# AIO: Reconstructing and Testing Scientific Literatures on Common Substrates

## Evidence mapping, executable signatures, matched reconstruction, and information characterization

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

Technical Report v1.0 – 4 September 2026

**Abstract**

AIO (AI On...) is the evidence-mapping and common-substrate reconstruction layer of CMH Scientific. It converts literature claims into auditable empirical signatures, reconstructs compatible studies under matched conditions, compares reported and reproduced evidence, and can invoke the protected NPIS capability internally to determine which distinctions a supplied substrate supports. The revised public package is a generic callable client rather than a distribution of the proprietary research engine. Frozen AI-on-Alzheimer validation screened 95,756 publications, identified 10,884 data-bearing/hypothesis-relevant rows, recovered 183 recurring paper-linked biomarker configurations, executed 204 ADNI-supported signatures across four historical cutoffs and two cohort definitions, and produced 1,632 evaluation cells (about 1.224 million scheduled model operations). Across 327 papers, median reported AUC was 0.860 versus 0.678 under matched common-substrate reconstruction, with 271 mismatches.

**Public-release note.** This revision aligns the technical report with the September 2026 two-layer software release: a generic callable public interface for new analyses and a separate frozen validation layer for previously reported evidence. Protected operational implementations are not distributed.

## 1 Purpose

AIO turns a scientific literature into an auditable executable evidence map. Language models may assist with reading and structuring public text, but deterministic scientific computation is kept separate from language generation. AIO maps claims, organizes recurring hypothesis structure, builds empirical signatures, aligns them to a supplied substrate, reconstructs compatible studies under matched protocols, compares reported with reproduced results, and calls the protected NPIS capability when information characterization is requested.

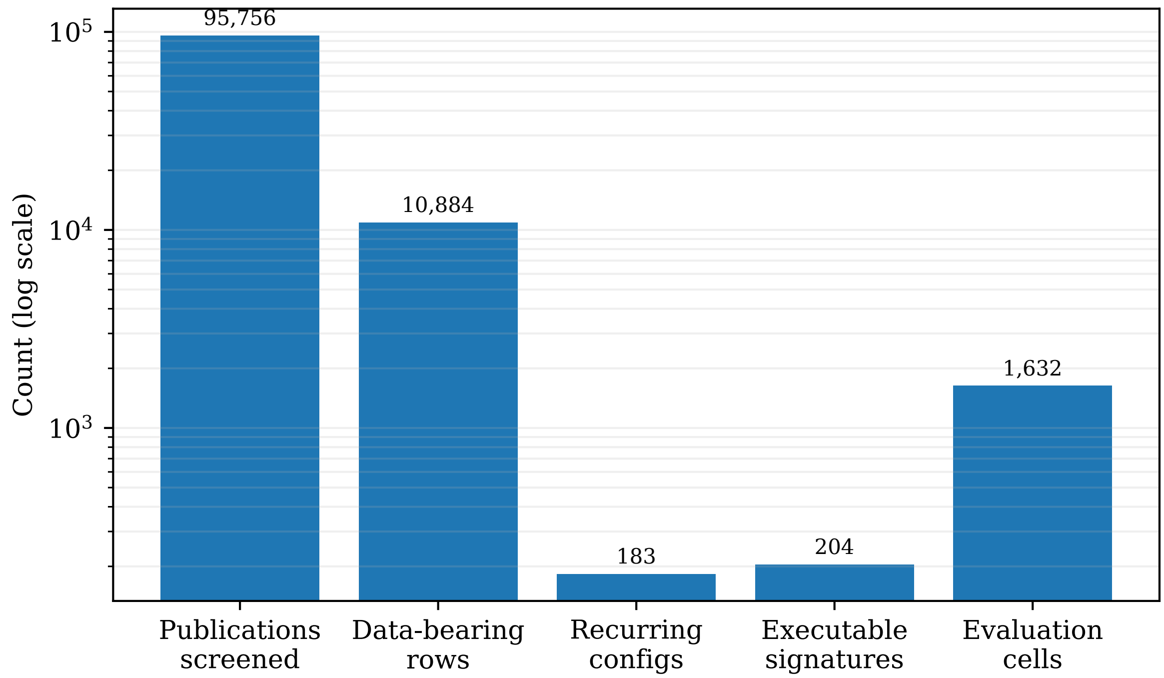


**Figure 1:** Scale of the frozen AI-on-Alzheimer precursor validation.

## 2 Two-layer public release

The public `AIOClient` accepts a domain, research question, literature corpus/reference, optional empirical data, mapping, target, grouping definition, and adjustable options. It returns an `AIOResult` from an authorized engine endpoint. The repository intentionally does not distribute the original `aio\_engine/` implementation or its vendored NPIS engine. Frozen Alzheimer audit values are isolated in `validation/`.

