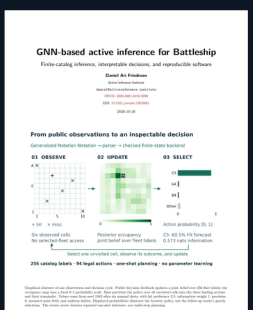
A game you can play, a model you can inspect

CHAPTER 01 A PUBLISHED, INSPECTABLE SYSTEM


GNN-based active inference for Battleship

Finite-catalog inference, interpretable decisions, and reproducible software

Daniel Ari Friedman

Active Inference Institute

Active Inference Institute, Japan & Irvine

ORCID: 0000-0001-6232-9098

DOI: 10.5281/zenodo.23676081

2026-04-26

From public observations to an inspectable decision

Generalized Notation Notation — parser → checked finite-state backend

01 OBSERVE

02 UPDATE

03 SELECT

Six observed cells
No selected fleet access

Posterior occupancy
joint belief over fleet labels

C5: 60.5% hit forecast
0.573 nats information

Select one unrivaled cell, observe its outcome, and update

256 catalog labels · 94 legal actions · one-shot planning · no parameter learning

Graphical abstract of one observation-and-decision cycle. Public hit/miss feedback updates a joint belief over 256 fleet labels; the posterior occupancy map and the selected-fleet belief are then used to choose the next cell. Six observed cells, 94 legal actions, and fleet remainder. Values come from real GNN after six scored shots, with full preference C3, information weight 1, precision 8, untrained prior. With real inference labels. Displayed probabilities illustrate the learned policy; not the follow-up study's greedy benchmark. The version never observes exposed one-shot inference, not multi-step planning.

PUBLIC SOURCE

github.com/ActiveInferenceInstitute/active_inference_battleship

MANUSCRIPT ARCHIVE

DOI 10.5281/zenodo.23676081

PLAY + READ

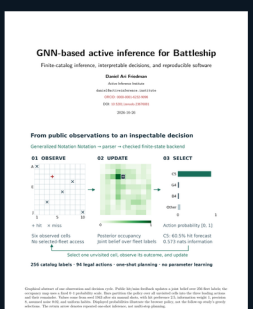
activeinferenceinstitute.github.io/active_inference_battleship/

The publication connects a playable interface to an explicit model.

Friedman (2026), p. 1 + released source | DOI 10.5281/zenodo.23676081 | Original explanatory animation

One release, connected public artifacts

CHAPTER 01 A PUBLISHED, INSPECTABLE SYSTEM


GNN-based active inference for Battleship

Finite-catalog inference, interpretable decisions, and reproducible software

Daniel Ari Friedman

Active Inference Institute

daniel@activeinference.institute

ORCID: 0000-0001-6232-9098

DOI: 10.5281/zenodo.23676081

2026-04-26

From public observations to an inspectable decision

Generalized Notation Notation — parser → checked finite-state backend

01 OBSERVE

02 UPDATE

03 SELECT

Six observed cells
No selected fleet access

Posterior occupancy
joint belief over fleet labels

C5: 60.5% hit forecast
0.573 nats information

Select one unrivaled cell, observe its outcome, and update

256 catalog labels · 94 legal actions · one-shot planning · no parameter learning

Graphical abstract of one observation-and-decision cycle. Public ish/ans feedback updates a joint belief over 256 fleet labels; the posterior occupancy map and the C5 hit forecast are computed from that belief; the selected cell is then observed, and the loop repeats. Yellow crosses mark sunk-ship cells, while hit predictions C5, information weight I, posterior b, and entropy H are used to select actions. Displayed probabilities illustrate the learned policy; see the follow-up study's greedy ablation. The results assume observe-learned one-shot inference, not sequential planning.

PUBLIC SOURCE

github.com/ActiveInferenceInstitute/active_inference_battleship

MANUSCRIPT ARCHIVE

DOI 10.5281/zenodo.23676081

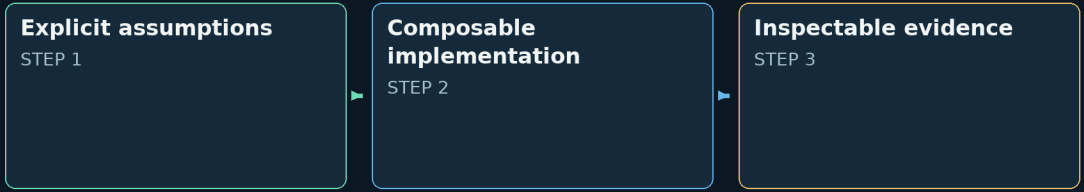
PLAY + READ

activeinferenceinstitute.github.io/active_inference_battleship/

The DOI identifies the manuscript and deposited companions.

What this work contributes

CHAPTER 01 A PUBLISHED, INSPECTABLE SYSTEM



Read the assumptions before interpreting the output.

The contribution is transparent software, not a new inference principle.

Friedman (2026), p. 6 + released source | DOI 10.5281/zenodo.23676081 | Original explanatory animation

An illustrated route through the paper

CHAPTER 01 A PUBLISHED, INSPECTABLE SYSTEM



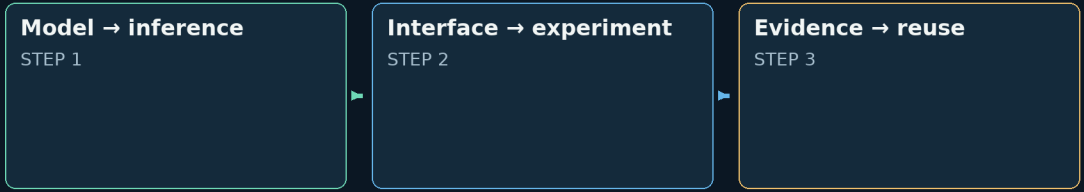
Follow the inputs, transformation, and resulting evidence.

