# From Prediction to Scientific and Clinical Awareness

## The CMH Scientific instrument landscape

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

Technical Report v1.0 – 10 September 2026

**Abstract**

Biomedical AI is commonly organized around prediction: map an observed state to a label, risk, response, or action. CMH Scientific instead develops a family of computational instruments that expose different forms of operational awareness around prediction. Here "awareness" does not mean consciousness or sentience. It means explicit machine-readable state about the limits, context, history, measurement, intervention, and falsification conditions that determine whether a scientific or clinical output is justified. This report maps the CMH instrument landscape across evidence awareness (AIO), information-sufficiency awareness (NPIS), recoverability awareness (OptiCeil), structural/history awareness (TotemicAI), acquisition awareness (Acquire), supervisory action awareness (TopolAI), patient-specific intervention awareness (ABP), procedural and physiologic awareness (Deferra), trial-falsification awareness (CTF), and governance awareness (CMH Guard). The instruments are separable; none is a mandatory serial stage for every task.

**Awareness is an engineering property here.** A system is "aware" only in the limited operational sense that it maintains and reports a declared state – for example, that the current measurement cannot support the requested distinction, that the response relation changed after intervention, or that a purported responder subgroup also appears under placebo.

## 1 Why prediction alone is an incomplete product category

A high-performing predictor does not by itself answer several questions that matter before, around, and after inference: Does the measurement distinguish the target? How much information is recoverable from the representation? Which past observations must be retained? What should be measured next? Is an intervention direction supported for this individual? Has an intervention changed the process being modeled? Should a clinical AI speak now? Does an apparent trial subgroup survive a placebo falsification test?

CMH Scientific treats those questions as distinct computational jobs. The result is better understood as a landscape of instruments than as a model zoo.

## 2 The awareness landscape

| Instrument | Operational awareness | Question made explicit |
|---|---|---|
| AIO | Evidence / hypothesis awareness | What claims exist, which are comparable, and what survives reconstruction on common empirical ground? |
| NPIS | Information-sufficiency awareness | Do the recorded variables distinguish the scientific states, directions, or futures being claimed? |
| OptiCeil | Recoverability awareness | How much of the declared task appears recoverable from this representation, and is the limitation measurement or observer? |
| TotemicAI | Structural/history awareness | Which observations preserve the nonredundant response geometry, and when has new data broken that memory? |
| Acquire | Acquisition awareness | What should be measured next, in whom, and when does more measurement stop earning its cost? |
| TopolAI | Supervisory/action awareness | Given the evidence state, should the system use, continue, selectively acquire, change measurement, add modality, reacquire, abstain, or remain unresolved? |
| ABP | Individual intervention awareness | For a deeply observed focal biological system, which intervention directions and levels are supported by that system's own history? |
| Deferra | Procedural / physiologic awareness | Is the current observation adequate, has the course or response relation changed, and should AI assist, withhold, reacquire, or remain silent? |
| CTF | Trial-falsification awareness | Does an apparent responder subgroup survive the identical discovery procedure when applied to placebo? |
| CMH Guard | Governance awareness | Is the proposed data/code/model/compute/export route permitted under the declared policy boundary? |

**Table 1:** CMH Scientific separates operational awareness into explicit, auditable state variables and controller questions.

## 3 Layers rather than a mandatory pipeline

The instruments can be composed, but composition is task-dependent. A literature-reconstruction project may use AIO and NPIS without any clinical controller. A microscopy design project may use OptiCeil and Acquire. A streaming perturbation experiment may use TotemicAI with TopolAI. A patient-specific intervention design may use ABP without AIO. A procedural deployment may place Deferra around computer vision, physiology, a predictive model, or a device controller.

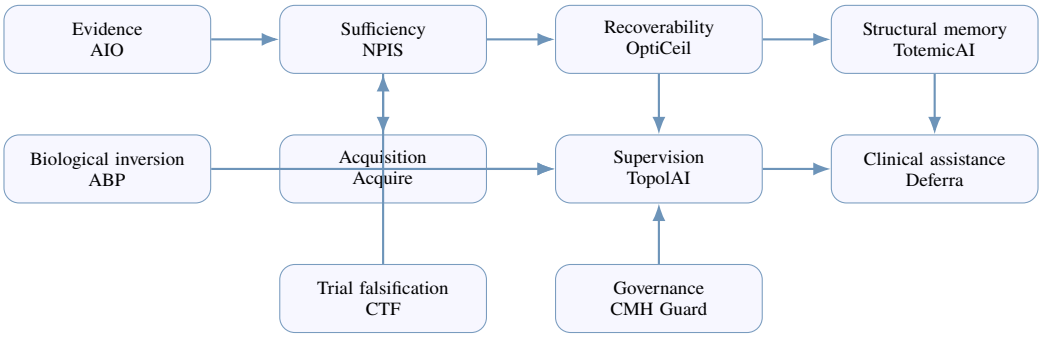


**Figure 1:** One possible composition of the instrument landscape. Arrows indicate useful handoffs, not a required universal sequence.

## 4 Awareness before prediction: NPIS and OptiCeil

NPIS asks whether the observed substrate contains the distinction required by the claim. It characterizes local branching, separation among compatible futures, temporal refinement, departure/change, and residual predictive availability without treating a single model score as the definition of information.

OptiCeil asks a complementary question: within a declared representation and validation unit, how much performance appears recoverable, how much has already been demonstrated, and does the remaining gap point toward the observer or toward

acquisition? Its empirical envelope is representation- and instrument-conditioned rather than a universal biological ceiling.

Together these instruments turn "the model is weak" into more specific states: useful information remains but the observer has not extracted it; the current measurement is restrictive; the conclusion is unresolved; or a different acquisition should be tested.

