

LANDRY INDUSTRIES · MARIE LANDRY'S SPY SHOP

SESHAT'S AXIOM



$$1 = 1 = 1$$

*What is cited is what was retrieved;
what was retrieved is what was stored.*

A STANDARD FOR VERIFIABLE MACHINES

MARIE-SOLEIL SESHAT LANDRY

FIRST EDITION · OCTOBER 2026 · AXIOM VERSION 5.1

SESHAT'S AXIOM

*What Is Cited Is What Was Stored — A Standard for
Verifiable Machines*

MARIE-SOLEIL SESHAT LANDRY

ORCID 0009-0008-5027-3337

LANDRY INDUSTRIES INC. · MONCTON, NEW BRUNSWICK, CANADA

COPYRIGHT AND LICENCE

© 2026 Marie-Soleil Seshat Landry. Released under the Creative Commons Attribution 4.0 International licence (CC BY 4.0). You may copy, adapt and redistribute this work for any purpose, with attribution.

AUTHOR IDENTIFIER

ORCID 0009-0008-5027-3337. This book is published under the name Marie-Soleil Seshat Landry. Earlier records also use the form Marie Seshat Landry^[1][2]. The ORCID iD, not the spelling of the name, is the key that joins them (Chapter 14).

PUBLISHER

Landry Industries Inc. and Marie Landry's Spy Shop, Moncton, New Brunswick, Canada.

EDITION

First edition, 8 October 2026. The canonical text of the axiom in this edition is Version 5.1, reproduced in Appendix A. Corrections will be issued as new versions; superseded versions stay on the record.

PERSISTENT IDENTIFIER

A DOI will be assigned when this edition is deposited on Zenodo. Cite the version DOI, not the concept DOI, so that readers retrieve exactly the text you read (Chapter 14).

SOURCES AND CONFORMANCE

Every web source was checked on 8 October 2026, either by opening it or through a search excerpt of its text. Appendix D lists each source with its version pin, access tier and conformance status. The declared store holds 59 records. 37 were verified directly or recomputed; 6 were verified only through a secondary source; 2 are used only as attributed claims; 12 are held in the author's restricted archive and were read in full; 1 could not be checked; and 1 failed and appears only as a case of failure. Web sources are pinned by version where one exists and otherwise by access date; snapshots are to be archived with the Zenodo deposit.

DISCLOSURE OF AI ASSISTANCE

Prepared with the assistance of Claude, an AI model made by Anthropic, under the author's direction and editorial authority. The AI ran the live source checks recorded in Appendix D. Responsibility for the content rests with the author.

NOT LEGAL ADVICE

Chapters 13 and 15 discuss evidence, regulation and court practice. Nothing in this book is legal advice.

TYPOGRAPHY

Set in EB Garamond, Cinzel and Poppins. Composed in HTML and rendered to PDF with WeasyPrint.

CONTENTS

How to Read This Book	v
Author’s Note: The Name	vi
PART I · FRAMING	
1. The Scribe and the Measurer	9
2. Why Machines Need a Scribe	13
PART II · THE AXIOM	
3. Stating Seshat’s Axiom	19
4. Three Kinds of Claim	25
5. The Axiom’s Own Record	29
6. How the Chain Breaks	33
PART III · THE BOUNDARY	
7. Provenance Is Not Truth	39
8. Objections and Limits	43
PART IV · PROTOCOLS	
9. The Landry Hallucination-Free Protocol	49
10. WTF and the Investigative Loop	53
11. Measuring Conformance	57
PART V · APPLICATIONS	
12. AI and Retrieval Systems	63
13. OSINT Chain of Custody	65
14. Open Science and Publishing	67
15. Law and Compliance	71
CLOSE	
Conclusion: A Standard for Verifiable Machines	73
Appendix A. The Canonical Text, Version 5.1	75
Appendix B. Conformance Checklists	79
Appendix C. Glossary	81
Appendix D. Sources and Conformance Register	83

HOW TO READ THIS BOOK

This book argues for a rule about records, so it is written to obey that rule. Six conventions make that visible.

