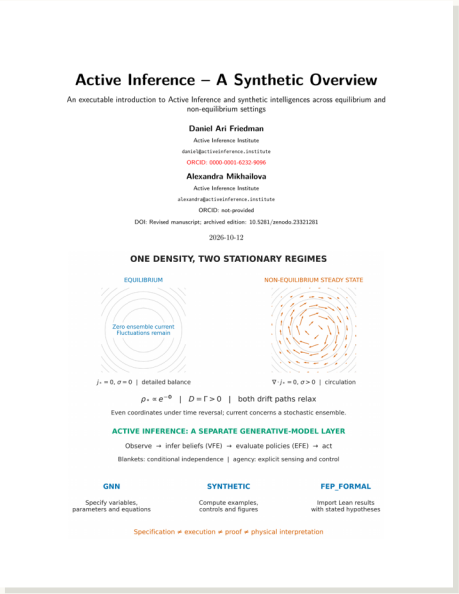


Active Inference – A Synthetic Overview

An executable introduction to Active Inference and synthetic intelligences across equilibrium and non-equilibrium settings

Daniel Ari Friedman
Active Inference Institute
daniel@activeinference.institute
ORCID: 0000-0001-6232-9096

Alexandra Mikhailova
Active Inference Institute
alexandra@activeinference.institute
ORCID: not-provided

DOI: Revised manuscript; archived edition: 10.5281/zenodo.23321281

2026-10-12

ONE DENSITY, TWO STATIONARY REGIMES

EQUILIBRIUM

Zero ensemble current
Fluctuations remain

NON-EQUILIBRIUM STEADY STATE

$j_\pi = 0, \sigma = 0$ | detailed balance

$\nabla \cdot j_\pi = 0, \sigma > 0$ | circulation

$\rho_\infty \propto e^{-\Phi}$ | $D = \Gamma > 0$ | both drift paths relax

Even coordinates under time reversal; current concerns a stochastic ensemble.

ACTIVE INFERENCE: A SEPARATE GENERATIVE-MODEL LAYER

Observe → infer beliefs (VFE) → evaluate policies (EFE) → act

Blankets: conditional independence | agency: explicit sensing and control

GNN
Specify variables, parameters and equations

SYNTHETIC
Compute examples, controls and figures

FEP_FORMAL
Import Lean results with stated hypotheses

Specification ≠ execution ≠ proof ≠ physical interpretation

SYNTHETIC / REVISED PAPER AND LECTURE

Active Inference -- A Synthetic Overview

Active Inference Synthetic

A worked introduction to beliefs, decisions and dynamics

Revised paper, lecture and sources

GitHub release: synthetic-teaching-v0.6.0

FEP_Formal published reference

zenodo.org/records/23220027

PDF cover reproduced without changing its content

Scan in the video or click in the slide PDF

Three questions, three mathematical objects

INFERENCE	DECISION	DYNAMICS
What explains an observation?	Which action comes next?	What is stationary?
Posterior belief	Expected free energy	Density and current

READ THE RESULT

- Inference: posterior belief
- Decision: expected free energy
- Dynamics: density and current

KEEP THE OBJECTS DISTINCT

Three questions; three objects

A generative model makes assumptions explicit

PRIOR	LIKELIHOOD	POSTERIOR
Before observing	How observations arise	After observing
p(x)	p(o | x)	p(x | o)

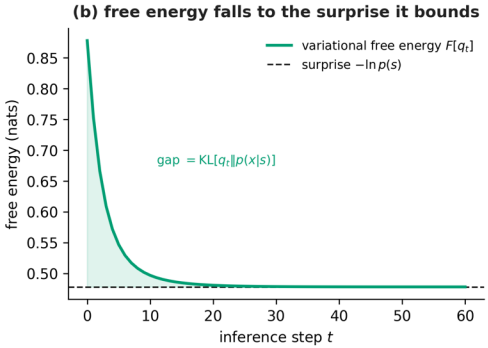
READ THE RESULT

- Declare hidden states
- Specify how observations arise
- Condition on actual evidence

