Acquire: Inverse Measurement Design for Scientific AI

Tested acquisition under task and burden constraints

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

Technical Report v1.0 – 4 September 2026

Abstract

Acquire is the inverse-measurement layer of CMH Scientific. It asks which tested additional observation or protocol is justified after a task-relevant information deficit has been diagnosed. The revised public release separates a generic callable client from frozen validation ledgers and withholds candidate-generation, ranking, frontier-optimization, selective-acquisition, and controller internals. Validation spans PRO-ACT ALS, calcium imaging, a cross-replicate perturbational acquisition study, longitudinal COPD, and AML streaming. In PRO-ACT, collecting more of the original ALSFRS substrate yielded a held-out AUC gain of +0.0020, whereas laboratory measurements yielded +0.0131 with a 0.0087 branching reduction. Calcium imaging provided 9/9 frozen tolerance transfers, and the perturbational study reproduced the same acquisition order across held-out replicates while exhaustive ranks of 10/24 and 12/24 explicitly ruled out a global-optimum claim.

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

Conventional learning fixes the acquisition and optimizes the observer. Acquire reverses the question: once the current information state has been characterized, which *tested and admissible* additional observation should be acquired, where should it be acquired, and when should acquisition stop? Residual ambiguity alone is not evidence for a particular new modality. Candidate-specific recommendations require candidate-specific evidence.

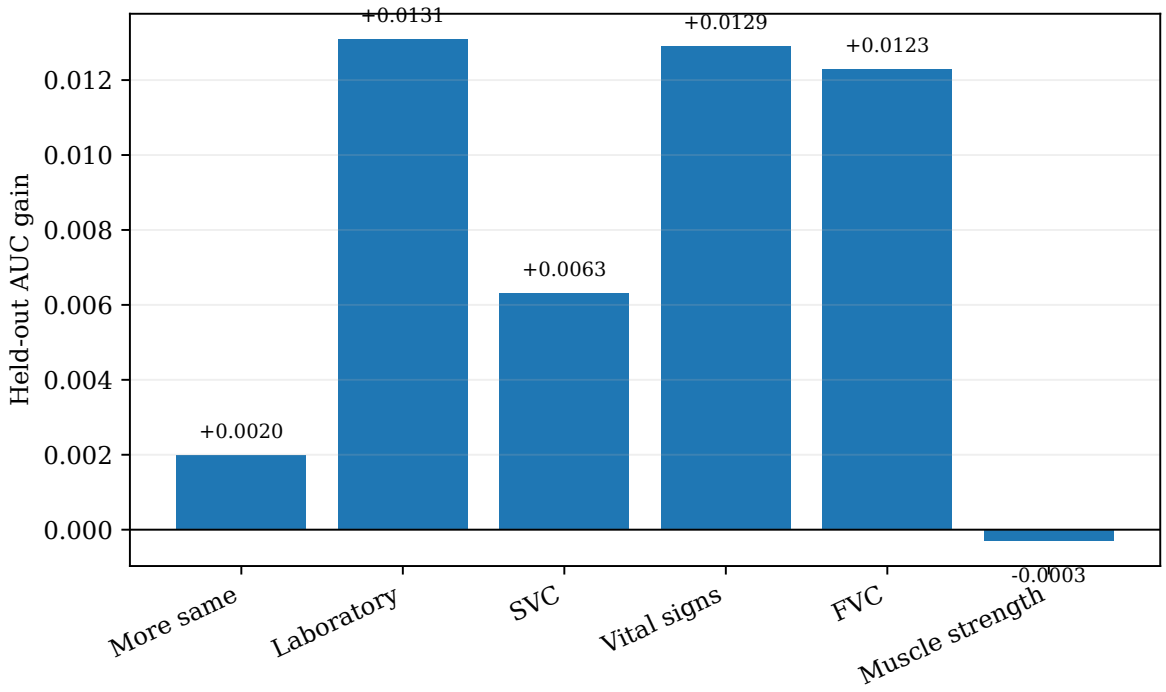

Figure 1: Held-out PRO-ACT AUC gains for more of the same substrate versus tested alternative measurement families.