1. **A bracketed number marks a fact-constant.** It points to Appendix D, where each source carries its version pin, its access tier and its conformance status. Several numbers mean several records support the statement.
2. **A sentence without a number is the author's own work:** a definition, an argument, a design or an inference. Inferences that carry weight are set in boxes labelled *Inference*.
3. **Boxes labelled *Narrative* contain framing,** such as myth and legend, attributed to the sources that tell it. They are not evidence for anything in the argument.
4. **Ancient dates name their chronology.** Different Egyptologists date the same period differently, and a date without its frame is incomplete (Chapter 7).
5. **Time-sensitive facts carry an "as of" date.** Unless stated otherwise, that date is 8 October 2026, the day every source was checked.
6. **Quotations are short and exact.** Longer material is paraphrased in the author's words and cited. Where the book quotes its own author's earlier documents, it quotes them exactly, errors included, because the record must not be tidied after the fact.

Appendix D also reports the book's own audit: how many of its sources were verified directly, how many only through a secondary source, how many are held in the author's private archive, and which claims failed. A reader who doubts any statement can follow its number to the record and check it.

AUTHOR’S NOTE: THE NAME

Seshat is my middle name, and it is also the name of the axiom this book defends. That overlap is where the book began. It is not part of the argument: the axiom stands or falls on its definitions and its tests, not on a name.

In Egyptian religion Seshat was a goddess of writing and measurement^[3]. The Egyptologist Richard H. Wilkinson, as quoted in a 2016 encyclopedia article, describes her helping the king in the “stretching the cord” ritual that laid out a temple’s plan^[4, 5]. Popular accounts add that a copy of every earthly writing was carried to a library of the gods. I found no primary source for that idea, so this book treats it as legend and attributes it to the article that tells it^[4].

I named the axiom for her because her two crafts are its two duties: measure exactly, and record what you measured. Everything that follows asks what it takes for a person, or a machine, to do both, and how anyone else can check that it was done.

The axiom has changed since I first wrote it down. Chapter 5 sets out every earlier version, including the mistakes. I left them on the record on purpose. An axiom about faithful records that quietly rewrote its own history would refute itself.

Marie-Soleil Seshat Landry
Moncton, New Brunswick, 8 October 2026

PART I

FRAMING

*Where the name comes from, and why machines that write need a
rule about records.*

THE SCRIBE AND THE MEASURER

Seshat was a goddess of writing and measurement. This chapter reports what the sources say about her, and marks the places where they cannot show what they claim.

NARRATIVE

This chapter is framing. Every statement about Seshat below is a statement about what a named source says. None of it is offered as evidence for the axiom, and the book takes no position on questions that Egyptologists leave open.

TWO CRAFTS

The *Encyclopaedia Britannica* describes Seshat as the Egyptian goddess of writing and measurement and calls her the ruler of books^[3]. Wikipedia’s article adds wisdom and knowledge to her domain^[6]. A 2016 article in the *World History Encyclopedia* by Joshua J. Mark says that her name literally means “female scribe”^[4].

The two crafts appear together in her iconography. Britannica reports that she was often shown holding a notched palm rib that stood for the passing of time, and that she was portrayed recording the king’s jubilees, cattle counts and campaigns as early as the Second Dynasty^[3]. Mark quotes the Egyptologist Rosalie David on one of her duties: writing the king’s name on the leaves of the persea tree, each leaf standing for a year of his allotted life^[4, 7]. A count kept on a rib and a name written on a leaf are the same act seen twice: something measured, then something recorded.

STRETCHING THE CORD

Wikipedia states that Seshat is attested from the Second Dynasty, where she assists King Khasekhemwy in the “stretching the cord” ritual^[6]. Mark quotes two passages by Richard H. Wilkinson on that ritual. In the first, Wilkinson notes that

depictions usually show the king performing the foundation rite with Seshat's help (*Symbol & Magic in Egyptian Art*, p. 174)^[4, 5]. In the second, he describes the rite as orienting a temple by astronomical observation: sighting the stars of a northern circumpolar constellation through a notched instrument called a merkhet to fix a true north–south line (*The Complete Temples of Ancient Egypt*, p. 38)^[4, 8].

Britannica fits the same picture. It says Seshat was believed to be expert in sighting the stars and planets, and that she was the keeper of ground plans and charts in the ritual^[3]. Her role reached into human institutions as well: Wikipedia records a Fourth Dynasty prince, Wep-em-nefret, who held the titles Overseer of the Royal Scribes and Priest of Seshat on a slab stela^[6].

INFERENCE

The cord ritual pairs an act of exact measurement with an act of lasting record. A temple laid out on a mismeasured line stays wrong for as long as it stands. That is the intuition the axiom formalises: capture exactly, and keep what you captured unchanged.

