TopolAI: A Supervisory Control Architecture for Scientific AI

Sufficiency, recoverability, structural memory, acquisition, and reassessment

M. Antony Ewing, PhD
Conquer Medical Health – CMH Scientific
ORCID: 0009-0005-1550-1718

Technical Report v1.1 – 10 September 2026

Abstract

TopolAI is the supervisory scientific-AI architecture of CMH Scientific. It coordinates distinct evidence states – information sufficiency, tested-representation recoverability, structural memory, inverse measurement design, and reassessment – without collapsing them into a single score. The September 2026 public release exposes a generic callable client while protected routing, estimators, compiler logic, and acquisition policies remain outside the public distribution. This revision clarifies that TotemicAI is a fitted structural-memory model with two policy modes: TotemicAI-Specified, in which the landmark policy is fixed and the memory is compiled from data, and TotemicAI-Learned, in which policy parameters are fitted on training data before landmark compilation. Current AML evidence validates the specified mode. TopolAI can integrate evidence-supported actions including use, continue, selectively acquire, change measurement, add modality, reacquire, abstain, or remain unresolved.

v1.1 clarification. TopolAI is a supervisory controller, not a predictive model family. Its component instruments can be used independently. TotemicAI-Specified is the currently validated structural-memory embodiment; TotemicAI-Learned is an architectural extension requiring separate nested held-out validation.

1 Purpose and architecture

TopolAI changes what happens around a predictive model. Rather than beginning and ending with model performance, it keeps separate the evidence required to answer: does the substrate distinguish the task, how much appears recoverable, what structural history must be retained, what additional measurement has value, and what should the scientific system do next?

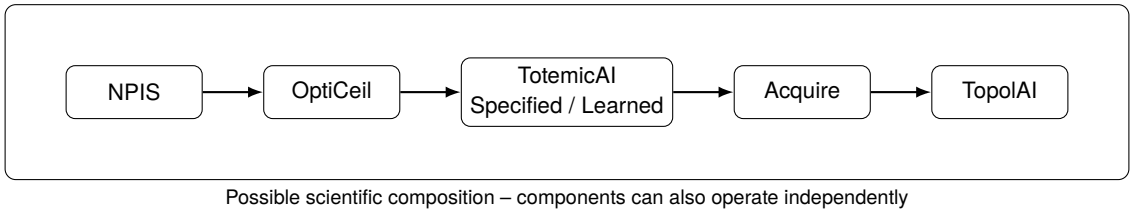

Figure 1: A common TopolAI composition. AIO, ABP, CTF, Deferra, and CMH Guard can connect where the task requires them; they are not mandatory serial stages.

The controller preserves several distinctions. Structural stability is not information sufficiency. A high recoverability envelope does not mean the current observer has reached it. Residual ambiguity does not identify an untested modality as the remedy. A biologically plausible intervention is not automatically causally identified. A scientifically useful action may still be blocked by governance policy.

2 Component states

Component	Primary question	Supervisory contribution
NPIS	Does the observed substrate distinguish the claimed states or futures?	Branching, indeterminacy, temporal refinement, departure/change, residual predictive availability.
OptiCeil	How much performance appears recoverable from the declared representation?	Demonstrated lower performance, conservative empirical envelope, extraction gap, transport diagnosis.
TotemicAI	Which observations preserve nonredundant structure?	Fitted landmark memory, coherence/challenge state; specified or learned policy mode.
Acquire	What should be observed next?	Ranked measurement family, protocol, patient, condition, window, modality, or acquisition frontier.
TopolAI	What should the system do next?	USE, CONTINUE, CONTINUE_SELECTIVELY, CHANGE_MEASUREMENT, ADD_MODALITY, REACQUIRE, ABSTAIN, UNRESOLVED.

Table 1: TopolAI supervises distinct evidence objects rather than replacing them with a single scalar.

3 TotemicAI policy modes inside TopolAI

TotemicAI-Specified uses a policy fixed before held-out evaluation; the response field determines which observations become landmarks. The resulting landmark memory is fitted from data. This is the mode supported by the released AML validation.

TotemicAI-Learned permits policy parameters to be estimated on training data before memory compilation. A learned-policy result must therefore be nested inside the training side of each held-out fold. TopolAI should report the mode explicitly because "structural memory validated" must not silently transfer the specified-policy evidence to a learned policy that has not yet been tested.

A TopolAI result involving structural memory should therefore include at least: policy mode, retention budget, validation unit, structural endpoints, comparator, coherence/challenge state, and provenance.

4 Two-layer public release

The public `TopolAIClient` accepts user-defined data, task, context, and options and returns a standardized `TopolAIResult`. The public package does not distribute internal routing tables, NPIS estimators, recoverability estimators, structural-memory compiler internals, acquisition-ranking logic, stopping policies, or sensor-selection policies. Frozen validation evidence is kept separate from generic user analysis.