BUILD THE MODEL

Prior × likelihood → posterior

Variational free energy bounds surprise



READ THE RESULT

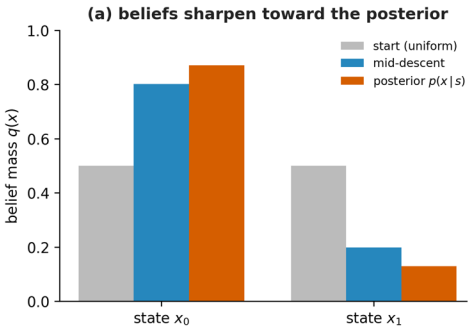
- The curve approaches its bound
- The gap is belief error
- The identity needs its support domain

READ THE GAP

F[q] = surprise + posterior KL

Source: Synthetic / 02_generative_model.md

Read the belief update as a numerical experiment



READ THE RESULT

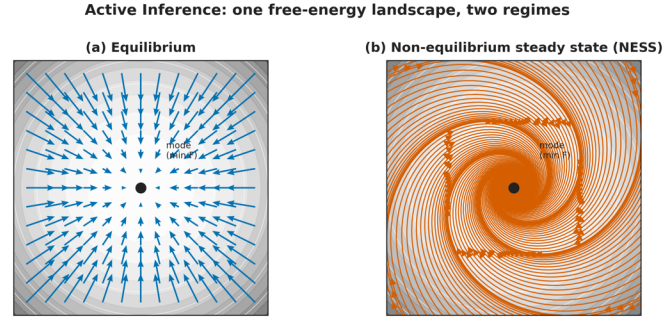
- Uniform starting belief
- Posterior calculated independently
- Descent needs convergence checks

CHECK THE ANSWER

Compare belief with exact Bayes

Source: Synthetic / 02_generative_model.md

One landscape, two components of motion



READ THE RESULT

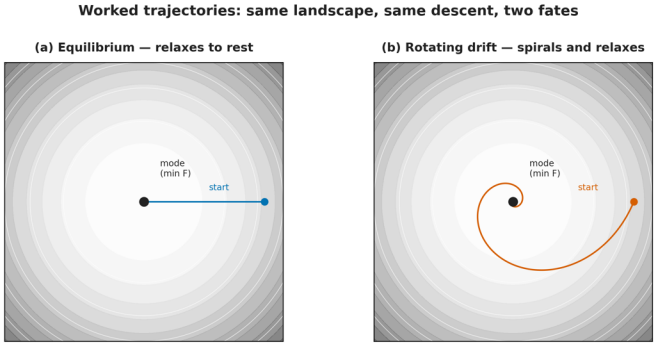
- Dissipation lowers the landscape
- Antisymmetry gives tangent motion
- Both terms belong to the drift

READ THE FLOW

Downhill + along-contour motion

Source: Synthetic / 03_worked_regimes.md

A spiral is a drift trajectory, not a steady state



READ THE RESULT

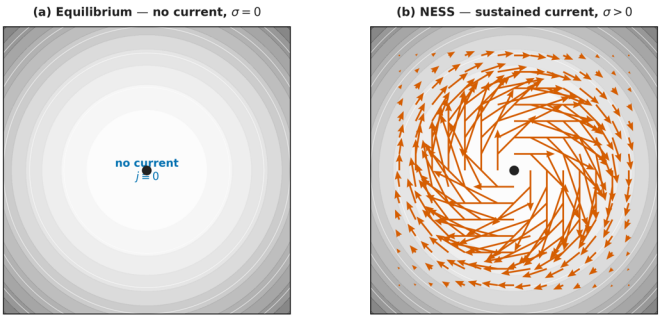
- Paths approach the mode
- Rotation depends on the time window
- Steady state describes an ensemble