The first four pages connect the publication to the working interface.

Friedman (2026), p. 1 + released source | DOI 10.5281/zenodo.23676081 | Original explanatory animation

A map for the next twenty minutes

CHAPTER 01 A PUBLISHED, INSPECTABLE SYSTEM



Read the assumptions before interpreting the output.

Ask what is known, what is optimized, and what is actually supported.

Friedman (2026), p. 6 + released source | DOI 10.5281/zenodo.23676081 | Original explanatory animation

The Battleship world

CHAPTER 02 THE GAME AND THE INFORMATION BOUNDARY

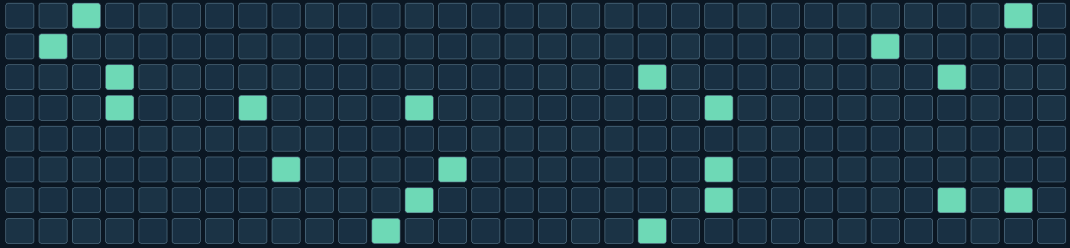


Shots reveal occupied or empty cells; fleets are hidden.

Friedman (2026), p. 8 + released source | DOI 10.5281/zenodo.23676081 | Original explanatory animation

A finite catalog defines the hypothesis space

256 COMPLETE-FLEET LABELS



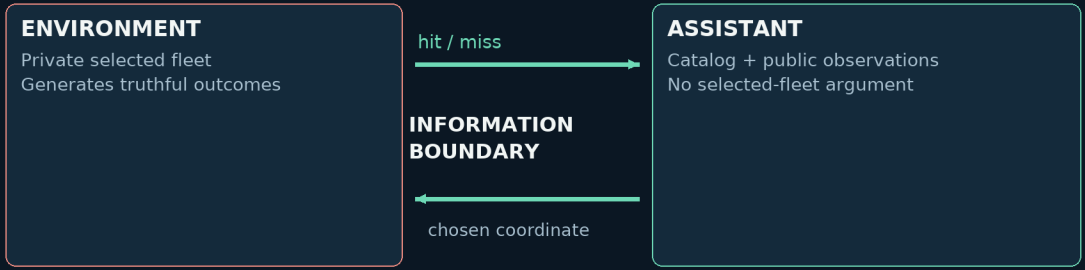
Relative weight / checkpoint maximum. These labels are whole fleets, not board locations.

Exact inference is conditional on the declared finite catalog.

Friedman (2026), p. 8 + released source | DOI 10.5281/zenodo.23676081 | Original explanatory animation

The environment knows; the assistant infers

CHAPTER 02 THE GAME AND THE INFORMATION BOUNDARY

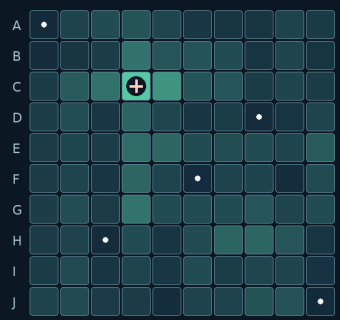


The assistant does not receive the selected fleet.

Friedman (2026), p. 11 + released source | DOI 10.5281/zenodo.23676081 | Original explanatory animation

A real six-shot observation history

CHAPTER 02 THE GAME AND THE INFORMATION BOUNDARY



OBSERVATIONS: 6 / 6

Joint-label entropy 3.511 nats

Effective states 24.8

Legal actions 94

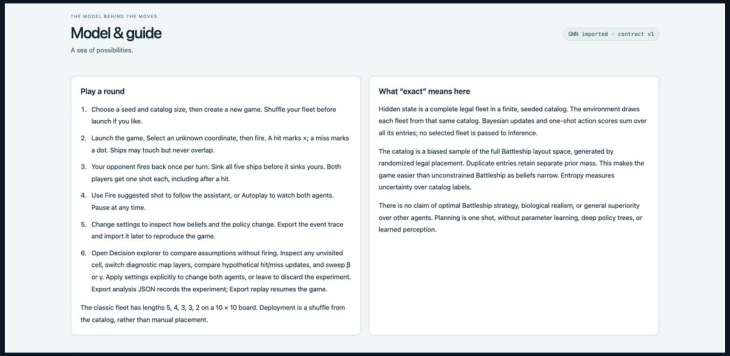
Exact checkpoints; interpolated visual transitions.

These belief states are recomputed from the published core.

Friedman (2026), p. 3 + released source | DOI 10.5281/zenodo.23676081 | Original explanatory animation

Play, inspect, pause, and export

CHAPTER 02 THE GAME AND THE INFORMATION BOUNDARY



THE MODEL BEHIND THE MOVES

Model & guide

A sea of possibilities.

GNN imported · contract v1

Play a round

1. Choose a seed and catalog size, then create a new game. Shuffle your fleet before launch if you like.
2. Launch the game. Select an unknown coordinate, then fire. A hit marks ×; a miss marks a dot. Ships may touch but never overlap.
3. Your opponent fires back once per turn. Sink all five ships before it sinks yours. Both players get one shot each, including after a hit.
4. Use Fire suggested shot to follow the assistant, or Autoplay to watch both agents. Pause at any time.
5. Change settings to inspect how beliefs and the policy change. Export the event trace and import it later to reproduce the game.
6. Open Decision explorer to compare assumptions without firing. Inspect any unvisited cell, switch diagnostic map layers, compare hypothetical hit/miss updates, and sweep β or γ. Apply settings explicitly to change both agents, or leave to discard the experiment. Export analysis JSON records the experiment; Export replay resumes the game.

