Open Access

Milestone 009 — From hypothesis sprawl to a theorem ledger

Tyler J.H.G.

Abstract

How can checked mathematics constrain the next architecture experiment?

Research period: September 3, 2026 onward
Historical anchors: #196, #198–#200, later #257

The hypersphere work opened a much larger design space. Should hierarchy use hyperbolic geometry? Could a concept be a subspace rather than a vector? Should relationships be transformations? Would product spaces, graphs, or sheaf-like transport preserve distinctions that an ordinary vector collapses?

I could turn each question into an architecture and start training. Before doing that, though, I wanted to know whether the proposed construction actually had the property motivating it. A failed training run would be a poor way to learn that the mathematical premise was impossible.

The theorem ledger records that earlier reasoning: what can be established about the claim, what it assumes, and which empirical question remains once those parts are settled.

1. Four kinds of evidence

I separated the claims into four categories:

- ✓ **Formal checkpoint** — a mathematical statement established under explicit assumptions.
- ● **Observation** — a measurement under a specified empirical protocol.
- ◇ **Assumption** — a premise the argument uses without establishing it.
- ? **Hypothesis** — an explanation or prediction still open to testing.

Consider three statements from the sphere investigation:

```
normalization removes infinitesimal radial perturbations → formal
checkpoint
```

```
the unit model had lower loss at the #164 endpoint → observation
```

```
radial variation is nuisance information → hypothesis
```

The first describes the map, the second describes a measured model, and the third proposes a task-dependent explanation. Connecting them requires an argument that

Published Sep 29, 2026

establishes the missing premises. Recording their categories makes those premises visible.

2. The ledger

The ledger connects claims to their assumptions, evidence, consequences, and dependent experiments. If a premise changes, I can see which later claims need to be reconsidered.

A proof can sometimes eliminate a proposed experiment. A counterexample can rule out a universal architecture claim. In other cases, the mathematics reduces a broad architecture question to one remaining empirical premise, giving the experiment a more specific target.

The public Theorem Library contains the reusable formal results, including machine-checked Lean proofs where appropriate. The ledger tracks how those results affect the investigation and its next decisions.

> **Sticky note — counterexample:** one valid case that violates a universal claim disproves that claim as stated. Finding it before implementation can remove an entire branch of architecture search.

= Try a small example

A small dictionary labels three claims by the evidence supporting them.

```
claims = {
    'normalization radial derivative is zero': 'formal',
    'unit model had lower NLL at #164 endpoint': 'observation',
    'radial direction is nuisance': 'hypothesis',
}
for claim, kind in claims.items(): print(f'{kind:11s} | {claim}')
```

Saved output:

```
formal      | normalization radial derivative is zero
observation | unit model had lower NLL at #164 endpoint
hypothesis  | radial direction is nuisance
```

3. Research context

The growing design space crossed geometric deep learning, information theory, dynamical systems, representation learning, and formal methods. Many architecture proposals contained subproblems that those fields had already studied.

Reading that work helped separate established mathematics from unresolved behavior in the model. A short derivation or counterexample could also expose an assumption I had omitted. The literature became a way to identify where the experimental question begins, as well as a source of constructions to test.

4. Interpretation

A proof establishes its conclusion under its encoded assumptions. Applying it to a trained model requires evidence that the relevant assumptions describe that model.

The ledger makes this chain explicit. Mathematical results and observations can then constrain the next experiment without silently promoting a proposed explanation into an established result.

5. Sources and chronology

- Milestone notebook 009 — preserved chronological source
- Research chronology
- Publication dispositions and Atlas-to-article map

The matching Atlas theme is **Claim-type distinctions**. The earlier Visual Intuition Atlas notebook preserves the source visualization.

The next chronological article is **Milestone 010 — What state should a reasoner preserve?**.