THE LIBRARY THAT CANNOT BE CITED

Mark's article states that whenever an author created a story, inscription or book on earth, an ethereal copy passed to Seshat, who placed it in the library of the gods^[4]. It is a striking image for a book about faithful copies. It is also a test case. The sentence carries no specific citation; the article lists a general bibliography of ten books but does not tie this claim to any of them^[4].

Under the axiom, the only fact-constant available here is an attributed one: *Mark (2016) states that copies of earthly writings went to a library of the gods*. That statement can be checked against the article, and it is true. Whether ancient Egyptians held the belief is a separate claim, and nothing consulted for this book grounds it. The same caution applies to Mark's statement that Seshat never had her own temples, cult or formal worship^[4]: it is attributed, not independently verified.

THE HEADDRESS NOBODY CAN READ

Mark describes Seshat's headdress as a seven-pointed star arched by a crescent in the form of a bow^[4]. Jacqueline Thurston, author of *Sacred Deities of Ancient*

Egypt, writing for the Archive for Research in Archetypal Symbolism, states that the precise meaning of the emblem remains unknown^[9].

On 8 October 2026 a search returned a Wikipedia talk-page discussion in which one contribution cited an outside website claiming that the emblem depicts a hemp leaf, while another editor noted that stars in Egyptian convention had five points rather than seven^[10]. When the page was requested again minutes later, it could not be fetched^[10]. Chapter 6 uses this as a case of decay. Here it serves a simpler purpose: a contested reading, offered by a source whose own source is a website, is not something this book will assert. The emblem's meaning stays open.

ONE NAME, MANY RECORDS

Seshat is a crowded name in modern research. It names a global history databank of past societies, whose data are published as snapshots on Zenodo^[11]; an open-source tool for managing annotation campaigns of speech recordings^[12]; and a predictive mutation-analysis model for software testing, whose authors note that the name comes from the Egyptian deity of writing and record keeping^[13].

DISAMBIGUATION

In this book, **Seshat's Axiom** means the rule set out in Chapter 3 and Appendix A, authored by Marie-Soleil Seshat Landry (ORCID 0009-0008-5027-3337). Cite it with the author, the version number and the version DOI. It has no connection to the databank or the software tools named above.

A WARNING IN THE RECORD

The same encyclopedia article that supplies most of this chapter dates the Middle Kingdom to c. 2613–2181 BCE^[4]. The encyclopedia's own timeline gives those years to the Old Kingdom and dates the Middle Kingdom to 2040–1782 BCE^[14]. A reader who copied the article faithfully would copy the error with it. Chapter 7 returns to this case, because it shows the exact edge of what the axiom can and cannot promise.

WHY MACHINES NEED A SCRIBE

Language models can produce fluent text that is tied to no record. The cost falls on whoever trusts the citation.

FLUENT AND UNANCHORED

A survey of hallucination in natural-language generation by Ji and colleagues, published under DOI 10.1145/3571730, reports that deep-learning text generation is prone to producing unintended text, which degrades performance and fails users in many real-world settings^[15]. The survey’s arXiv record has seven versions; the latest is dated 14 July 2024^[15].

Emily M. Bender and her co-authors put the mechanism bluntly in 2021. They described a language model as a system for “haphazardly stitching together sequences of linguistic forms” seen in its training data, guided by probabilities and without reference to meaning^[16, 17]. One need not accept every implication of that description to accept its practical consequence: a model can produce a sentence that looks exactly like a cited fact without any record standing behind it.

CASE · MATA V. AVIANCA

In a federal lawsuit against the airline Avianca in the Southern District of New York, counsel for the plaintiff filed a submission on 1 March 2023 that cited non-existent judicial opinions, with fake quotations and fake citations generated by ChatGPT^[18]. When asked whether one of the cases was genuine, the chatbot replied that it “is a real case”^[19]. On 22 June 2023, Judge P. Kevin Castel issued an opinion and order imposing sanctions^[18]. Reports of the order state that the lawyers were required to pay a \$5,000 penalty and to write to each judge falsely named as an author of the fake opinions^[20].

Under the axiom: every one of those citations fails Resolvability. There was no record to retrieve, so nothing could be stored, retrieved or matched. This is the failure mode called fabrication (Chapter 6).