The classic fleet has lengths 5, 4, 3, 3, 2 on a 10 × 10 board. Deployment is a shuffle from the catalog, rather than manual placement.

What "exact" means here

Hidden state is a complete legal fleet in a finite, seeded catalog. The environment draws each fleet from that same catalog. Bayesian updates and one-shot action scores sum over all its entries; no selected fleet is passed to inference.

The catalog is a biased sample of the full Battleship layout space, generated by randomized legal placement. Duplicate entries retain separate prior mass. This makes the game easier than unconstrained Battleship as beliefs narrow. Entropy measures uncertainty over catalog labels.

There is no claim of optimal Battleship strategy, biological realism, or general superiority over other agents. Planning is one shot, without parameter learning, deep policy trees, or learned perception.

Launch or shuffle

Fire manually or follow advice

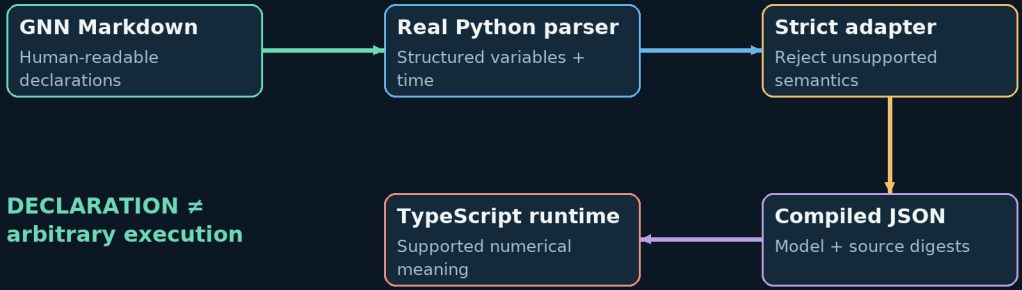
Pause autoplay; export replay

A replay resumes a game; an analysis export explains a state.

Friedman (2026), p. 3 + released source | DOI 10.5281/zenodo.23676081 | Original explanatory animation

GNN means Generalized Notation Notation

CHAPTER 03 FROM GNN TO EXECUTABLE MATHEMATICS

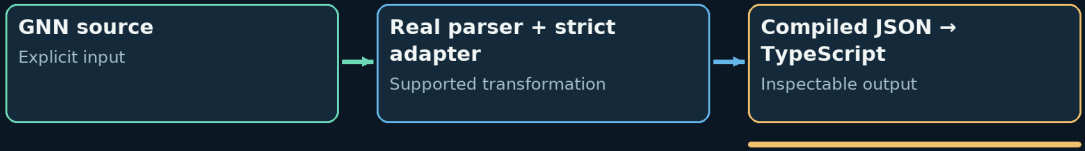


GNN specifies the supported model; it is not a decorative label.

Friedman (2026), p. 10 + released source | DOI 10.5281/zenodo.23676081 | Original explanatory animation

The specification-to-runtime path

CHAPTER 03 FROM GNN TO EXECUTABLE MATHEMATICS

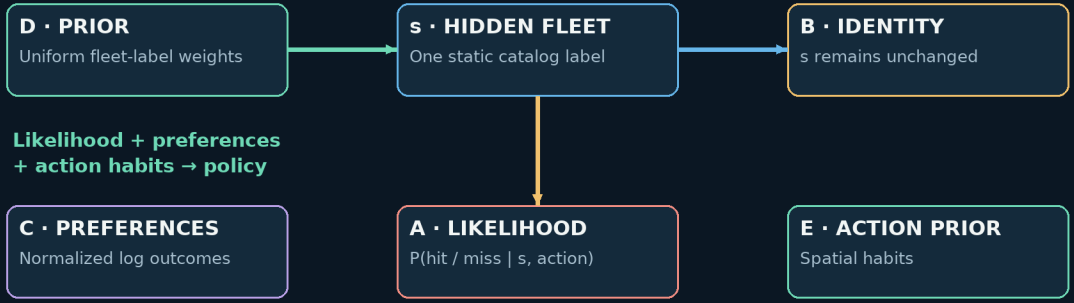


Follow the inputs, transformation, and resulting evidence.

Declaration, validation, compilation, and execution are separate stages.

Friedman (2026), p. 11 + released source | DOI 10.5281/zenodo.23676081 | Original explanatory animation

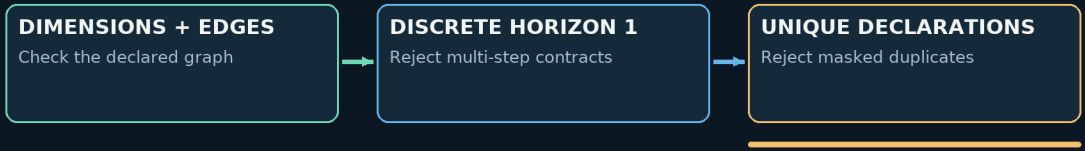
Five factors, one concrete interpretation



The factor names refer to this bounded static-fleet model.

Friedman (2026), p. 8 + released source | DOI 10.5281/zenodo.23676081 | Original explanatory animation

A contract that rejects incompatible models



Follow the inputs, transformation, and resulting evidence.

A source-bound build must reject unsupported semantics.

Friedman (2026), p. 10 + released source | DOI 10.5281/zenodo.23676081 | Original explanatory animation

A reusable core beyond the browser



Follow the inputs, transformation, and resulting evidence.

