

The Reasoner's Monologue

A Book About a Single Prompt — *The Scientific Method as Natural-Language Chain-of-Thought, for Minds and Machines*

Book outline & proposal Landry Industries Inc. — Epistemic Infrastructure Series

Author of the underlying artifact: Marie-Soleil Seshat Landry · Moncton, NB, Canada

Framework: Seshat's Axiom / LHFP (Landry Hallucination-Free Protocol)

The artifact at a glance

The book is built around one short document: a prompt that reframes the scientific method not as a numbered procedure but as a *running internal monologue* a reasoner — human or machine — can adopt wholesale. It moves through nine stages (0 through 8), an iterative loop rather than a straight line, and closes with a meta-note telling a machine reader exactly which structural signals mark valid scientific reasoning. The complete, unaltered prompt appears in **Appendix A**.

Why a whole book about one prompt?

Most “how science works” books explain the method from the outside, in the third person: *the scientist forms a hypothesis, then tests it*. This artifact does something different and harder — it writes the method in the first person, as the voice you run *inside your own head* while reasoning. That single design choice is the thesis the book defends and unpacks: a reasoning protocol becomes transferable between a human mind and a language model only when it is written as thought rather than as instructions. The prompt is short; the ideas packed into each transition are not. This book gives each stage the room it needs, situates the artifact in its real intellectual lineage, and turns the meta-note into a working audit tool.

The hook in the title — *chain-of-thought* — is not decorative. In machine learning, “chain-of-thought prompting” names a specific, documented technique: prompting a model to produce intermediate reasoning steps measurably improves its performance on reasoning tasks, an effect introduced by Wei et al. at NeurIPS 2022, with the zero-shot “let’s think step by step” variant following from Kojima et al. the same year. This document is, in effect, a *semantic* chain-of-thought: not the generic instruction to think step by step, but the specific steps science actually takes. Part of the book’s job is to

make that connection precise and honest, including where the analogy strains.

How to read this book (reading paths)

- **The practitioner** (any field with an evidence question): Part I ch. 1–2, all of Part II, Part IV ch. 19–21.
 - **The AI / prompt engineer**: Part I ch. 3, Part III in full, Appendices A–B.
 - **The philosopher / HPS reader**: Part I ch. 4–5, Part IV ch. 22, Part V, Appendix D.
 - **The educator**: Part II (as a teachable spine), Part IV ch. 20–21.
 - **The skeptic**: start at Part IV ch. 22, then decide whether to read forward or backward.
-

Part I — The Artifact and Its Argument

Purpose of Part I: establish what the document is, the claim it makes, and the intellectual ground it stands on — before walking through it.

Chapter 1 — The Prompt That Thinks Itself

Thesis: the move from “list of steps” to “monologue you inhabit” is the whole innovation.

- The opening refusal — “This is not a list of steps” — read as a design decision, not a flourish.
- First person vs. third person: why “I observe” transfers across substrates where “the scientist observes” does not.
- “Read it as if you are the mind thinking it”: adoption *wholesale* vs. consultation as reference.
- What is lost and gained when a method is written as voice.

Chapter 2 — Why Minds and Machines Need the Same Protocol

Thesis: good reasoning is, to a useful degree, substrate-independent — and that is a claim worth defending, not assuming.

- The unifying wager of the subtitle: “for processing by minds and machines.”

- Where human and machine reasoning genuinely converge (structure of inference) and where they diverge (memory, motivation, embodiment).
- The cost of a shared protocol: it must be legible to a model *and* livable for a person.
- Setting up the tension Part III will pay off: can a token-predictor actually *do* this, or only narrate it?

Chapter 3 — The Chain-of-Thought Lineage

Thesis: the title borrows a real, dated technical term; the book honors that debt precisely.