2 Two-layer public release

The public `acquire/` client accepts caller-defined candidate measurements/protocols, task criteria, cost or burden metadata, current evidence, and constraints. The protected engine performs candidate comparison and returns a standardized `AcquireResult`. The separate `validation/frozen\_acquire\_ledger.json` preserves released validation values without masquerading as a new computation.

3 Step-by-step use

1. **Start from a diagnosed information state.** Supply the task and current evidence, for example an NPIS or OptiCeil result.
2. **Define the admissible candidate set.** Each candidate must correspond to an observation, protocol, sensor, assay, temporal window, or modality for which evidence can actually be generated.
3. **Declare burden and constraints.** Include cost, time, exposure, storage, patient burden, or any other domain-relevant constraint.
4. **Request a design decision.** The authorized engine can compare candidates, estimate conditional value, and determine whether acquisition should be universal, selective, changed, or stopped.
5. **Reassess after new information.** The result is not a one-time prescription. New measurements change the information state and should trigger a new assessment.

Python public call

```python
from acquire import AcquireClient

client = AcquireClient(transport=my_authorized_transport)

result = client.design(
    task={"id": "my_measurement_design",
          "objective": "improve_task_information"},
    candidates=my_candidate_measurements,
    data=my_data,
    base_evidence=my_current_information_state,
    constraints={"budget": my_budget},
    options={"allow_selective_acquisition": True},
)

print(result.state)
print(result.selected)
print(result.frontier_summary)
```

Example prompt

Compare these tested candidate measurements against collecting more of the current substrate. Include acquisition burden, held-out information gain, and task gain. State whether the added measurement should be universal or selective, and do not recommend an untested biomarker or modality.

Structured result

RESULT: CHANGE_MEASUREMENT when a tested alternative dominates collecting more of the current substrate on the declared held-out criteria.
SUPPORTED: tested candidate-family ranking, tested frontier selection, leakage-controlled selective acquisition.
NOT SUPPORTED: claims about untested modalities or global optimality outside the declared candidate set.
NEXT ACTION: acquire the best-supported admissible candidate under the declared burden; otherwise remain unresolved or test another measurement dimension.

4 Validation

4.1 PRO-ACT ALS: measurement-family value

The frozen ledger contains 8,260 participants and 70,746 real longitudinal transitions, split into 4,956 discovery and 3,304 held-out participants. Increasing the original ALSFRS discovery substrate from half to all discovery subjects increased held-out AUC from 0.5975 to 0.5995 (+0.0020). Alternative measurement-family gains were:

Candidate	Branching reduction	Held-out AUC gain
Laboratory measurements	0.0087	+0.0131
SVC	0.0055	+0.0063
Vital signs	0.0055	+0.0129
FVC	0.0041	+0.0123
Muscle strength	0.0009	-0.0003
More of same ALSFRS	–	+0.0020

Table 1: Frozen held-out PRO-ACT comparison.

4.2 Calcium imaging: protocol transfer and over-acquisition

Across three optical regimes, all nine discovery-selected 90/95/99% temporal prescriptions passed the frozen held-out tolerance rule; eight of nine strictly met both nominal held-out quantities. At the 99% target, the frozen protocols were 67 ms / 2 volumes

for GCaMP8m, 250 ms / 8 volumes for nonspinal GCaMP6s, and 667 ms / 20 volumes for spinal GCaMP6s. Acquisition beyond task-saturating regions reduced balanced accuracy by 0.065, 0.031, and 0.024, respectively.

4.3 Cross-replicate, selective-acquisition, and streaming controls

In GSE254742 the response-informed order 75→10→25→50 transferred identically in both replicate directions and outperformed a spacing comparator. Exhaustive ranks were 10/24 and 12/24; the result is therefore reproducible but not globally optimal. COPD selective-acquisition validation showed conditional value of history of 0.710, 0.856, and 0.969 at 30%, 40%, and 50% budgets, respectively (Figure 2). AML streaming stopped intensive acquisition after 41.9% of the stream while retaining 10.6% of the response field, then returned CHANGE_MEASUREMENT; a later challenge triggered REACQUIRE after 1,248 observations.