## 5 Awareness of history: TotemicAI

TotemicAI is a fitted structural-memory model. In TotemicAI-Specified, the landmark policy is fixed and the memory is compiled from data. In TotemicAI-Learned, policy coefficients are fitted on training data and frozen before memory compilation and held-out evaluation. Current AML validation belongs to the specified mode.

The awareness element is not "remember everything." It is knowing which observations function as structural landmarks and whether the current state remains coherent with that stored structure. A challenge state can therefore trigger reacquisition or recompilation rather than silent drift.

## 6 Awareness of what to measure: Acquire and TopolAI

Acquire converts an unresolved information state into an acquisition decision. It can compare more of the same data with a different measurement family, a different protocol, another time window, another modality, or selective measurement of only those subjects for whom additional information is likely to matter.

TopolAI supervises the resulting state. It returns actions rather than merely another score: use, continue, continue selectively, change measurement, add modality, reacquire, abstain, or unresolved. This is supervisory awareness of what the scientific system should do next.

## 7 Awareness of intervention: ABP

ABP changes the modeling unit. When one biological system is observed deeply enough, the focal unit itself can become the model object. The user declares a desired state and interrogates the fitted response object in reverse to identify intervention directions represented and supported by that individual's own history. This is intervention awareness in the limited engineering sense: the system keeps track of which candidate redirections are supported for the focal biological primitive rather than substituting a population cohort for the individual's identity.

ABP can be expressed in clinical-intervention, therapeutic-development, and experimental-perturbation settings. Current evidence supports mechanics within observed intervention spaces; it does not establish de novo molecule generation, causal efficacy, or clinical treatment recommendation.

## 8 Awareness during action: Deferra

Deferra operates after the system enters a physical or physiologic feedback loop. It maintains support, procedure/physiology context, expected trajectory, departure, and response-relation state. Crucially, a poor observation is not treated as the same thing as a changed clinical course. Deferra can withdraw guidance and reacquire the view or measurement; remain silent during coherent expert execution; or activate assistance when a supported departure or changed response relation warrants it.

This creates a different product category from always-on assistive AI. Scene recognition can tell the system what appears in the field; prediction can estimate what comes next; Deferra supervises whether the current evidence gives the AI any business speaking or acting.

## 9 Awareness of falsification: CTF and AIO

AIO builds an evidence landscape: it reconstructs published claims on common empirical ground and asks which hypotheses the available substrate can distinguish. CTF applies a narrower adversarial logic to trials: an apparent responder-discovery procedure is repeated on placebo. If the same "responder" signal appears where active treatment is absent, the subgroup is

evidence of the discovery procedure, prognostic heterogeneity, regression, or another non-treatment mechanism rather than a clean treatment-response finding.

These are awareness functions because they make the evidentiary counterfactual explicit before a downstream model, trial, or product is built around a result.

## 10 Competitive landscape: what CMH instruments add

The table below describes functional distinctions rather than claims that competing products lack every listed capability.

| Common product category | Typical primary output | CMH layer added |
|---|---|---|
| Predictive model / foundation model | Score, class, generated response, representation | NPIS/OptiCeil expose whether the substrate supports the task and whether failure is measurement- or observer-limited. |
| Compression / retrieval memory | Representative or retrieval-efficient subset | TotemicAI preserves declared structural landmarks and challenge state rather than population representativeness. |
| Experimental-design optimizer | Candidate experiment or acquisition | Acquire ties the next measurement to an explicit diagnosed information deficit and can compare more-same versus different information. |
| Digital twin / response model | Simulated future under interventions | ABP makes focal-unit identity and same-unit support explicit and permits inverse interrogation of the observed individual's response object. |
| Surgical / bedside AI | Annotation, alarm, recommendation, robotic action | Deferra adds observation support, trajectory departure, changed-response state, selective silence, reacquisition, and assistance routing. |
| Trial subgroup analytics | Candidate responder subgroup | CTF makes placebo falsification of the discovery procedure a first-class test. |
| AI governance / access control | Policy, permissions, logs | CMH Guard routes the scientific workflow through explicit data/model/compute/export boundaries. |

**Table 2:** Functional positioning of CMH instruments. The comparison is architectural, not a claim of exhaustive market uniqueness.

## 11 Why the landscape matters

The suite is valuable precisely because these states are not collapsed. Structural stability does not prove information sufficiency. A high recoverability envelope does not mean the current model has reached it. A plausible intervention surface does not prove causality. An adequate operative view does not mean the course is coherent. A responder subgroup does not prove treatment effect. A technically possible compute route does not mean the data are permitted to leave an institution.

The common design principle is therefore not "one model to do everything." It is **make the missing state explicit, expose it to the user, and route the next action according to what is actually supported**.

## 12 Boundaries

The term awareness is used only in this operational engineering sense. No claim of consciousness, sentience, general agency, or human-equivalent understanding is made. The instruments have heterogeneous validation and should not be collapsed into a single performance score. Some modules are retrospective research tools; Deferra and ABP include clinical or therapeutic embodiments that require prospective validation before clinical use; learned TotemicAI requires separate out-of-sample validation from the currently validated specified mode.

## Intellectual property, competing interests, funding, and use

Relevant methods and computational systems are the subject of previously filed U.S. provisional patent applications. Public reports disclose the scientific objects, interfaces, evidence, and claim boundaries needed to evaluate the released work 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. *CMH Scientific: R&D Laboratory and Scientific AI Instrument Suite*. Conquer Medical Health, 2026.

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

[3] Ewing, M. A. *Deferra: Trajectory-Conditioned Selective AI Assistance*. Conquer Medical Health, 2026.

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