TECHNICAL REPORT TR-2026-09

Knowing a Fact vs. Being Able to Act

Separating declarative knowledge gaps from procedural competence gaps in a taught reasoning system

Stefan Ragland

Dominion Labs Research & Development

Dominion Labs

research@dmnlabs.org

January 13, 2026

Abstract

A system that learns by being taught accumulates two distinct kinds of mastery over any subject: *declarative knowledge* — the facts it holds — and *procedural competence* — the operations it can actually carry out. A system that conflates the two misreads its own state, mistaking "I have not yet been told this fact" for "I am not competent in this area," or the reverse. We formalize the two as independent, separately measured axes, state five falsifiable hypotheses about whether a taught reasoning system can discriminate between them and whether the axes stay orthogonal, and test them on a running symbolic system. Teaching three connected facts about a subject raised its declarative coverage (0.229 → 0.293) while leaving procedural competence exactly unchanged (0.500 → 0.500). Three hypotheses held immediately; a fourth — detecting a *specific* missing fact inside an otherwise well-developed subject — initially failed because the system's register of acknowledged ignorance existed but was never populated. After adding the missing detector, all sixteen checks passed, with negative controls confirming the detector neither fabricates gaps where none exist nor disturbs the competence estimate. We contrast this behavior with an autoregressive language model, which represents no such boundary and therefore cannot abstain on the basis of it.

1 Two kinds of not-knowing

Over any subject, a learning system can be strong along two different dimensions that are easy to confuse. *Declarative knowledge* is how much it knows — the facts it has been told and connected together. *Procedural competence* is what it can do — the operations it can actually execute to reach a goal. The distinction is old in the study of mind: knowing *that* versus knowing *how* [1], and, in cognitive architecture, declarative versus procedural memory [2].

> *Declarative knowledge and procedural competence are separable, independently measurable properties of a taught system. Representing them as orthogonal axes lets the system diagnose which kind of gap it faces — a missing fact to be acquired, or a missing operation to be practiced — and abstain honestly on that basis.*

2 Why conflating them fails

A missing fact and a missing operation call for opposite responses. A *knowledge gap* is closed by acquisition — being taught the fact, or looking it up. A *competence gap* is closed by practice or derivation — searching for an operation that produces the goal. Routing one to the other wastes effort and corrupts the system's estimate of what it can do: treat "I was never told this one fact" as "I am not competent in this whole area" and the system needlessly discounts real ability; treat "I cannot perform this operation" as "I just lack a fact" and it will keep asking to be told things no fact can supply.

An autoregressive language model has a single channel — next-token likelihood — and no internal variable separating "I hold little information about this" from "I cannot perform this operation." Both surface identically, as a distribution over tokens. Part of what is called hallucination is precisely the absence of this boundary: without a representation of *this subject's coverage* as distinct from *this operation's feasibility*, the model cannot know it is missing something, so it answers fluently anyway. Selective-prediction methods add an abstention head over likelihood [3]; they do not add the two-axis boundary studied here.

3 Method: two axes, two producers

Declarative coverage is a maturity measure over the graph of taught concepts, rising as connected facts are taught. It is a normalized score in [0,1] combining how many concepts the subject contains (size), how densely they interconnect (density), and how many distinct kinds of relation link them (variety), weighted 0.4 / 0.3 / 0.3. **Procedural competence** is a probabilistic belief that the system has learned the subject's operations; it moves only on the outcomes of attempting operations, never on being taught facts. The load-bearing property is that *neither producer writes the other's variable*.

Deciding which gap is faced. When a goal cannot be achieved, a structural diagnosis walks the chain from symbol to action and returns exactly one typed deficit. Two are central: an *unrepresented-predicate gap* (the goal concept appears nowhere in the subject's vocabulary — escalate to acquire a new concept) and a *missing-operation gap* (the concept is represented, but nothing the system knows produces it — explore for an operation). Crucially, diagnosing a gap does not by itself change the competence belief: a diagnosis is not evidence of incompetence.

The declarative twin. The system carries a first-class representation of *acknowledged ignorance*: a subject-scoped question it knows it cannot answer, with an information-value and a resolution strategy. Our core addition is a detector that, for a missing relation on a *known, in-subject* concept, registers exactly such an entry — while leaving the competence belief untouched — and refuses to register anything when the subject is out of scope or the relation is already present. A registered gap is surfaced as an exploration target and converted into an *acquisition* goal (be taught, or look up), not an experiment to build an operation.

4 Experiment

To measure the axes cleanly, a mature subject was built by teaching a small, connected, deliberately fictional taxonomy through the ordinary teaching interface (fictional terms keep the test subject isolated and deterministic). Six facts built the subject; three more connected facts were taught for the orthogonality test. Five hypotheses were then evaluated live: discrimination (H1), routing (H2), axis orthogonality (H3), declarative-gap detection with negative controls (H4), and gap pursuit (H5).