FOLLOW THE TRAJECTORY

A spiral still relaxes

Source: Synthetic / 03_worked_regimes.md

Watch a transient density relax



READ THE RESULT

Mean starts displaced

Covariance stays fixed

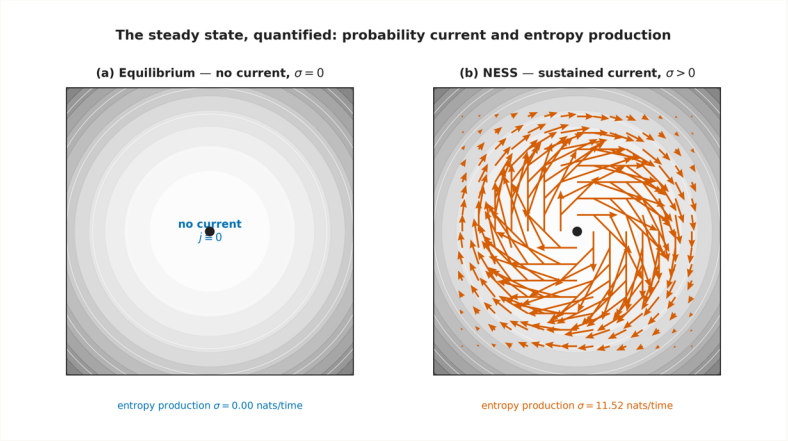
Contours show an ensemble

WATCH THE TRANSIENT

Changing density is not stationarity

Source: Synthetic / 03_equilibrium_concepts.md

The same stationary density, different currents



READ THE RESULT

Equilibrium: current zero everywhere

NESS: divergence-free circulation

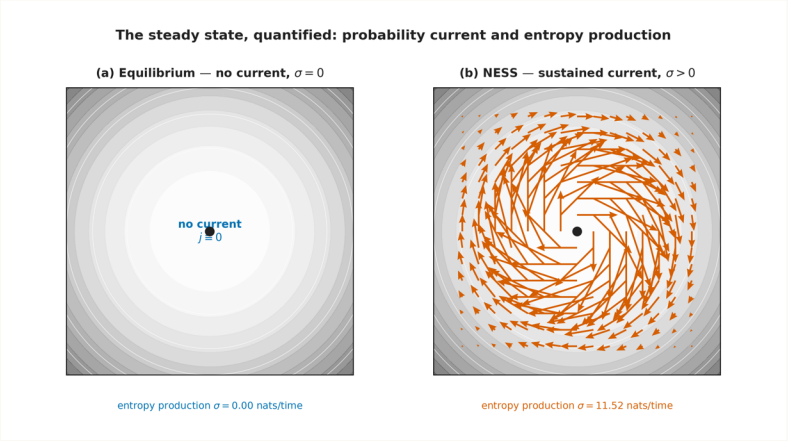
Density stays the same across this sweep

SWEEP THE DRIVE

A parameter sweep is not a time trajectory

Source: Synthetic / 03_equilibrium_concepts.md

Entropy production: what the formula measures



READ THE RESULT

Circulation reversal preserves cost

Physical units need conventions

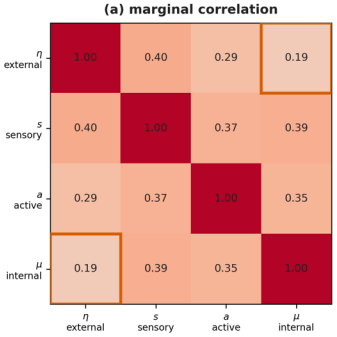
Positive rate does not prove cognition

KEEP THE UNITS

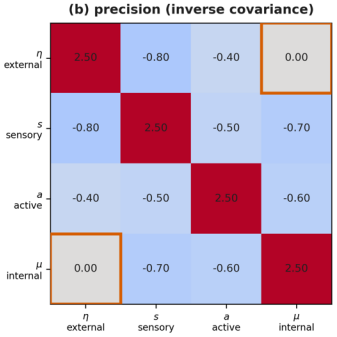
Irreversibility in a specified model

Source: Synthetic / 11_thermodynamic_interpretation.md

The Markov blanket is a conditional statement



η and μ are correlated: corr $= +0.19$



but $\Pi_{\eta\mu} = 0 \;\Rightarrow\; \eta \perp \mu \mid \{s, a\}$
(partial corr given blanket $= +0.00$)

