) = \mathbb{E}_{q(x)}\big[\ln q(x) - \ln p(x, y)\big]$$

Action selection via EFE (planning)

$$G(a) = \mathbb{E}_q\big[-\ln p(o)\big] - \mathbb{E}_q\Big[D_{\mathrm{KL}}\big(q(s \mid o, a) \,\|\, q(s \mid a)\big)\Big]$$

Inference loop (each timestep)

1. Observe y_t (noisy, partial)
2. Infer x_t by minimising $\mathcal{F}$ (Laplace)
3. For each a: rollout, compute $G(a)$
4. Execute $a_t = \arg\min_a G(a)$, get y_{t+1}

Pong specifics: observations = noisy ball/paddle coords; action = paddle Δy; control aims to bring ball–paddle contact into high-probability states.

Why FEP/Active Inference? *A biological must-do*

Core idea

Any system that *persists* - a cell, animal, robot... must keep its sensory states within viable bounds. If its encounters are too chaotic or surprising, it disintegrates. So it must *predict and control* its future sensations well enough to stay alive.

What follows from this

- The system needs an **internal model** of how its sensations arise from hidden causes in the world.
- It must **update beliefs** to better explain incoming data (perception/learning).
- It must **act** to make incoming data more predictable and in line with what keeps it viable (control/policy selection).

Why it matters for AI

This gives a **single normative objective** for perception, learning and action rather than stitching separate objectives together. It also builds in **curiosity** (information-seeking) as part of staying within viable bounds.

Free Energy Principle (FEP): *inference as a survival tool*

The FEP

Under the FEP, a persisting agent behaves as if it *minimises variational free energy*: a quantity that upper-bounds how *surprising* (unlikely) its sensations are under its world model.

Equation

$$F(q) \;=\; \mathbb{E}_{q(s)}\big[\ln q(s) \;-\; \ln p(s,o)\big]$$

Components:

- o: sensory observations (*what the agent senses*)
- s: hidden states that cause sensations (*what's out there*)
- $p(s,o)$: agent's *generative model* linking hidden states to observations
- $q(s)$: agent's *approximate posterior* beliefs about hidden states
- $\mathbb{E}_{q(s)}[\cdot]$: expected value (average) over states s according to current beliefs $q(s)$
- $F(q)$: *variational free energy*; smaller F means beliefs better explain data

Immediate consequence

Minimising F is (approximately) equivalent to Bayesian inference: improve beliefs so sensations are less surprising under the model.

FEP: distance-to-posterior + evidence

Equivalent decomposition

$$F(q) \;=\; D_{\mathrm{KL}}\big(q(s) \,\|\, p(s\,|\,o)\big) \;-\; \ln p(o)$$

Interpretation:

$D_{\mathrm{KL}}\big(q(s) \,\|\, p(s\,|\,o)\big)$: how far the agent's beliefs $q(s)$ are from the exact Bayesian posterior.

$-\ln p(o)$: (negative) model evidence for the observed sensations under the generative model.

Mean

Minimise the gap to the true posterior while making your sensations more expected under your model. That's how a biological system avoids disorganising ("surprising") states.

Active Inference: one objective for perception & action

Definition

Active Inference extends FEP to include *action*: the agent both

- **updates beliefs** (perception/learning) to reduce free energy, and
- **selects policies** (action sequences) that are expected to reduce future free energy.

Embodiment

Intelligence is *situated*: the agent's body, sensors, and effectors define the interface to the world (often framed via a *Markov blanket*). Morphology and environment constrain what the agent can sense, predict, and do; that is, *which* policies are even feasible.

Why this helps

Unlike pipelines that glue a perception module to a control module, Active Inference uses one principle for both. That makes **exploration** (information seeking) and **goal pursuit** two sides of the same coin.

Expected Free Energy (EFE): decisions that learn & achieve goals

Operational objective for choosing policies (fixed size, wrapped)

$$G(\pi) \;=\; \underbrace{\mathbb{E}_{q(o,s\,|\,\pi)}\big[-\ln p(o)\big]}_{\textbf{pragmatic value: prefer outcomes you want}} \;-\; \underbrace{\mathbb{E}_{q(o,s\,|\,\pi)}\Big[D_{\mathrm{KL}}\big(q(s\,|\,o,\pi)\,\|\,q(s\,|\,\pi)\big)\Big]}_{\textbf{epistemic value: expected information gain}}\,.$$

Components:

π: a *policy* (candidate sequence of actions)

$p(o)$: *preferences* over outcomes ($\uparrow$ when outcomes are good/viable)

$q(o, s\,|\,\pi)$: predicted outcomes and states if the agent follows π

$q(s\,|\,o, \pi)$ vs. $q(s\,|\,\pi)$: posterior vs. prior beliefs about states under π

$G(\pi)$: *expected free energy*; choose policies that $\min G(\pi)$

Meaning

Pragmatic term: "Head toward outcomes I prefer/that keep me viable."
Epistemic term: "Also choose actions that teach me most about hidden states (reduce uncertainty)."

One-minute recap

Biological must-do: Persisting systems must avoid "surprising" (disorganising) sensory states.

FEP: Do approximate Bayesian inference by minimising variational free energy.

Active Inference: Use the same principle to *act*, not just perceive.

EFE: Choose policies that both reach preferred outcomes *and* reduce uncertainty: *goal-seeking that learns*.

Cover image adapted from: Pezzulo, G., Rigoli, F., Friston, K. J. (2024), *Biological Psychology*, 175, 108741. DOI:10.1016/j.biopsycho.2023.108741