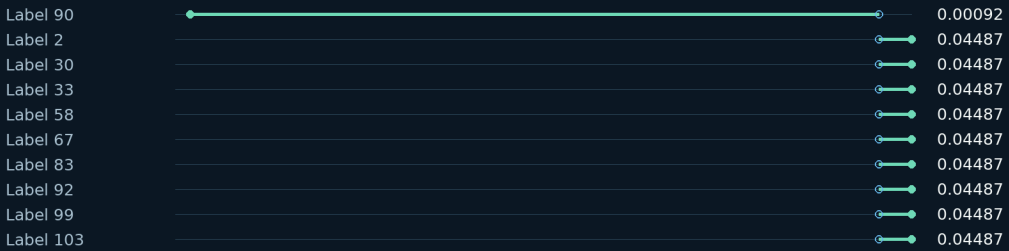
The browser is one interface to a composable public-information core.

Friedman (2026), p. 11 + released source | DOI 10.5281/zenodo.23676081 | Original explanatory animation

Bayesian updating over complete fleets

q_new(s) ∝ q_old(s) × P(observation | s, action)

Ten largest absolute changes · all 256 labels are normalized



○ Before shot 6 ● After shot 6 Scale: 0 to 0.0449

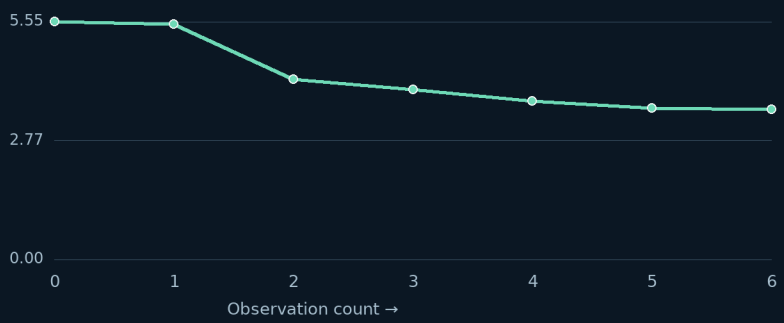
A heatmap is a marginal of a joint fleet posterior.

Friedman (2026), p. 8 + released source | DOI 10.5281/zenodo.23676081 | Original explanatory animation

Entropy, effective support, and surprise

CHAPTER 04 BELIEFS, INFORMATION, AND PREFERENCES

LABEL ENTROPY · REAL SIX-SHOT HISTORY



3.511 nats
Current label entropy

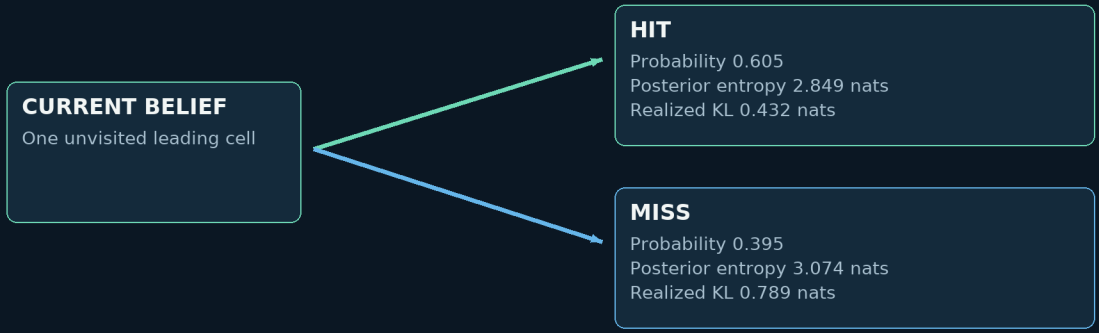
24.8
Effective states

These diagnostics describe different properties of the same belief.

Friedman (2026), p. 8 + released source | DOI 10.5281/zenodo.23676081 | Original explanatory animation

Expected information is not realized information

CHAPTER 04 BELIEFS, INFORMATION, AND PREFERENCES



Expected information averages possible outcomes before acting.

Friedman (2026), p. 10 + released source | DOI 10.5281/zenodo.23676081 | Original explanatory animation

Preferences and the one-step objective

CHAPTER 04 BELIEFS, INFORMATION, AND PREFERENCES

$$G(a) = R(a) - \beta\, I(a)$$

R: preference cost
Which outcomes do we prefer?

I: expected information
How much could the observation teach?

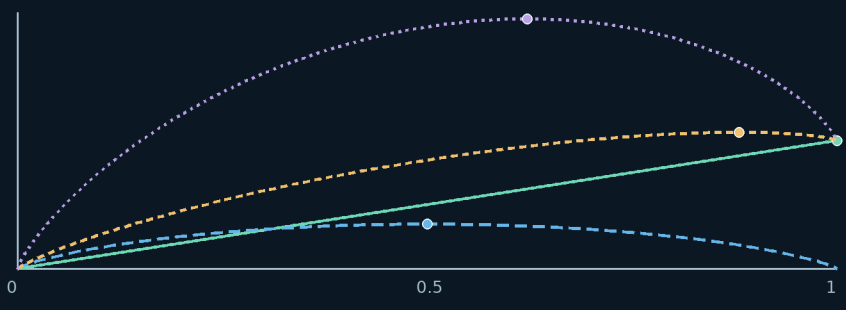
β: epistemic weight
How strongly does information matter?

Epistemic weight changes the trade-off, not the observation history.

Friedman (2026), p. 9 + released source | DOI 10.5281/zenodo.23676081 | Original explanatory animation

Why the selectors can prefer different cells

CONTINUOUS ANALYTIC UTILITIES — NOT EXPERIMENTAL OUTCOMES



Pragmatic: 2p
p* = 1.0000

Epistemic: h(p)
p* = 0.5000

Combined: 2p + h(p)
p* = 0.8808

Weighted: 2p + 4h(p)
p* = 0.6225

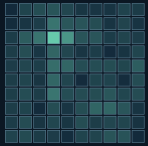
Continuous objective curves explain trade-offs, not measured dominance.