READ THE RESULT

External and internal can correlate

Full conditioning removes dependence

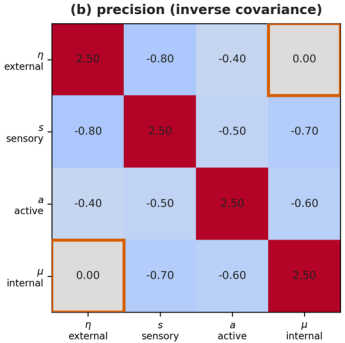
Gaussian precision exposes screening

CONDITION ON THE BLANKET

Marginal ≠ conditional dependence

Source: Synthetic / 05_markov_blanket.md

What a blanket does not establish



READ THE RESULT

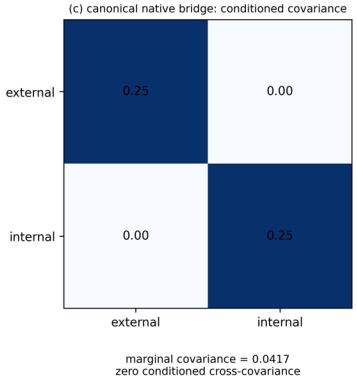
- Partial conditioning may not screen
- Causal direction needs more assumptions
- Representation is a separate claim

TEST THE BOUNDARY

Screening is a statistical statement

Source: Synthetic / 05_markov_blanket.md

A native Gaussian bridge has its own model



READ THE RESULT

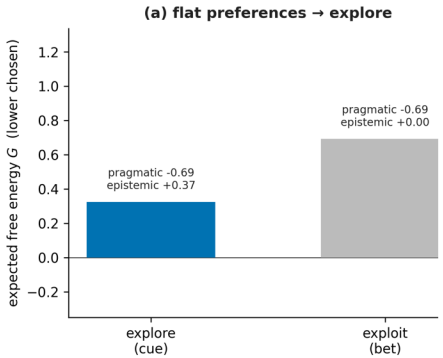
- A distinct specified precision matrix
- Conditioned cross-covariance is zero
- Perturbation supplies a counterexample

MATCH THE EXACT MODEL

One native Gaussian bridge

Source: Synthetic / 05_markov_blanket.md

Act before the outcome is known



READ THE RESULT

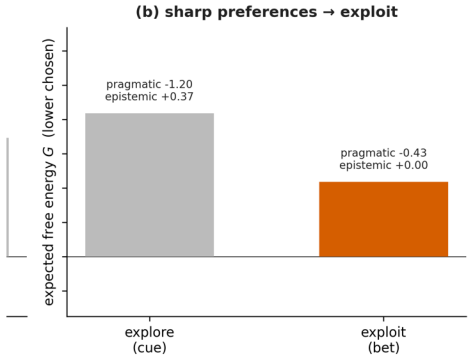
- Predict outcomes before observing
- Preferences score desirable outcomes
- Information scores posterior change

SCORE POSSIBLE ACTIONS

$$G = -\text{pragmatic} - \text{epistemic}$$

Source: Synthetic / 04_expected_free_energy.md

Preferences shift the explore-exploit balance



READ THE RESULT

- Flat preference: cue wins
- Concentrated preference: bet wins
- Matrices define this demonstration

CHANGE PREFERENCES

The same objective changes the choice

Source: Synthetic / 04_expected_free_energy.md

Expected information gain has a concrete meaning

PREDICT	CONDITION	AVERAGE
Possible observations	A posterior for each	Weight every outcome
q(o | action)	q(x | o, action)	Information gain

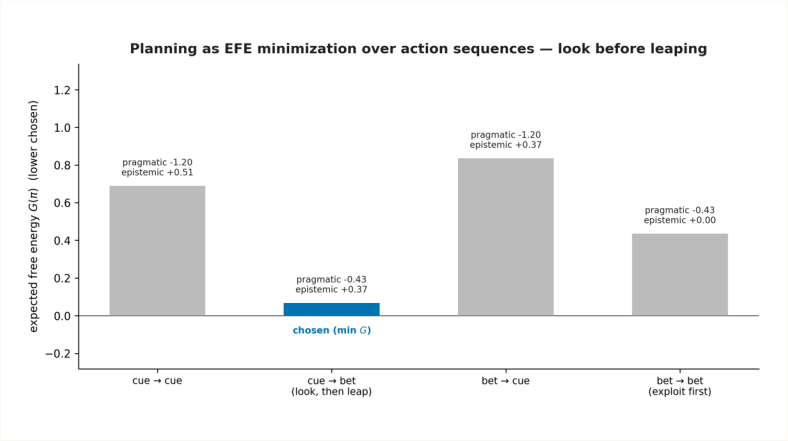
READ THE RESULT

- Predict every possible observation
- Condition a posterior for each
- Average divergence from the prior