5 Step-by-step use

1. Declare the scientific task and supplied substrate.
2. Audit whether the representation supports the target distinction.
3. Estimate tested recoverability and observer headroom when relevant.
4. Preserve structural history when the task requires streaming or sparse memory; report whether TotemicAI-Specified or TotemicAI-Learned is being used.
5. Rank only tested candidate acquisitions or return an unresolved state.
6. Route the evidence into an explicit next action.
7. Reassess after new observation, intervention, or structural challenge.

Python public call

```python
from topolai import TopolAIClient

client = TopolAIClient(transport=my_authorized_transport)
result = client.assess(
    data=my_data,
    task={"id": "what_should_i_measure_next",
          "objective": "assess_information_support"},
    context={"data_mode": "longitudinal"},
    options={"return_component_evidence": True},
)
print(result.state)
print(result.components_reported)
```

Example prompt

Assess this task end to end. Separate substrate sufficiency from observer headroom, identify whether structural memory is required, compare only tested acquisition options, and return UNRESOLVED rather than inventing a measurement or intervention that has not been evaluated.

6 Validation program

Substrate	Released supervisory evidence
Calcium imaging	914 recording segments, 312 neurons, 146,962 balanced samples; the acquisition program showed task- and indicator-dependent recoverability and validated temporal acquisition prescriptions under held-out neurons/datasets.
Primary AML Perturb-seq	TotemicAI-Specified structural selections outperformed equal-size random memory on direction, relational geometry, and topology across all nine primary comparisons. The learned-policy variant is not included in this validation claim.
PRO-ACT ALS	8,260 participants and 70,746 transitions; increasing the original ALSFRS discovery substrate produced little held-out gain, while tested additional measurement families added more information.
Cross-replicate perturbation	Response-informed acquisition order reproduced across replicate directions and beat a spacing comparator but was not globally optimal under exhaustive ordering.
Longitudinal COPD	Leakage-safe selective acquisition showed heterogeneous value of additional longitudinal and integrated information and concentrated much of that value in selected subjects.

Table 2: Component evidence inherited by the TopolAI supervisory architecture.

7 Connections to the wider CMH Scientific suite

TopolAI is not the commercial wrapper around every CMH product. AIO maps and reconstructs literature claims; CTF applies placebo falsification to responder-discovery logic; ABP performs individual-first inverse intervention interrogation; Deferra governs selective assistance in a physical or physiologic feedback loop; CMH Guard applies data/code/model/compute/export policy. Any of these can exchange structured evidence with TopolAI where the task requires it.

In particular, Deferra is a downstream clinical-assistance controller rather than a mandatory part of scientific measurement design. ABP is a patient-specific biological/intervention model rather than simply another acquisition stage. This separation keeps product roles intelligible.

8 Controller-state discipline

The public architecture recognizes USE, CONTINUE, CONTINUE_SELECTIVELY, CHANGE_MEASUREMENT, CHANGE_PROTOCOL, ADD_MODALITY, REACQUIRE, ABSTAIN, and UNRESOLVED as semantic output states. No single metric forces a state. TopolAI should expose the evidence supporting the state and the domain in which that evidence was tested.

9 Boundaries

TopolAI does not prove causal mechanism from observational structure, guarantee that an untested sensor will repair a deficit, establish a universal information ceiling, or imply that a learned structural-memory policy is validated because the specified policy performed well. The architecture is intended to make those unresolved states explicit.

Intellectual property, competing interests, funding, and use

Relevant methods and computational systems are the subject of previously filed U.S. provisional patent applications. This report discloses scientific objects, public interfaces, empirical validation logic, and results needed to evaluate the public research architecture while withholding protected operational implementations. Publication does not grant a patent or commercial software license.

The author is Founder of Conquer Medical Health and has financial and intellectual-property interests in the technologies described. No external funding supported preparation of this technical report.

Use of AI-assisted tools

AI-assisted tools were used for editorial and software-documentation support. The author determined the scientific claims, supplied the underlying analyses and validation artifacts, and reviewed the numerical content and final report.

References

[1] Ewing, M. A. *Widespread Nonidentifiability in Patient Data Limits AI Prediction Regardless of Model Sophistication*. 2026.

[2] Ewing, M. A. *TotemicAI: Fitted Structural Memory for Scientific AI*. Conquer Medical Health, 2026.

[3] Ewing, M. A. *OptiCeil: Task-Conditioned Recoverability*. Conquer Medical Health, 2026.

[4] Ewing, M. A. *Acquire: Inverse Measurement Design for Scientific AI*. Conquer Medical Health, 2026.

[5] Ewing, M. A. *CMH Scientific: R&D Laboratory and Scientific AI Instrument Suite*. Conquer Medical Health, 2026.