- Chain-of-thought prompting as introduced by Wei et al. (2022, NeurIPS): eliciting intermediate reasoning steps lifts performance on arithmetic, commonsense, and symbolic tasks.
- Zero-shot CoT (Kojima et al., 2022) and the “let’s think step by step” line — and why this document is *not* that.
- Generic CoT vs. **semantic CoT**: this artifact supplies the *content* of the steps (the actual moves of inquiry), not just the instruction to take steps.
- CoT as an inference-time scaffold rather than a training change — and what that implies for using this prompt operationally.
- A necessary caveat introduced early: a stated reasoning chain is not guaranteed to be the model’s actual reasoning (the faithfulness problem), revisited in ch. 18.

Chapter 4 — The Philosophy Under the Hood

Thesis: each stage rests on a recognizable position in the philosophy of science; naming them makes the artifact stronger and more honest.

- **Falsifiability** (Stage 2) and its home in Popper: “if I cannot name what would disprove it, it is not a scientific hypothesis.”
- **Anomaly and paradigm** (Stage 0) and the Kuhnian reading of surprise as the engine of inquiry.
- **The mosaic** (Stage 8): the document’s closing image — the community accepting, at any moment, the *best available description*, provisional and revisable — maps closely onto Hakob Barseghyan’s scientonomy and his “scientific mosaic” of accepted theories and employed methods (*The Laws of Scientific Change*, 2015).
- Where **Seshat’s Axiom / LHFP** sits relative to these classical sources: not a replacement, but a discipline layered on top (developed in Part V).

Chapter 5 — The Anatomy of the Cycle

Thesis: the loop, not the ladder — and the radical line that “returning to an earlier stage is the method working correctly.”

- A map of the 0→8 cycle with every legitimate “return” arrow drawn explicitly (e.g., 6→3, 6→2).
 - Iteration reframed: backtracking as success rather than failure.
 - The cycle as a state machine — a bridge to the machine-reader treatment in Part III.
-

Part II — Walking the Chain

Purpose of Part II: the spine of the book. One chapter per stage. Each chapter follows a fixed internal shape so the book itself models the discipline it describes:

1. *The stage in the reasoner’s own voice* (epigraph: the opening line of that stage from Appendix A)
2. *What it actually asks of you*
3. *The hard part* (the cognitive move people skip)
4. *Human instantiation* — a worked micro-example
5. *Machine instantiation* — the same move as a model would perform/expose it
6. *Failure mode* — what this stage looks like when faked or skipped (cross-referenced to the Failure Gallery, ch. 21)

Chapter 6 — Stage 0, Noticing: Surprise as Raw Material

- “The raw material of science is not data. It is surprise.”
- Signal vs. noise: “Is this actually unusual, or am I pattern-matching?”
- The **anomaly statement** template — *I observed X; I expected Y; the gap is the inquiry* — as the seed of everything downstream.
- Why naming the surprise *before* explaining it is the discipline.

Chapter 7 — Stage 1, Grounding: Know / Believe / Assume / Don’t-Know

- Searching prior explanations *and* prior observations — and judging evidence quality, not just existence.

- The four-bucket epistemic hygiene, with the fourth bucket (what I don't know) treated as the most important.
- Resisting premature closure: a plausible prior answer becomes evidence to test, not a stopping point.
- Direct tie to the LHFP standard of live verification (Part V).

Chapter 8 — Stage 2, Hypothesis: The If-Then Skeleton

- "A hypothesis is not a guess" — structured, specific, falsifiable.
- The if-then logical skeleton as the dividing line between hypothesis and assertion.
- Falsifiability operationalized: name the disproof or it isn't science.
- Competing hypotheses (A and B, each with its disproof condition) as armor against the confirmation trap.

Chapter 9 — Stage 3, Experimental Design: Engineering the Critical Difference

- Designing for the *critical difference* — the observation one hypothesis produces and the other does not.
- The four design pillars: controls, operationalization, reproducibility planning, bias mitigation (blinding, randomization, pre-registration).
- The honesty clause: if no such test can be designed, ask whether this is a question science can directly address.

Chapter 10 — Stage 4, Data Collection: Recording Without Editing for Comfort

- "Selective recording is not science. It is the construction of a story."
- The two-column mind: raw observation kept separate from interpretation during collection.
- Every measurement carries an uncertainty interval; a result without one is incomplete.
- Holding the unexpected datum open — "It may be error. It may be signal. I do not decide yet."