## 3 Step-by-step use

1. **Declare the research field and adjudication question.** Define what competing claims or empirical distinctions should be compared.
2. **Supply or identify the literature corpus.** Literature acquisition and extraction can be performed by the protected engine; the public client accepts a corpus or reference.
3. **Supply the common substrate and scientific mapping.** Map literature variables/endpoints to the empirical dataset; unresolved mappings must remain unresolved rather than be fabricated.
4. **Declare target and independence unit.** Define outcome, horizon, grouping, and validation design.
5. **Run matched reconstruction.** The protected AIO engine performs extraction, signature construction, model comparison, statistical adjudication, and optionally NPIS internally.
6. **Interpret the evidence map.** Distinguish claims that reproduce under the common protocol from those that collapse, fail to map, or cannot be distinguished by the substrate.

**Python public call**

```python
from aio import AIOClient

client = AIOClient(transport=my_authorized_transport)

result = client.analyze(
    domain="autism",
    question="Which hypotheses reproduce on a common substrate?",
    literature=my_literature,
    data=my_data,
    mapping=my_mapping,
    target={"column": "outcome"},
    grouping={"column": "patient_id"},
    options={
        "cv_folds": 5,
        "information_characterization": True,
    },
)

print(result.state)
print(result.evidence)
print(result.summary)
```

**NPIS dependency.** AIO users do *not* need to install the standalone public NPIS package merely because AIO uses NPIS. The authorized AIO engine can invoke the protected NPIS capability internally. The separate NPIS repository exists for researchers who want NPIS as a standalone analysis.

**Example prompt**

AI On Autism: map the major competing hypotheses, identify recurring quantitative signatures, reconstruct compatible studies on my supplied cohort under one endpoint and validation protocol, compare reported and reproduced results, and state which distinctions the cohort can actually support.

**Structured result**

**RESULT:** Competing literature claims are converted into executable signatures and evaluated on common empirical ground; some distinctions survive matched reconstruction while others collapse into the same information support.
**SUPPORTED:** claims that reproduce and remain distinguishable under the declared common protocol.
**NOT SUPPORTED:** unmapped variables/endpoints, causal mechanism not represented by the substrate, or numerical conclusions generated by language-model prose alone.
**NEXT ACTION:** when key claims remain unresolved, pass tested candidate measurements to Acquire.

## 4 AI-on-Alzheimer precursor validation

The frozen program screened 95,756 publications, identified 10,884 data-bearing or hypothesis-relevant rows, compressed them into 183 recurring paper-linked biomarker configurations, and executed 204 ADNI-supported signatures across four historical cutoffs (2012, 2016, 2020, 2025) and two cohort definitions, producing 1,632 evaluation cells. The preserved execution trace contains 10,884 trace rows and 183 trace signatures and recovers six hypothesis camps.

A grouped six-marker panel reached AUC 0.845 versus 0.767 for the strongest single canonical marker, a difference of +0.078 with bootstrap interval [0.019, 0.135]. Across 327 papers, median reported AUC was 0.860 while matched common-substrate AUC was 0.678, with 271 mismatches (Figure 2).