AVERAGE OVER OBSERVATIONS

Information gain is expected updating

Planning evaluates whole observation branches



READ THE RESULT

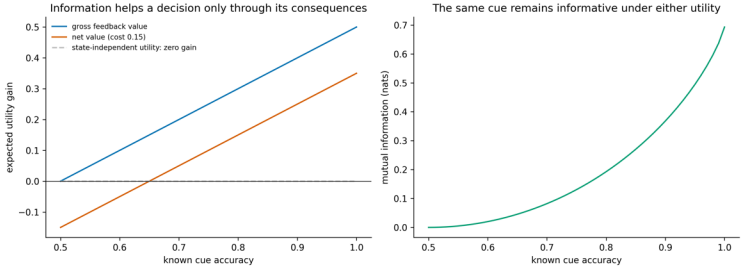
- Enumerate each observation branch
- Update beliefs along the branch
- Average after evaluating the branch

FOLLOW WHOLE HISTORIES

Terminal value + history information

Source: Synthetic / 08_planning.md

When information changes a decision



READ THE RESULT

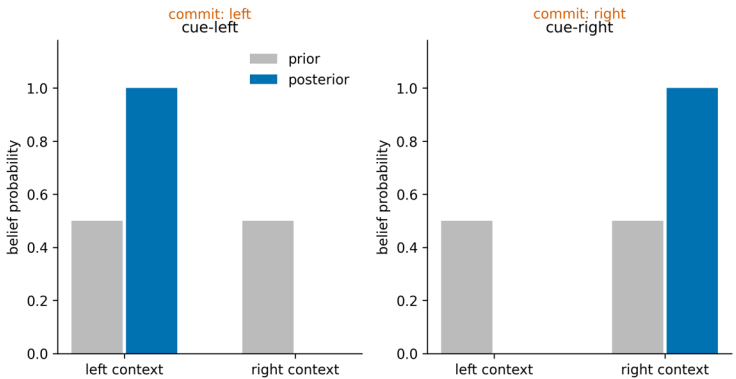
- Compare immediate and feedback choices
- Subtract the observation cost
- Keep utility and nats separate

CHECK INSTRUMENTAL VALUE

Information can change a commitment

Source: Synthetic / 09_value_of_information.md

Close the perception-action loop



READ THE RESULT

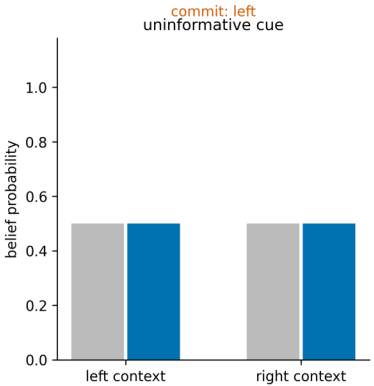
- Commitment has state-dependent payoff
- A cue changes the posterior
- A later decision uses that evidence

CLOSE THE LOOP

Observe → update → choose again

Source: Synthetic / 09_synthetic_agent.md

A useful control removes information



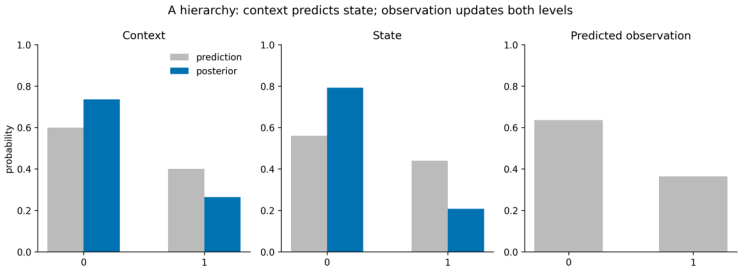
READ THE RESULT

- Replace the cue with a fair coin
- The posterior remains the prior
- Useful evidence changes beliefs