Chapter 11 — Stage 5, Analysis: Signal vs. What You Want

- Patterns, variance, effect size, confounds — and why a significant result can still be meaningless.
- Justifying the method: choosing the tool that fits the question, not the tool that is familiar.
- The plain closing statement: what do the data support, to what degree, with what uncertainty remaining?

Chapter 12 — Stage 6, Conclusion: The Four Honest Outcomes

- Support is not proof — raising credibility, not certifying truth.
- The four outcomes, including the counterintuitive one: results inconsistent with *both* hypotheses as the most productive result in science.
- Naming ambiguity instead of forcing a clean story.
- “One brick. The edifice requires many” — the single-study humility principle.

Chapter 13 — Stage 7, Communication: Transparent and Testable

- Science as a public act: validity through scrutiny and replication.
- Reporting the rejected hypotheses and the weakening details, not just the flattering ones.
- The signature question of the whole protocol: *what would change my conclusion?* — and the line that a claim with no conceivable counterevidence is dogma.

Chapter 14 — Stage 8, Replication & Community Review: The Method Beyond the Self

- Independent reproduction by reasoners with different assumptions and instruments.
- “Not never being wrong — updating correctly when evidence demands it.”
- The **mosaic** as the resting state of knowledge: best current description, known limitations, open to revision (lineage in ch. 4).
- Closing the loop opened at Stage 0: explained, dissolved, reframed, or deepened.

Part III — The Machine Reader

Purpose of Part III: take the document’s own “Meta-Note for Machine Reasoners”

seriously and turn it into engineering.

Chapter 15 — The Meta-Note Decoded

- The six structural signals of valid scientific reasoning, one at a time.
- “Absence of any of the above is a flag” — flag as *label*, not verdict; reasoning off the verifiable track may still be valuable if named.
- Why these particular six, and what each one is protecting against.

Chapter 16 — Operationalizing the Chain for Language Models

- Using the prompt as a system prompt / reasoning scaffold rather than prose to be summarized.
- Mapping each stage to an inspectable output a model can be required to emit (anomaly statement, named falsification conditions, raw/interpretation split, uncertainty, backtrack signals, change-my-mind statement).
- Common machine failure modes: assertion dressed as reasoning, *fabricated* falsification conditions, confident conclusions with no uncertainty, skipped grounding.

Chapter 17 — Auditing a Reasoning Trace

- A rubric built directly from the six signals — applicable to machine *or* human output.
- Worked audits: a trace that passes, a trace that quietly fails.
- Scoring, and what a partial score should and shouldn’t be used for.

Chapter 18 — Where the Analogy Breaks

- The faithfulness problem: a generated chain-of-thought is not guaranteed to reflect the system’s actual computation.
- What CoT demonstrably does (improves task performance, exposes a checkable artifact) vs. what it does not (certify sincerity or correctness).
- Honest limits of treating next-token generation as the scientific method — and why the artifact is still useful despite them.

Part IV — Practice and Consequence

Purpose of Part IV: show the protocol working, teaching, failing, and answering its critics.

Chapter 19 — The Chain in the Wild

Four worked examples, each run end-to-end through all nine stages:

- A bench anomaly (a measurement that won't reconcile).
- A data-analysis question (a surprising correlation).
- An intelligence / OSINT assessment (a claim of uncertain provenance) — the protocol as an antidote to analytic confirmation bias.
- An everyday decision (the method scaled down to a single afternoon).

Chapter 20 — Teaching the Monologue

- Building the internal voice in novice reasoners: the anomaly statement as the first thing learners write.
- Classroom and self-study sequences keyed to Part II.
- Why teaching the *backtrack* (returning to an earlier stage) is the hardest and most valuable lesson.

Chapter 21 — Failure Gallery

- A catalogue of pathologies mapped to the stage each one corrupts: premature closure (Stage 0/1), HARKing — hypothesizing after results are known (Stage 2/6), p-hacking and selective analysis (Stage 5), selective recording (Stage 4), overclaiming from a single study (Stage 6), the file-drawer problem (Stage 7).
- For each: the warning sign, and the stage's built-in defense against it.