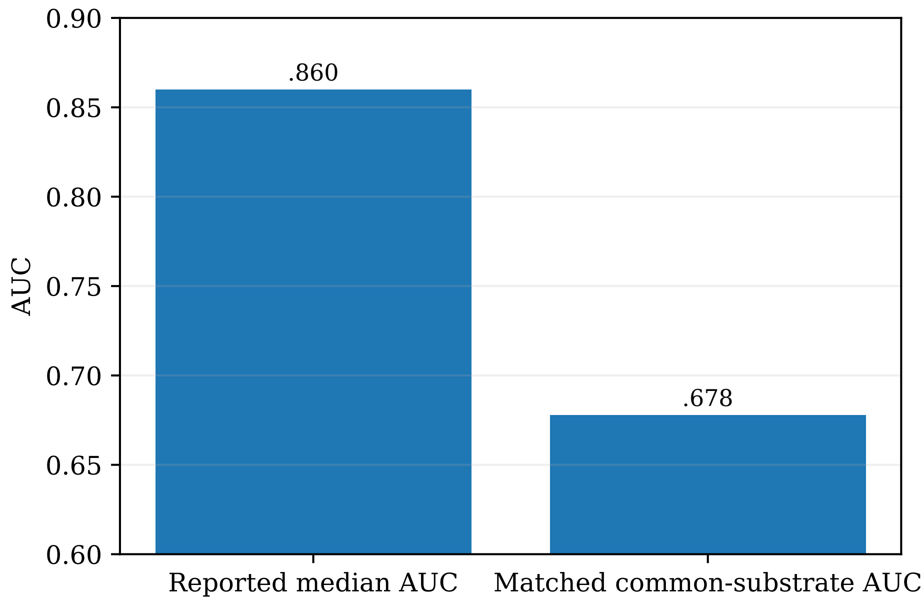


**Figure 2:** Reported versus matched common-substrate discrimination in the 327-paper transport comparison.

## 5 LLM use, compute routing, and boundaries

AIO can use local/open-weight or hosted language models for public-literature extraction, subject to license and data-policy review. Clinical/private data should remain inside the approved computational boundary by default. Literature-to-dataset variable alignment, endpoint/horizon equivalence, eligibility criteria, and hypothesis-cluster interpretation remain auditable scientific decisions. Unmapped signatures are flagged rather than fabricated. LLM-generated prose is not numerical evidence.

**Public repository:** `github.com/ConquerMedical/AIO`. The v0.2 repository exposes a generic client, result contract, synthetic call example, frozen audit fixture, and interface tests. Literature search, extraction prompts, clustering, signature construction, model zoo, matched-reproduction, statistical adjudication, and NPIS internals remain protected.

## 6 Installation, verification, and result contract

Python public call

```
git clone https://github.com/ConquerMedical/AIO.git
cd AIO
python -m venv .venv
source .venv/bin/activate
pip install -e ".[test]"
pytest -q
python examples/generic_call.py
python examples/frozen_validation.py
```

The public tests verify that the generic client accepts caller-defined analyses and that the original `aio\_engine/` and vendored `npis/` engine trees are absent.

| Result field | Meaning |
|---|---|
| `analysis\_id` | Identifier for the reconstructed literature/data analysis. |
| `state` | Current scientific-evidence state. |
| `evidence` | Structured quantitative or provenance evidence. |
| `artifacts` | Reports, traces, tables, or other returned analysis artifacts. |
| `tested\_domain` | Declared domain/question and validation scope. |
| `summary` | Concise interpretation of the matched reconstruction. |

**Table 1:** Public AIO result contract.

## 7 Reconstruction discipline

The deterministic/LLM split is a core scientific boundary. The language model may structure text, but numerical comparison, model fitting, confidence intervals, matched protocol construction, and information characterization must be executed by deterministic analysis code. Each executable signature should retain provenance back to its paper-linked claim, mapped variables, endpoint/horizon, inclusion criteria, and validation design. If a required variable or endpoint cannot be mapped, the signature is marked unmapped rather than completed by plausible language.

## 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 and use the public research architecture. Proprietary operational implementations, calibration internals, candidate-generation procedures, search/ranking policies, deployment orchestration, and other system components not necessary to evaluate the public scientific claims remain outside the scope of disclosure. 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. *TopolAI: Information-First 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. *AI on Alzheimer Disease: Common-Substrate Reconstruction and Adjudication*. 2026.

[6] Cover, T. M. and Hart, P. E. Nearest neighbor pattern classification. *IEEE Transactions on Information Theory* 13(1), 21–27 (1967).

[7] U.S. Department of Health and Human Services. Guidance Regarding Methods for De-identification of Protected Health Information in Accordance with the HIPAA Privacy Rule.

[8] Kaggle. Notebooks documentation.

[9] Hugging Face. Hub repositories, licenses, Paper Pages, Collections, and Spaces documentation.

[10] Modal. Documentation for serverless compute, GPU workloads, and secrets.