Friedman (2026), p. 20 + released source | DOI 10.5281/zenodo.23676081 | Original explanatory animation

Six maps answer six different questions

CHAPTER 05 USING THE DECISION EXPLORER

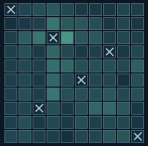
Occupancy



0 to 1

Higher value → brighter

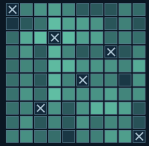
Hit forecast



0 to 1

Higher value → brighter

Information (nats)



0 to 0.693

Higher value → brighter

Preference cost



1.07 to 2.5

Lower is better

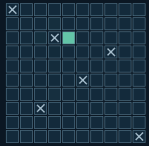
Objective G



0.494 to 2.46

Lower is better

Policy probability



0 to 1

Higher value → brighter

A color map is meaningful only with its quantity, units, and direction.

Friedman (2026), p. 4 + released source | DOI 10.5281/zenodo.23676081 | Original explanatory animation

Read an action receipt before firing

CHAPTER 05 USING THE DECISION EXPLORER

INSPECT C5 · BEFORE FIRING

Ship occupancy	**60.94%**
Predicted hit	**60.50%**
Expected information	**0.5729 nats**
Preference cost R	**1.0664 nats**
Objective G = R − βI	**0.4935 nats**
Choice probability	**87.83%**

IF HIT: 60.5%

Posterior entropy
2.849 nats

IF MISS: 39.5%

Posterior entropy
3.074 nats

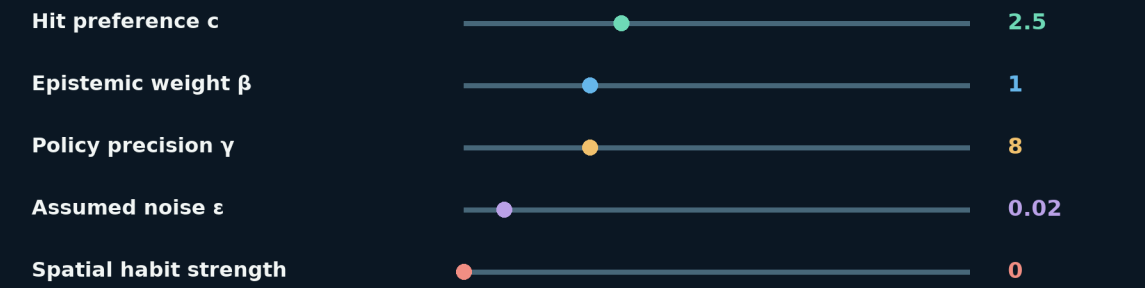
Hypotheses, not observed outcomes

The receipt explains an inspected action; inspection does not fire it.

Friedman (2026), p. 4 + released source | DOI 10.5281/zenodo.23676081 | Original explanatory animation

Five controls, distinct effects

CHAPTER 05 USING THE DECISION EXPLORER



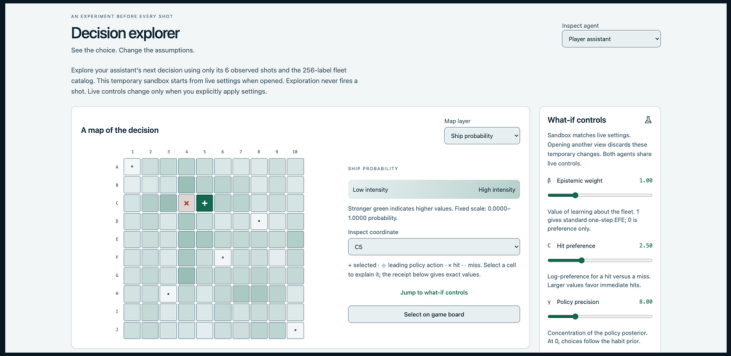
Recorded illustrative settings, not a calibration of the real world.

Controls encode assumptions and objectives, not calibrated truths.

Friedman (2026), p. 10 + released source | DOI 10.5281/zenodo.23676081 | Original explanatory animation

A sandbox separates exploration from execution

CHAPTER 05 USING THE DECISION EXPLORER



Copy live controls

Compare temporary settings

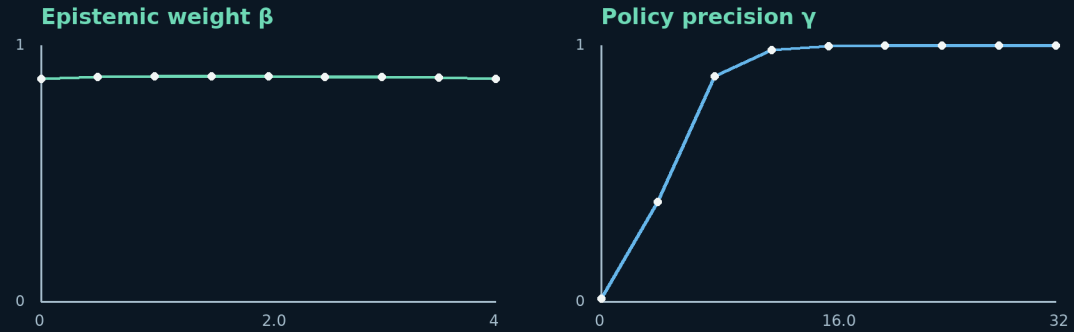
Apply explicitly or discard

A what-if result is not a new observation or experiment outcome.