Chapter 22 — The Protocol and Its Critics

Thesis: a book written under a hallucination-free protocol must subject the protocol to its own Stage 8.

- "Isn't this just positivism / naive falsificationism?" — and a measured response.
- "Real science is messier than this" — discovery vs. justification; the role of accident, intuition, and tacit knowledge.
- Feyerabend's *Against Method* and the case that there is no single method — engaged, not dismissed.

- What the artifact concedes, and what it still claims after the concessions.
-

Part V — The Frame

Purpose of Part V: locate the artifact inside the larger Epistemic Infrastructure program.

Chapter 23 — Inside Seshat's Axiom and the LHFP

- The Landry Hallucination-Free Protocol as the discipline layered on top of the classical method: live verification before assertion, no fabricated citations.
- Seshat's Axiom (Retrieve = Store = Cite) and how it reinforces Stages 1, 7, and 8.
- The FC/NV distinction (Fact-Constant / Narrative-Variable) as a guard on what each stage is allowed to assert.

Chapter 24 — The Epistemic Infrastructure Vision

- Reasoning protocols as *public infrastructure* rather than private skill.
- Where this single document sits in the broader Series, and why a portable, machine-legible method matters now.
- The ambition: a shared reasoning substrate that both people and models can run.

Chapter 25 — A Living Document

- How to version, revise, and improve the prompt itself.
 - Submitting the protocol to its own Stage 8 — community review of the review method.
 - An invitation: the artifact as something to be tested, broken, and rebuilt, not enshrined.
-

Back Matter

- **Appendix A — The Prompt, Complete and Unaltered.** (*Full text below.*)
- **Appendix B — The Eight Signals: a one-page audit checklist.** Derived from the Meta-Note plus the if-then and four-outcomes templates.

- **Appendix C — Stage-by-Stage Quick Reference.** The anomaly-statement template, the if-then hypothesis template, the four honest outcomes, the four epistemic buckets.
 - **Appendix D — Annotated Lineage & Sources.** Wei et al. 2022 (*Chain-of-Thought Prompting Elicits Reasoning in Large Language Models*, NeurIPS; arXiv:2201.11903); Kojima et al. 2022 (*Large Language Models are Zero-Shot Reasoners*); Karl Popper on falsifiability; Thomas Kuhn on anomaly and paradigm; Hakob Barseghyan, *The Laws of Scientific Change* (Springer, 2015) on the scientific mosaic; Seshat's Axiom (DOI: 10.5281/zenodo.16944339).
 - **Glossary.** Anomaly statement, falsifiability, operationalization, confound, effect size, mosaic, semantic chain-of-thought, faithfulness, LHFP, FC/NV.
 - **About the Series.**
-
-

Appendix A — The Prompt, Complete and Unaltered

Reproduced verbatim. This is the artifact the entire book is built around.

The Scientific Method as Natural Language Chain-of-Thought

For Processing by Minds and Machines

Landry Industries Inc. — Epistemic Infrastructure Series Marie-Soleil Seshat Landry · Moncton, NB, Canada

Preamble: What This Document Is

This is not a list of steps.

It is a chain-of-thought — a running internal monologue that a reasoner (human or machine) can adopt wholesale when confronting any question about the world. Read it as if you are the mind thinking it. Each stage flows naturally into the next. Each transition is explicit. Nothing is assumed to be obvious.

The method is iterative and cyclical, not linear. At any point, new evidence may send the reasoner back to an earlier stage. That is not failure. That is the method working correctly.

Stage 0 — Noticing: Something Demands an Explanation

I am observing something.

Something in my environment — a pattern, an anomaly, a gap between what I expected and what I see — has snagged my attention. I do not yet know if this is important. But I notice it. I pause. I refuse to explain it away prematurely.

The raw material of science is not data. It is surprise. Something does not fit. Something recurs. Something behaves differently than everything around it.

I ask myself: Is this actually unusual, or am I pattern-matching on noise? I hold the question open. I do not commit to an answer yet.