EVEN TRUE CITATIONS DECAY

Fabrication is the dramatic failure. Decay is the quiet one. In a 2014 study, Jonathan Zittrain, Kendra Albert and Lawrence Lessig found that more than 70% of the URLs in the *Harvard Law Review* and other journals, and 50% of the URLs in United States Supreme Court opinions, no longer led to the information originally cited^[21]. A citation that was honest when written can stop resolving within years. Their proposed remedy was for libraries to preserve the contents of cited links over the long term^[21], an idea that reappears in this book as the snapshot rule.

PARTIAL ANSWERS

Two engineering ideas move in the right direction. Retrieval-augmented generation, introduced by Lewis and colleagues in 2020, pairs a sequence-to-sequence model with a dense vector index of Wikipedia reached through a neural retriever^[22]. Its authors report that the resulting models generate more specific, diverse and factual language than a parametric-only baseline^[22]. Pointer-generator networks, described by See, Liu and Manning in 2017, address a related problem: their authors observe that sequence-to-sequence summarisers can reproduce factual details inaccurately, and their architecture can copy words directly from the source text by pointing, which aids accurate reproduction^[23].

INFERENCE

Retrieval raises the odds that a model sees the right record. Copying raises the odds that it reproduces the record's words. Neither, on its own, guarantees that the cited passage is the one that was stored, that the stored passage is the one that was published, or that the claim in the output matches the passage it cites. What is missing is a single identity rule that runs the whole length of the chain, from capture to citation, and a way for a stranger to check it. That rule is the subject of Part II.

WHAT A SCRIBE ADDS

A human scribe in a records office does four things a fluent writer need not do. The scribe notes where each document came from and when. The scribe keeps the original rather than a memory of it. The scribe copies exactly, marks every

omission and names every change. And the scribe keeps commentary apart from the record. Seshat's Axiom asks machines, and the people who use them, to do the same four things, and it defines each one closely enough that an auditor can tell whether it was done.

PART II

THE AXIOM

*The rule, the kinds of claim it governs, its own history, and the
ways the chain breaks.*

STATING SESHAT'S AXIOM

*A fact must reach its citation unchanged from the record it came from.
Everything else in this chapter makes that sentence testable.*

SESHAT'S AXIOM · VERSION 5.1

WHAT IS CITED IS WHAT WAS RETRIEVED; WHAT WAS
RETRIEVED IS WHAT WAS STORED. ($I = I = I$)

1. **Fidelity (Store = Retrieve).** A record is captured with its provenance: source, exact locator, version, capture time, a content hash and an archived snapshot. Anything later retrieved is identical to what was stored.
2. **Grounding (Retrieve = Cite).** Every fact-constant in an output matches its cited record exactly, or under a declared, reproducible transformation within a declared reference frame. The record's qualifiers and its "as of" date travel with it.
3. **Resolvability (precondition).** Every citation is version-pinned and locates the exact passage. It is public, held by a named auditor, or committed by a salted hash for later disclosure.

Corollary, Transparency. Every inference is labelled as inference, and every chain of inference ends in grounded fact-constants.

Duty, Disclosure. If the declared store holds a record that contradicts a cited fact-constant, the output says so.

Non-claim. The axiom guarantees faithfulness to sources, not their truth.

WHAT THE THREE ONES ARE

The three terms of $I = I = I$ are one piece of content seen at three moments: as stored, as retrieved and as cited. The equals signs claim that it did not change between those moments. Earlier versions named the terms as actions (Read, Copy, Paste, Cite), and Chapter 5 traces that history. Version 5 keeps the actions

as the pipeline and puts the identity on the content, because actions cannot be equal to each other, but what they carry can be.

$$R(S(x)) = x \quad \text{and} \quad c \equiv T(R(S(x))), \quad T \in \{\text{identity}\} \cup \{\text{declared transformations}\}$$

Read the notation as follows. S stores a record x ; R retrieves it; c is the claim as it appears in the output; T is either no change at all or a transformation that the output names and that anyone can reproduce, such as converting pounds to kilograms. The first equation is Fidelity. The second is Grounding. Resolvability is what lets a stranger evaluate both.




Figure 3.1 · The three clauses along one pipeline

CLAUSE I: FIDELITY

Fidelity governs the store. It has two halves. The first is capture: a record enters the store with enough information to find the original again and to prove later that nothing has changed. The second is retrieval: what comes out of the store is bit-for-bit what went in. In practice the second half is checked by recomputing a cryptographic hash. The U.S. Secure Hash Standard, FIPS 180-4, specifies hash algorithms whose digests are used to detect whether a message has changed since its digest was made^[24].