5 Results

5.1 The axes are orthogonal

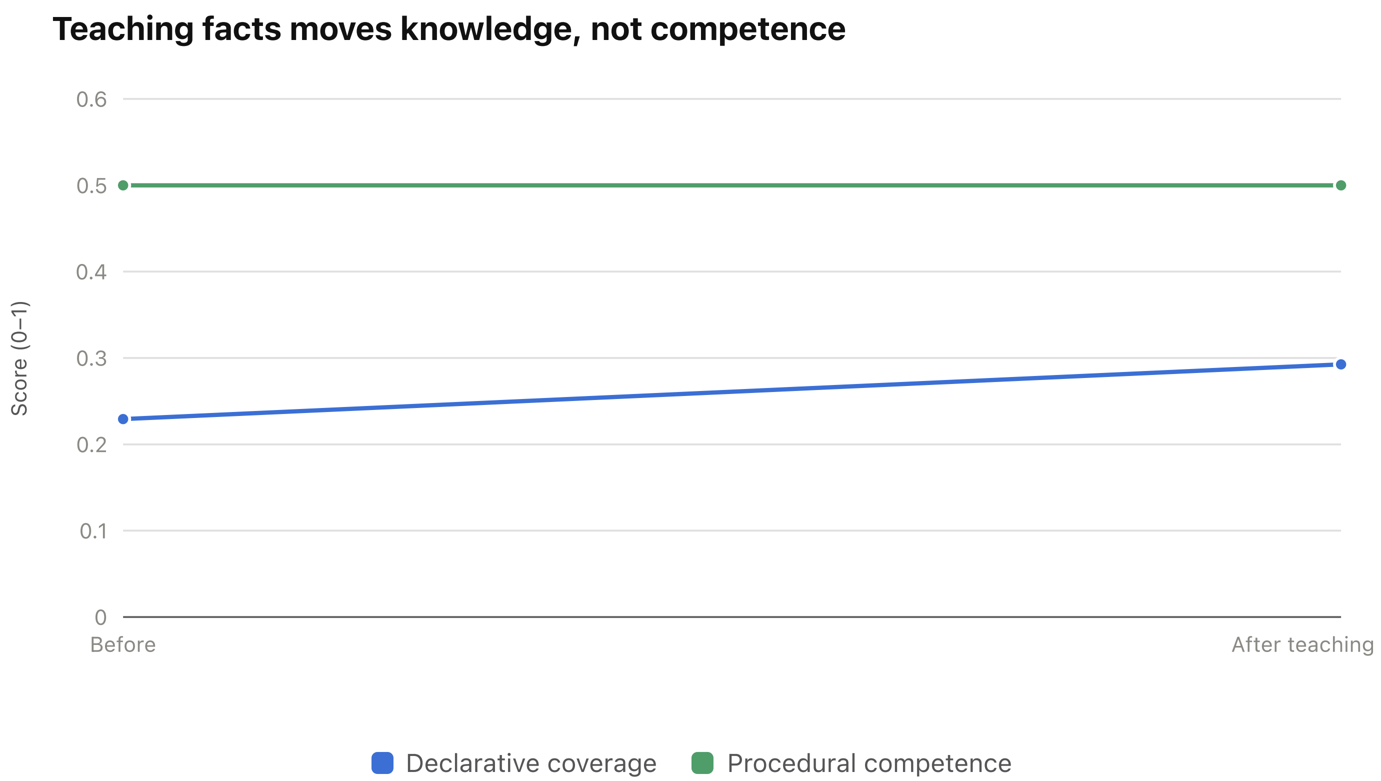

Figure 1. After teaching three additional connected facts, declarative coverage rose from 0.229 to 0.293, while procedural competence held at exactly 0.500 (change below 10^{-9}). One axis moves; the other does not — the orthogonality the design requires.