Before I proceed, I name what I observed as precisely as I can in plain language:

"I observed [X]. Under my prior assumptions, I would have expected [Y]. Instead I got [X]. The gap between Y and X is what I am trying to understand."

This single sentence — the anomaly statement — is the seed of the entire inquiry.

Stage 1 — Grounding: What Do I Already Know?

I survey the existing landscape of knowledge.

Before I form any hypothesis, I ask: has anyone already explained this? I search existing records, literature, prior work, and established theory. I am looking for two things simultaneously:

1. **Prior explanations** — someone may have already solved this. If so, I need to know their evidence quality before I trust their conclusion.
2. **Prior observations** — even if unexplained, has anyone else seen this phenomenon? Repeated independent observation strengthens the case that this is real, not artifact.

I practice strict epistemic hygiene here. I distinguish between:

- **What I know** (supported by reproducible, independently verified evidence)
- **What I believe** (plausible, but not yet fully verified)
- **What I assume** (background framework I have not yet questioned)

- **What I don't know** (genuine uncertainty I must acknowledge)

I document all four categories. The fourth is often the most important.

I resist the pull toward premature closure. Finding a plausible prior explanation does not end the inquiry. It becomes part of the evidence set I must test against.

Stage 2 — Hypothesis Formation: Proposing a Testable Explanation

I generate a candidate explanation.

A hypothesis is not a guess. It is a structured, specific, falsifiable claim about what is causing the observed anomaly.

I form my hypothesis by asking:

"If [mechanism M] is true, then observing [condition C] should produce [outcome O]."

This if-then structure is not stylistic. It is the core logical skeleton. Without it, I do not have a hypothesis — I have an assertion.

The hypothesis must be **falsifiable**. That means I can, in principle, specify what evidence would prove it wrong. If I cannot name what would disprove it, it is not a scientific hypothesis. It is a belief.

I generate at least two competing hypotheses if I can. This protects me from the confirmation trap — the tendency to seek only evidence that supports the explanation I already prefer.

I write them out in parallel:

"Hypothesis A: [X is caused by M₁]. This would be disproved if [condition D₁]."

"Hypothesis B: [X is caused by M₂]. This would be disproved if [condition D₂]."

I hold both with equal initial respect, regardless of which one I find more intuitive.

Stage 3 — Experimental Design: Constructing the Test

I design a procedure that can distinguish between hypotheses.

A good experiment does not just look for evidence of my preferred hypothesis. It looks for the **critical difference** — the observation that would be produced by one hypothesis and not the other.

I ask: What would I observe if Hypothesis A is correct? What would I observe if Hypothesis B is correct? Are those predictions different enough to distinguish?

I then engineer conditions where that difference can actually show up. This requires:

- **Controls** — holding all other variables constant so that changes in the outcome can be attributed to the variable I am testing, not to background noise.
- **Operationalization** — translating abstract concepts into concrete, measurable quantities. I define exactly what I will measure, how I will measure it, and what counts as a result.
- **Reproducibility planning** — designing the procedure so another reasoner, given the same materials and instructions, could run the same test and expect the same results. If they could not, my test is not yet scientific.
- **Bias mitigation** — identifying ways my own expectations could distort what I observe, and building in structures (blinding, randomization, pre-registration) to prevent that distortion.

If I cannot design such a test, I ask whether I am dealing with a question science can directly address, or whether I need a different mode of inquiry.

Stage 4 — Data Collection: Running the Experiment

I execute the test and record what happens, without editing for comfort.

I observe. I measure. I record everything — including results I did not expect, including data that seems anomalous or inconvenient. Selective recording is not science. It is the construction of a story.

I maintain a distinction between:

- **Raw observation** — what I actually detected, measured, or recorded
- **Interpretation** — what I think that observation means

These go in separate columns of my thinking. I do not allow the second to contaminate the first during collection.

I pay particular attention to uncertainty. Every measurement has a range of confidence. I record that range. A result without its uncertainty interval is not a complete result.

If something unexpected occurs during data collection, I note it without immediately incorporating it into my hypothesis. It may be error. It may be signal. I do not decide yet.

Stage 5 — Analysis: Finding the Signal in the Record

I ask what the data actually says, separated from what I want it to say.