Friedman (2026), p. 4 + released source | DOI 10.5281/zenodo.23676081 | Original explanatory animation

Sensitivity is conditional on what is held fixed

CHAPTER 05 USING THE DECISION EXPLORER



Probability of inspected C5; same six-shot belief and other settings held fixed.

A sensitivity curve is local analysis, not contingent planning.

Friedman (2026), p. 10 + released source | DOI 10.5281/zenodo.23676081 | Original explanatory animation

From a demonstration to a paired follow-up

CHAPTER 06 FIVE STRATEGIES IN A PAIRED EXPERIMENT

M = 256 | 128 paired units



M = 4,096 | 128 paired units



Each tile is one catalog/truth unit. Its five stripes are the five profiles on that SAME unit.

The follow-up is prospective within the project, not publicly preregistered.

Friedman (2026), p. 12 + released source | DOI 10.5281/zenodo.23676081 | Original explanatory animation

Five strategies, shared cases

CHAPTER 06 FIVE STRATEGIES IN A PAIRED EXPERIMENT

Random	Pragmatic	Epistemic	Combined	Weighted
Uniform legal cell	Maximize hit forecast	Maximize information	Minimize G, β = 1	Minimize G, β = 4
Same paired case	Same paired case	Same paired case	Same paired case	Same paired case

Pairing compares profiles on the same catalog and selected truth.

Friedman (2026), p. 13 + released source | DOI 10.5281/zenodo.23676081 | Original explanatory animation

Three endpoints prevent a one-number verdict

CHAPTER 06 FIVE STRATEGIES IN A PAIRED EXPERIMENT

COMPLETION SHOTS

24

Lower is preferable

HITS AFTER 5 SHOTS

3

Higher is preferable

LABEL ENTROPY AT 5

0.5462

H / ln(M)

Lower is preferable

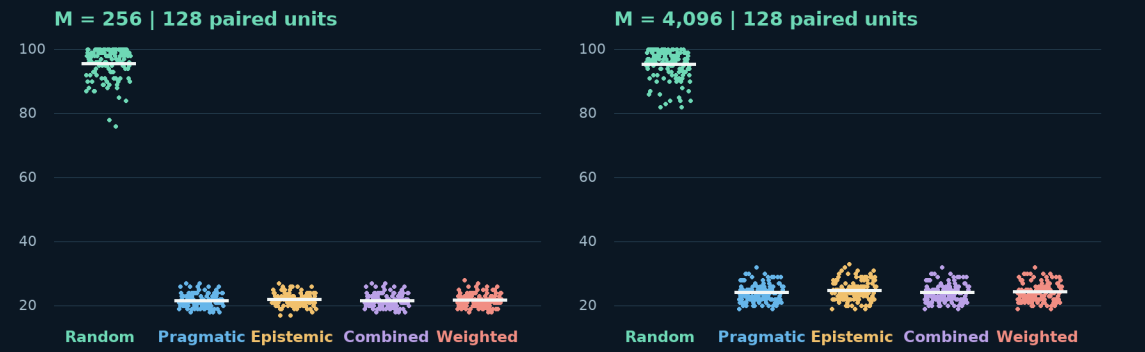
One illustrative retained unit: M = 4,096, case 0, pragmatic. Group analysis uses all 128 units.

Completion, exploitation, and discrimination are different goals.

Friedman (2026), p. 14 + released source | DOI 10.5281/zenodo.23676081 | Original explanatory animation

Completion: random versus informed selection

CHAPTER 06 FIVE STRATEGIES IN A PAIRED EXPERIMENT



In these catalog-defined cases, random search takes many more shots.

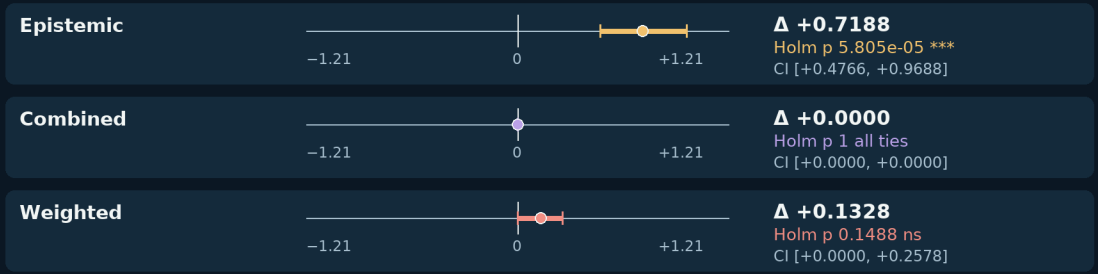
Friedman (2026), p. 15 + released source | DOI 10.5281/zenodo.23676081 | Original explanatory animation

The informative comparison is within informed profiles

CHAPTER 06 FIVE STRATEGIES IN A PAIRED EXPERIMENT

PROFILE − PRAGMATIC · COMPLETION | M = 4,096

Points: paired means. Lines: 95% descriptive percentile intervals.



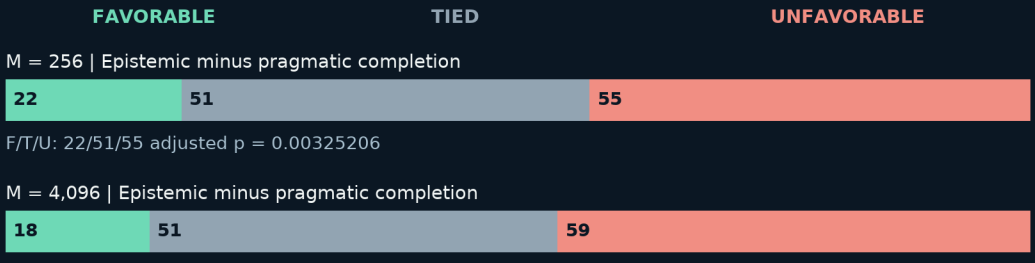
Completion rows share the same shots scale; sign balance and mean magnitude differ.