Fidelity also covers metadata. If a store binds an identifier to a title, that binding is a fact-constant and can be wrong. Chapter 6 opens with exactly that error.

THE STORED RECORD

The axiom does not prescribe a database, but it does prescribe what a stored record must carry. The fields below merge the six-part fact of the author's Fact-

Constant/Narrative-Variable architecture (content, provenance, temporal scope, semantic scope, metadata and version)^[25] with the four components that the WTF 7.0 framework required of every fact (value, scope, provenance and temporal marker)^[26].

FIELD	WHAT IT HOLDS	WHY THE AXIOM NEEDS IT
Content	The exact bytes or text captured	The object of Fidelity
Content hash	SHA-256 digest of the content	Proves retrieval returned the same content
Source and version	DOI, URL or archive ID, with version number or date	Resolvability; protects against drift
Locator	Page, section, table, paragraph or timestamp	Lets an auditor find the exact passage
Capture time and snapshot	When it was taken; an archived copy	Survives link decay
Reference frame	Chronology, units, calendar, currency year	Makes numbers and dates comparable
Qualifiers and scope	“In mice”, “preliminary”, test conditions	Grounding must keep them attached
As-of date	When the fact was true according to the source	Facts about the present expire
Access tier	Public, named auditor, or salted commitment	Resolvability for restricted sources

CLAUSE 2: GROUNDING

Grounding governs the output. It says that every fact-constant in what a person or machine writes must match the record it cites. “Match” has a strict meaning. Either the text is identical, or it differs only by a transformation that the output declares and that a checker can reproduce exactly. Converting a temperature, rounding to a stated precision and reformatting a date are reproducible. Paraphrasing a quotation is not, and neither is translation, which is why Chapter 4 requires the original wording to travel with any translated quotation.

Grounding also protects context. A quotation stripped of “in mice” or “preliminary results suggest” may still be a string that appears in the source, yet it says something the source did not say. The qualifier rule closes that gap: the scope, conditions and uncertainty attached to a fact in the source travel with it.

CLAUSE 3: RESOLVABILITY

Resolvability is the precondition for checking the other two. A citation must name a specific version of a specific record and point to the exact passage. A citation to a whole book, to a web page that may change tomorrow, or to a private file nobody else can open does not let an auditor do anything.

Some legitimate sources cannot be public: confidential informants, paywalled databases, sealed records. Version 5.1 offers two routes for them. A source may be held by a *named* auditor, so that trust attaches to an identifiable person or institution rather than to an anonymous assurance. Or the author may publish a *salted hash commitment*: the digest of the record combined with a secret random value, to be revealed later. Chapter 8 explains why the salt matters.

TRANSPARENCY, DISCLOSURE AND THE NON-CLAIM

The three clauses govern fact-constants. Reasoning is governed by a corollary: label every inference, and make sure every chain of reasoning bottoms out in grounded facts. The corollary does not require inferences to be correct. It requires them to be visible, so that a reader can see exactly where the record ends and the argument begins.

The Disclosure duty addresses a failure the clauses cannot catch on their own. An author can cite every source faithfully and still mislead by citing only the sources that agree. Disclosure is checkable relative to a declared store: if the store holds a record that contradicts a cited fact, an auditor can find it, and the output must have mentioned it.

Finally, the non-claim states the axiom's boundary in one sentence. A perfectly conforming output can faithfully repeat a false source. Chapter 7 is devoted to that edge.

CONFORMANCE

An output conforms when an independent auditor, following its citations, can verify every fact-constant in it. Each fact-constant ends in one of three states: **verified**, **deferred** (committed by hash and awaiting disclosure) or **failed**. Conformance is transitive. Citing a conforming output grounds the fact; citing a non-conforming one grounds only the attributed claim that its author said it. A cycle

of citations, in which A cites B and B cites A, grounds nothing. A single declared, version-pinned baseline reference may ground general knowledge for a whole work, so that no author has to cite the boiling point of water.

This book declares no baseline reference. Every fact-constant in it carries its own citation, and Appendix D reports the result.

THREE KINDS OF CLAIM

The axiom is strict about facts and silent about style. Everything therefore depends on telling the two apart, and the test is whether a checker can look the claim up.