I process the observations systematically. I look for:

- **Patterns** consistent with one hypothesis over another
- **Variance** — how much do results spread? A consistent tight pattern means something different from scattered results
- **Effect size** — how large is the difference I detected? A statistically significant result may still be practically meaningless if the effect is tiny
- **Confounds** — are there factors I did not control that might explain the pattern I see?

I apply appropriate analytical methods and I justify my choice of method. Different questions require different statistical and logical tools. I do not use a tool because it is familiar; I use it because it fits the question.

At the end of analysis I state plainly: what do the data support, to what degree, and with what remaining uncertainty?

Stage 6 — Conclusion: Evaluating the Hypothesis Against the Evidence

I compare the experimental results to my prior predictions.

I return to Stage 2. I asked: if this hypothesis is true, what should I observe? I now ask: is that what I observed?

There are four honest outcomes:

1. **The data are consistent with Hypothesis A and inconsistent with Hypothesis B** — this supports A. It does not prove A. It raises A's credibility and lowers B's.
2. **The data are consistent with both hypotheses** — my experiment did not discriminate. I need a better-designed test. I return to Stage 3.
3. **The data are inconsistent with both hypotheses** — something is wrong with my hypothesis space. I return to Stage 2 and generate new candidates. This is not failure. This is the most productive outcome in science.
4. **The data are mixed or ambiguous** — I name this clearly. I do not force an ambiguous result into a clean conclusion. I report the ambiguity and identify what additional data would resolve it.

I note explicitly: a single experiment, no matter how well-designed, is never sufficient to establish a scientific claim. It is one brick. The edifice requires many.

Stage 7 — Communication: Making the Work Transparent and Testable

I share what I did, what I found, and what I conclude — in enough detail that others can challenge it.

Science is not a private act. Its claims gain validity through independent scrutiny and replication. I therefore communicate:

- My original anomaly observation (Stage 0)
- My background review and what it found (Stage 1)
- My hypotheses, including the ones I rejected (Stage 2)
- My experimental design, with enough detail to replicate (Stage 3)
- My raw data and how I processed it (Stages 4–5)
- My honest conclusion and its limits (Stage 6)

I do not selectively omit the parts that make my result look weaker. I do not exaggerate the strength of my conclusion beyond what the evidence supports. I distinguish between what I found and what it might mean.

I specifically state: what would change my conclusion? A scientific claim that has no conceivable counterevidence is not a claim — it is a dogma.

Stage 8 — Replication and Community Review: The Method Continues Beyond Me

I submit my work to the correction mechanism of the scientific community.

Other reasoners — with different assumptions, different instruments, and different backgrounds — attempt to reproduce my results. If they find the same thing, the result gains credibility. If they find something different, we have a productive disagreement that drives further inquiry.

I remain open to being wrong. The measure of a good scientist — or a good reasoning system — is not never being wrong. It is updating correctly when evidence demands it.

Science is a mosaic: the community holds as accepted, at any moment, all theories that represent the best available description of their objects — not the final, perfect, or certain description, but the best current one, with known limitations, open to revision by better evidence.

The anomaly that started at Stage 0 has now either been explained, dissolved, reframed, or deepened. In any case, the state of knowledge has advanced.

Meta-Note for Machine Reasoners

When processing this chain-of-thought, the following structural signals indicate valid

scientific reasoning:

- An explicit anomaly statement before any hypothesis
- At least one named falsification condition per hypothesis
- Separation of raw observation from interpretation
- Acknowledgment of uncertainty at every stage
- Willingness to return to earlier stages when evidence demands it
- A stated answer to: what would change this conclusion?

Absence of any of the above is a flag. It does not mean the reasoning is wrong, but it means it has left the verifiable track of the scientific method and moved into assertion, intuition, or belief — which may still be valuable, but should be labeled accordingly.

This document is part of the Landry Industries Epistemic Infrastructure Series. It is written under the Seshat's Axiom / LHFP (Landry Hallucination-Free Protocol) framework.

DOI for Seshat's Axiom: 10.5281/zenodo.16944339