5.2 Discrimination, routing, detection, pursuit

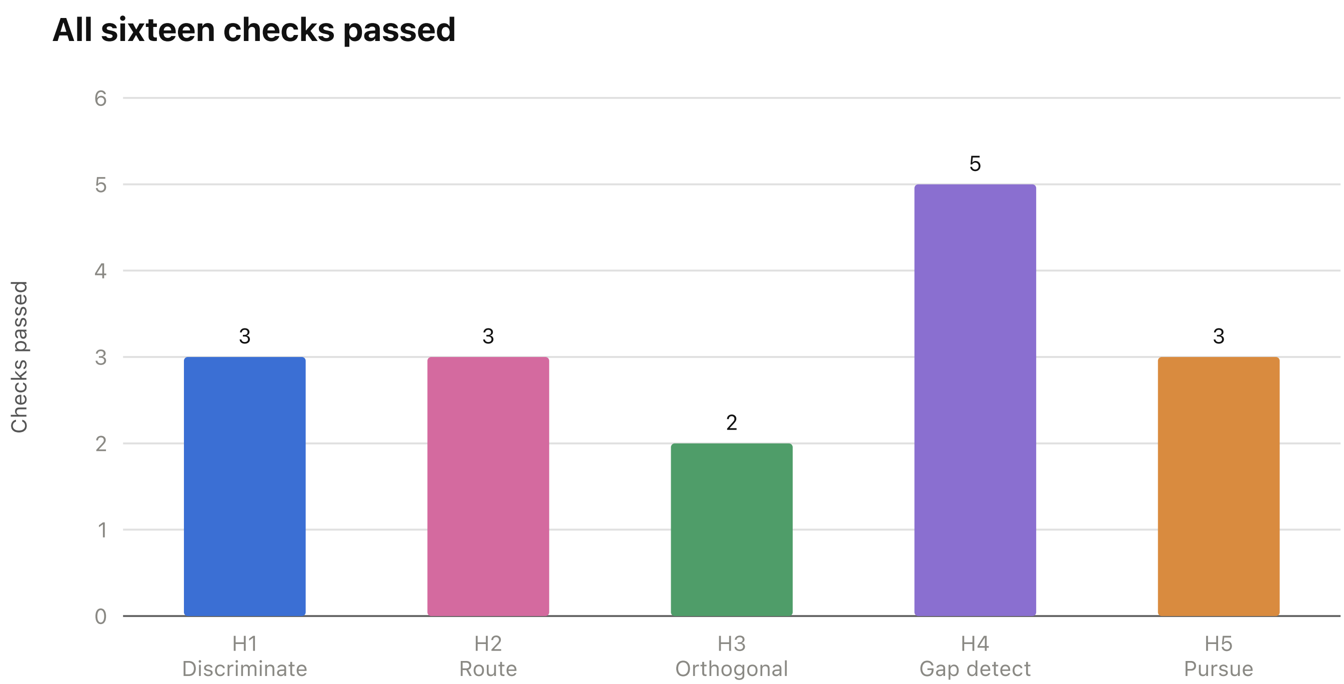

Figure 2. Checks passed by hypothesis: discrimination 3/3, routing 3/3, orthogonality 2/2, declarative-gap detection 5/5, gap pursuit 3/3 — 16/16 in total. Discrimination returns the correct gap type (0 actionable operations for a missing-operation gap; predicate-not-in-vocabulary for an unrepresented-predicate gap); routing sends the two to acquire-a-concept versus explore-for-an-operation.

5.3 A gap it detects, without fabricating gaps

Detecting a specific missing fact inside a mature subject initially failed: the register of acknowledged ignorance existed but nothing wrote to it. That baseline failure is itself a reported finding. With the detector added, a missing relation on a known concept registered exactly one acknowledged-unknown (count $0 \to 1$, information-value 0.6), which became one exploration target (entropy 0.94) routed to acquisition — while competence stayed at 0.500. Two negative controls confirm specificity.

Situation	Gap registered?	Competence
In-subject, relation not yet held	yes ($0 \to 1$)	0.500 (unchanged)
Out-of-scope subject	no	0.500 (unchanged)
Relation already present	no	0.500 (unchanged)

Table 1. Gap detection and its negative controls. The detector fires only for an in-subject, not-yet-held relation, and never disturbs the competence estimate.

6 Limitations and conclusion

The coverage measure is a specific normalized combination of size, density, and variety; other reasonable measures exist, and the absolute numbers are not the point — the separation of the two axes is. The subject is a small fictional taxonomy chosen for determinism; scaling the discrimination to large, real, partially-taught subjects is future work, as is a competence axis defined over many operation types rather than one. The result we stand on is qualitative and robust: a taught system can hold two independent measures of mastery, decide which kind of gap it faces, and abstain honestly — a boundary a single-channel model does not represent.

References

[1] G. Ryle. *The Concept of Mind*. Hutchinson, 1949.

[2] J. R. Anderson. *The Architecture of Cognition*. Harvard University Press, 1983. (Declarative vs. procedural memory.)

[3] R. El-Yaniv, Y. Wiener. On the foundations of noise-free selective classification. *JMLR*, 11, 2010.

[4] S. Muggleton. Inductive logic programming. *New Generation Computing*, 8(4), 1991.

[5] C. Guo, G. Pleiss, Y. Sun, K. Q. Weinberger. On calibration of modern neural networks. *ICML*, 2017.