REMOVE THE INFORMATION

Observation alone is not learning

Hierarchical inference predicts down and updates up



READ THE RESULT

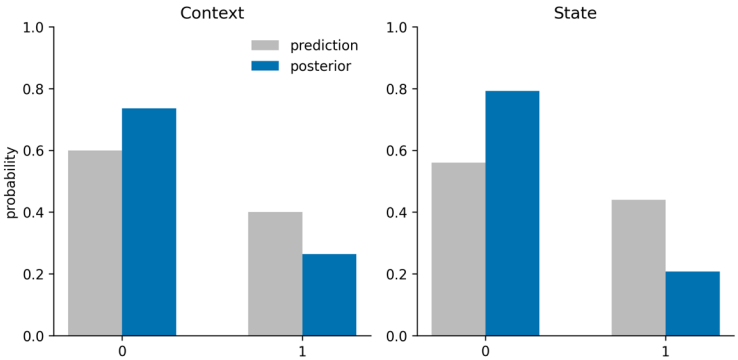
- Context predicts hidden state
- State predicts observation
- Joint conditioning revises both levels

TRACE THE TWO DIRECTIONS

Predict down; condition upward

Source: Synthetic / 10_hierarchical_inference.md

Learn from the limits of a hierarchy



READ THE RESULT

- Inspect both conditional models
- Check by direct joint enumeration
- Finite conditioning is not a neural claim

FIND THE MISSING LINK

An uninformative channel blocks updating

Finite path proofs state their support conditions

ACTUAL LEAN DECLARATION / PINNED FEP_FORMAL

```
theorem finite_path_integral_fluctuation {Path : Type*} [Fintype Path]
    (protocol : FinitePathProtocol Path)
    (hForward : ∀ path, 0 < protocol.forward path)
    (hReverse : ∀ path, 0 < protocol.reverseAligned path) :
    ∑ path, protocol.forward path *
      Real.exp (-pathwiseEntropyProduction protocol path) = 1 :=
  integralFluctuation_eq_one protocol hForward hReverse
```

READ THE RESULT

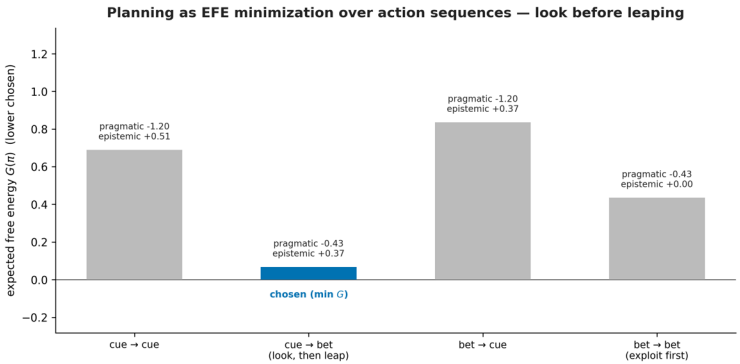
- Lean checks the stated hypotheses
- Numerics need a model correspondence
- Scientific meaning needs justification

LINK THE EVIDENCE

Theorem → model → interpretation

Source: Synthetic / 07_thermodynamic_transfer.md

A history channel bridges planning objectives



READ THE RESULT

- Treat observations as one outcome
- Normalize terminal preferences
- The added constant is policy-independent

NORMALIZE THE HISTORY

An explicit bridge preserves ranking

Source: Synthetic / 08_planning.md

GNN specifications meet numerical checks

CONTINUOUS OU

ou_regimes.md

x[2,type=float]

precision=1.0

gamma=1.0

omega_zero=0.0

omega_driven=2.4

FINITE CUE / BET

finite_agent.md

A_cue[2,2,type=float]

A_cue={(0.9,0.1),(0.1,0.9)}

A_bet={(0.85,0.85),(0.15,0.15)}

D={(0.5,0.5)}

C_probs={(0.5,0.5)}

Upstream parser + exact demo values + dimension checks

Negative controls: wrong gamma and state dimension reject

READ THE RESULT

Outcome rows; state columns

Drive Q_ou is not noise covariance

Wrong parameters and dimensions fail

READ THE SPECIFICATIONS

Parsing checks representation; execution needs a bridge

Your map of the Institute ecosystem

github.com/ActiveInferenceInstitute/

Prove	fep_formal
Specify	Generalized_Notation_Notation
Define	Active_Inference_Ontology
Learn	fundamentals
Study	courses
Compute	cognitive
Explore	ActiveInferenceJournal

READ THE RESULT

Synthetic: paper and Python examples

FEP_Formal: pinned Lean reference

GNN: optional model specification

YOUR NEXT STEP

Reproduce one model; change one assumption