Information seeking does not automatically minimize shots to completion.

Friedman (2026), p. 18 + released source | DOI 10.5281/zenodo.23676081 | Original explanatory animation

Stars need a question and a direction

CHAPTER 07 WHAT THE EVIDENCE DOES AND DOES NOT SAY



Stars indicate conditional sign imbalance — here it favors the pragmatic baseline.

A significance marker can favor the baseline, not the highlighted profile.

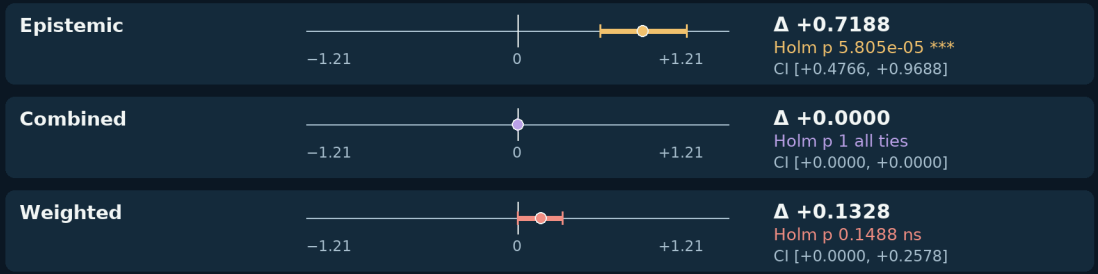
Friedman (2026), p. 13 + released source | DOI 10.5281/zenodo.23676081 | Original explanatory animation

Effect intervals and sign tests tell different stories

CHAPTER 07 WHAT THE EVIDENCE DOES AND DOES NOT SAY

PROFILE − PRAGMATIC · COMPLETION | M = 4,096

Points: paired means. Lines: 95% descriptive percentile intervals.



Completion rows share the same shots scale; sign balance and mean magnitude differ.

A mean interval is not the inversion of the conditional sign test.

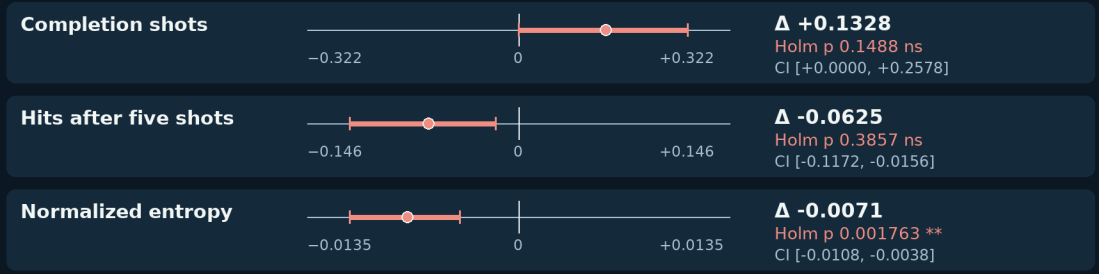
Friedman (2026), p. 16 + released source | DOI 10.5281/zenodo.23676081 | Original explanatory animation

Early uncertainty can improve without a completion claim

CHAPTER 07 WHAT THE EVIDENCE DOES AND DOES NOT SAY

WEIGHTED − PRAGMATIC | M = 4,096

Points: paired means. Lines: 95% descriptive percentile intervals.



Endpoint units differ; each row uses its own labeled scale.

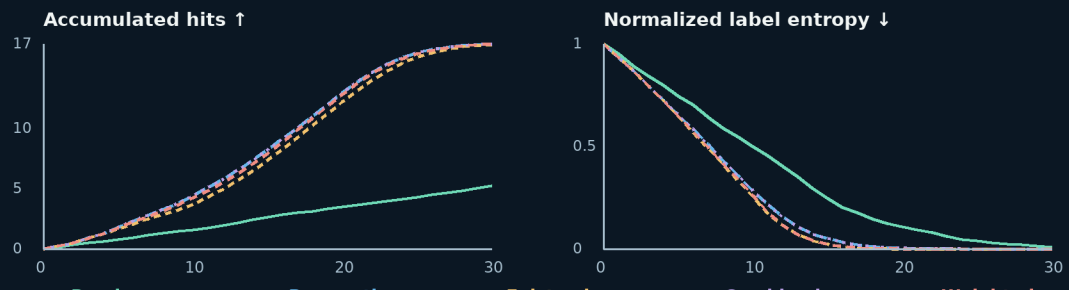
Do not transfer a result from one endpoint to another.

Friedman (2026), p. 18 + released source | DOI 10.5281/zenodo.23676081 | Original explanatory animation

Watch the trajectories, not just the finish line

CHAPTER 07 WHAT THE EVIDENCE DOES AND DOES NOT SAY

M = 4,096 | ALL 128 CASES | DESCRIPTIVE MEANS

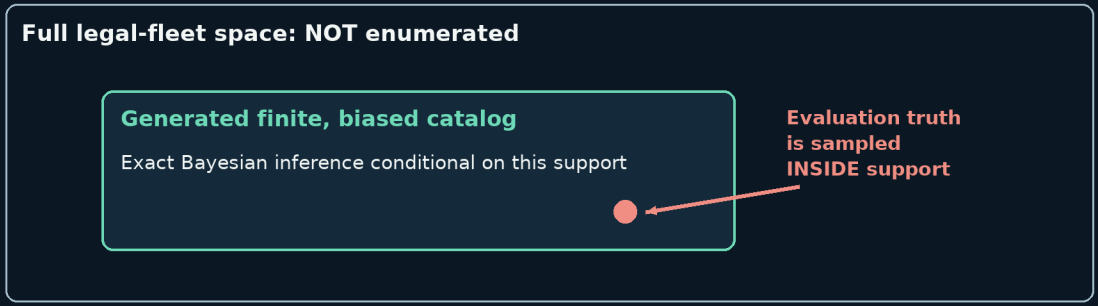


The curves are descriptive summaries, with no curve-wise significance test.