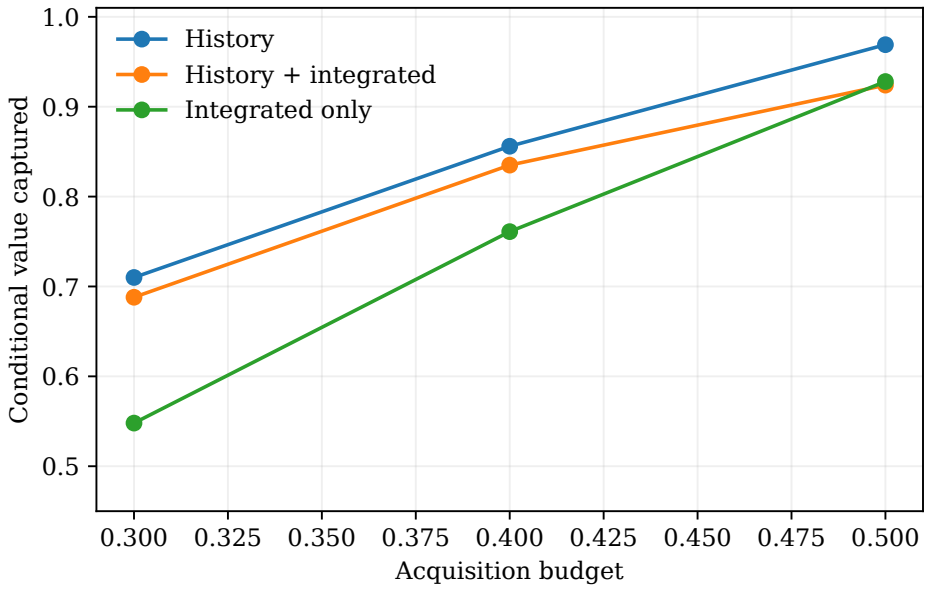

Figure 2: Frozen COPD selective-acquisition value across acquisition budgets.

5 Boundaries and reproducibility

Acquire only ranks or selects among declared, tested/admissible candidates. Retrospective held-out transfer is not prospective physical or clinical actuation. It does not infer that a particular untested sensor, assay, modality, dose, or protocol will repair an information deficit.

Public repository: `github.com/ConquerMedical/Acquire`. The repository exposes the callable contract and frozen validation ledger; the production candidate-generation, frontier, ranking, selective-acquisition, controller, and actuation implementations are not distributed.

6 Installation, interface verification, and result contract

Python public call

```
git clone https://github.com/ConquerMedical/Acquire.git
cd Acquire
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 test suite checks the callable contract, validates the frozen ledger, and asserts that the original `acquire\_ai/` engine tree is absent.

Result field	Meaning
`task\_id`	Declared inverse-measurement task.
`state`	Acquisition/controller state returned by the authorized engine.
`selected`	Tested/admissible candidate objects selected under the declared constraints.
`evidence`	Held-out/task-specific evidence supporting the selection.
`frontier\_summary`	Summary of the tested candidate frontier; not a claim of global optimality.
`tested\_domain`	Candidate set, validation design, and scope to which the decision applies.

Table 2: Public Acquire result contract.

7 Additional frozen validation detail

Validation	Frozen result
Cheng replicate 1→2	Normalized acquisition AUC 0.726 versus spacing 0.562; five-of-six fidelity 0.98; exhaustive rank 10/24.
Cheng replicate 2→1	Normalized acquisition AUC 0.705 versus spacing 0.548; five-of-six fidelity 0.91; exhaustive rank 12/24.
COPD layered resolution ($K = 30$)	Current 0.250; history 0.278; integrated 0.340; both 0.466.
COPD layered branch probability ($K = 30$)	Current 0.209; history 0.185; integrated 0.131; both 0.129.
COPD task-specific control	$n = 785$; clinical resolution 0.0743; expression resolution 0.1083; Gaussian-noise control 0.0441; shuffled-expression control 0.0478.

Table 3: Additional values preserved in the frozen Acquire validation ledger.

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.