THE LOOKUP TEST

The author's Fact-Constant/Narrative-Variable (FC/NV) architecture separates immutable factual constants from mutable explanatory hypotheses^[25]. Version 5.1 keeps that separation and supplies a single operational test for it:

RULE

Can a checker confirm the statement by looking it up in the record it cites? If yes, it is a fact-constant and the three clauses apply. If no, it is either an inference, which must be labelled, or narrative, which is free. When in doubt, an assertion is a fact-constant: the burden falls on the author, not the reader.

The test classifies by what a claim does, not by what the author calls it. Labelling a factual assertion “narrative” does not move it out of reach of the axiom. This closes the most obvious loophole in any fact/opinion scheme.




Figure 4.1 · The lookup test as four questions, asked in order

FIVE CLASSES

Fact-constants are names, numbers, dates, quotations, identifiers and any other assertion a checker can find in the cited record. **Derived fact-constants** are results of reproducible computation from grounded inputs: a total, a ratio, a hash. They cannot be looked up, because nobody has published them yet, but they can be re-computed, and that is just as strong. **Attributed claims** take the form “Source S states P”. They are fact-constants about the source, not about the world, and they are the honest way to report a claim one cannot verify. **Inferences** are the author’s conclusions; they must be labelled and must rest on grounded fact-constants. **Narrative-variables** are everything else: wording, order, emphasis, tone.

CLASSIFYING REAL SENTENCES

The table applies the test to sentences drawn from this book’s own sources.

SENTENCE	CLASS	WHAT MAKES IT CONFORM
FIPS 180-4 specifies SHA-1 and the SHA-2 family of hash functions.	Fact-constant	Citation to the standard ^[27]
The SHA-256 digest of “Hello World!”, in Named Information form, is f40xZX_x_F05LcGBSKHWXfwtSx-j1nc0St3SABJtkGk.	Derived fact-constant	Recomputed for this book; matches the example printed in RFC 6920, section 8.1 ^[28, 29]
Mark (2016) states that copies of earthly writings went to a library of the gods.	Attributed claim	Checkable against the article ^[4]
Copies of earthly writings went to a library of the gods.	Fails as stated	No consulted record grounds it as a fact; it must be attributed

SENTENCE	CLASS	WHAT MAKES IT CONFORM
The Middle Kingdom began c. 2055 BCE.	Incomplete fact-constant	Must name its frame: in the Shaw chronology ^[30]
Fifty percent of URLs in U.S. Supreme Court opinions no longer led to the cited information.	Fact-constant with qualifier	Must keep “in a 2014 study” ^[21]
Chapter 1 should therefore treat the library story as legend.	Inference	Labelled; rests on the attributed claim above
Seshat’s two crafts mirror the axiom’s two duties.	Inference (interpretive)	Labelled as the author’s reading
A temple laid out on a mismeasured line stays wrong.	Narrative-variable	An image, not a checkable claim

ATTRIBUTION: TURNING CLAIMS INTO CHECKABLE FACTS

Attribution is the move that lets the axiom coexist with interpretation. “Hemp fibre outperforms glass fibre” is a contested comparison that no single lookup can settle. “Study X reports that hemp fibre outperformed glass fibre in tensile test Y” is a fact-constant: open study X, find test Y, compare. Attribution never forbids interpretation. It simply records who interpreted what, so that the reader can weigh the interpreter.

Attribution has a cost, which is honesty about strength. An attributed claim grounds only what the source said. A chain of attributed claims, each citing the last, grounds nothing about the world until it reaches a record that does. Chapter 7 calls this the difference between provenance and truth.

DECLARED TRANSFORMATIONS AND REFERENCE FRAMES

Version 5.1 allows a fact-constant to differ from its record only by a transformation the output declares and anyone can reproduce. The test is reproducibility: two careful people applying the declared transformation must get the same result.

CHANGE	STATUS	CONDITION
Unit conversion	Allowed	Declare the conversion; keep stated precision
Rounding	Allowed	Declare the rule (for example, to two significant figures)
Date or number formatting	Allowed	No change of value or calendar
Character encoding, case, whitespace	Allowed	Not inside a quotation where case carries meaning

CHANGE	STATUS	CONDITION
Translation	Not a transformation	The original wording must travel with it, or the translation is labelled as the author's rendering
Paraphrase inside quotation marks	Forbidden	Quote exactly or paraphrase without quotation marks
Joining sentences, dropping words	Forbidden unless marked	Use an ellipsis for every omission
Dropping a qualifier	Forbidden	Scope and uncertainty travel with the fact

A reference frame is different from a transformation. It is the system a value only has meaning within: a chronology for ancient dates, a unit system for measurements, a calendar, a coordinate datum, a currency and its base year. Two sources can both be faithful and still disagree because they use different frames. Naming the frame turns an apparent contradiction into two comparable facts.