Friedman (2026), p. 17 + released source | DOI 10.5281/zenodo.23676081 | Original explanatory animation

The strongest limits are part of the result

CHAPTER 07 WHAT THE EVIDENCE DOES AND DOES NOT SAY



Explanatory set diagram; area is not a measured fraction of all possible fleets.

Conditional exactness does not establish unrestricted or real-world optimality.

Friedman (2026), p. 21 + released source | DOI 10.5281/zenodo.23676081 | Original explanatory animation

A connected scholarly lineage

CHAPTER 08 REPRODUCE, EXTEND, AND BUILD ON THE RELEASE

Shannon (1948)	Entropy and communication
Houlsby et al. (2011)	Bayesian active learning
Garnett et al. (2012)	Bayesian active search
Friston et al. (2015)	Epistemic value and active inference
Rothe et al. (2017–2019)	Questions, goals, and spatial learning

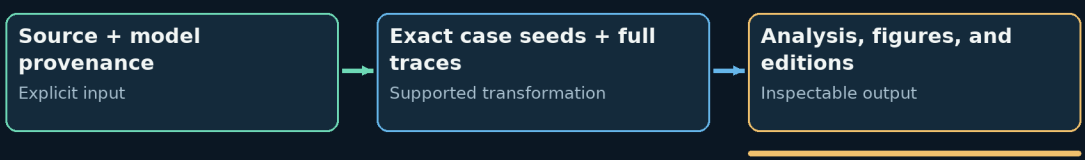
34 bibliography entries; the reading ledger discloses source-access and task limits.

Prior papers have different tasks, feedback, and objectives.

Friedman (2026), p. 6 + released source | DOI 10.5281/zenodo.23676081 | Original explanatory animation

Reproduce the right thing

CHAPTER 08 REPRODUCE, EXTEND, AND BUILD ON THE RELEASE



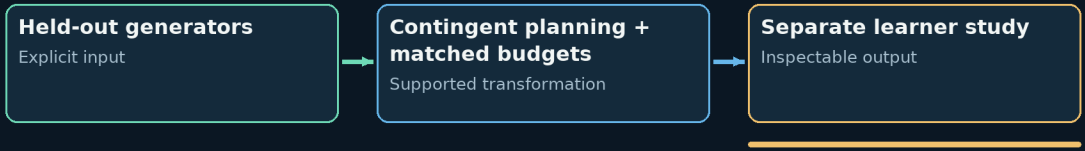
Follow the inputs, transformation, and resulting evidence.

A repeated deterministic run verifies the same cases; it does not add data.

Friedman (2026), p. 21 + released source | DOI 10.5281/zenodo.23676081 | Original explanatory animation

Design a stronger next experiment

CHAPTER 08 REPRODUCE, EXTEND, AND BUILD ON THE RELEASE



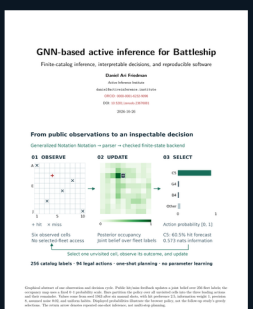
Follow the inputs, transformation, and resulting evidence.

Broader claims require broader tests, not louder visualization.

Friedman (2026), p. 22 + released source | DOI 10.5281/zenodo.23676081 | Original explanatory animation

Reuse the work, preserve the boundaries

CHAPTER 08 REPRODUCE, EXTEND, AND BUILD ON THE RELEASE


GNN-based active inference for Battleship

Finite-catalog inference, interpretable decisions, and reproducible software

Daniel Ari Friedman

Active Inference Institute

daniel@activeinference.institute

ORCID: 0000-0001-6232-9098

DOI: 10.5281/zenodo.19878085

2026-04-26

From public observations to an inspectable decision

Generalized Notation Notation — parser → checked finite-state backend

01 OBSERVE

02 UPDATE

03 SELECT

C1

C2

C3

C4

C5

0

1

× hit × miss

0

Action probability (0–1)

Six observed cells
No selected-fleet access

Posterior occupancy
Joint belief over fleet labels

C5: 40.5% hit forecast
0.573 nats information

Select one unrivaled cell, observe its outcome, and update

256 catalog labels · 94 legal actions · one-shot planning · no parameter learning

Graphical abstract of one observation-and-decision cycle. Public hit/miss feedback updates a joint belief over 256 fleet labels; the posterior occupancy map and the C5 hit forecast then select the next unrivaled cell. Six observed cells, no selected-fleet access, and 0.573 nats information. Palette: cream focus, coral C5 and hit/miss, teal posterior C2, ochre/amber C3, indeterminate weight C4, pewter A, outcome dots. Hit, miss, and uniform labels. Displayed pseudoMermaid illustrates the factored policy; not the follow-up study's grouped architecture. The version never observes opponent row-clue references, nor multi-step planning.

SOFTWARE
MIT license

MANUSCRIPT + ORIGINAL FIGURES
Creative Commons Attribution 4.0

CITATION
Manuscript DOI + exact source commit

The public archive supports reuse without enlarging the scientific claims.

Friedman (2026), p. 22 + released source | DOI 10.5281/zenodo.23676081 | Original explanatory animation

From a hidden fleet to an accountable decision

What does the agent know?

What does its objective value?

What would change your conclusion?

PLAY + INSPECT

Play, inspect, reproduce, and question the assumptions.

Friedman (2026), p. 22 + released source | DOI 10.5281/zenodo.23676081 | Original explanatory animation