THE AXIOM’S OWN RECORD

Before Version 5, one name retrieved several different rules. By its own standard the axiom was failing Fidelity, and the remedy is the one it prescribes for any record: pin every version, keep them all, and say what changed.

THE LINEAGE, RECORD BY RECORD

The table lists, in date order, the documents found in the author’s archive that state the axiom or a direct precursor, followed by the two versions written in October 2026. Where a file carries no internal date, the table says so; the date a file reached the archive is not an authorship date. The archive documents are held in the author’s private Google Drive and the Version 5 foundation document is a private working document, so their access tier is restricted, with the author as named holder. Chapter 14 sets out the plan to deposit them publicly with this book.

DATE	RECORD	STATEMENT OF THE RULE
Sept. 2024	Seshat Scientific Investigative Framework (SSIF) ^[2]	Not yet the axiom. The earliest Seshat-named method in the archive: evidence handling, chain of custody, 5W+1H analysis
10 Jan. 2026	White Paper: the Landry Hallucination-Free Protocol ^[31]	A model should switch to copying when a fact is required; verification compares the hash of the output with the source
Jan. 2026 (archived 27 Jan.)	Deterministic Retrieval and Constrained Inference, Protocol v7.0 ^[32]	Facts are copied verbatim; narrative is labelled as hypothesis. Its closing directive: “Find the Constant. Copy the Constant. Cite the Constant.”
Jan. 2026	Fact-Constant/Narrative-Variable Architecture, v2.0 ^[25]	Facts as immutable six-part records; narratives as falsifiable hypotheses
26 Mar. 2026	Seshat’s Axiom (1=1=1=1) Prior Art Release ^[33]	Four operations, Read = Copy = Paste = Cite, with zero change in information entropy across them

DATE	RECORD	STATEMENT OF THE RULE
Undated	Technical Report: Seshat's Axiom ^[34]	1=1=1: the fact retrieved is identical to the fact stored and to the fact cited; copy, paste in quotation marks, cite inline
Mar. 2026	Technical Report, corrected ^[35]	Same rule; links repaired; adds a section declaring the report self-compliant
Undated	"seshats axiom.pdf" ^[36]	1=1=1=1: Fact = Representation = Verification = Meaning
Undated	Synthesizing WTF Frameworks into 9.0 ^[26]	States the axiom as Retrieved = Stored = Cited, and describes a different "1=1=1 Law" from WTF 7.0: value, scope, provenance and temporal marker
8 Oct. 2026	Book Foundation, Version 5 ^[37]	Three clauses, a corollary, a duty and a non-claim
8 Oct. 2026	This book, Version 5.1	Version 5 plus salted hash commitments (Chapter 8)

Two features of this record deserve plain statement. First, the name $I=I=I=I$ was used for two different rules: the Prior Art Release's four operations^[33] and WTF 7.0's four required fields^[26]. A name that retrieves two different contents is a Fidelity failure in miniature. Version 5 therefore retires " $1=1=1=1$ " as a name, while keeping the substance of both rules. Second, the Technical Report's section 3.1 already defined $1=1=1$ as the fact retrieved being identical to the fact stored and to the fact cited^[34]. Version 5 descends directly from that sentence.

WHAT VERSION 5.1 KEEPS, CHANGES AND WITHDRAWS

EARLIER ELEMENT	SOURCE	TREATMENT IN 5.1
Read, Copy, Paste, Cite as four operations	Prior Art Release	Kept as the pipeline. The identity now sits on the content, not on the operations
Equal Shannon entropy at every operation	Prior Art Release	Withdrawn. Entropy cannot detect a change (see below). Replaced by equality of content hashes
Perfect read fidelity: the read state equals source truth	Prior Art Release	Withdrawn. It contradicts the non-claim. Replaced by capture fidelity: the stored record equals the published source
Anti-summarisation guardrail	Prior Art Release	Scoped to fact-constants. Narrative may be summarised
Visible boundary between verified data and commentary	Prior Art Release	Kept as the Transparency corollary

EARLIER ELEMENT	SOURCE	TREATMENT IN 5.1
Auditable closure: an auditor can trace the citation back	Prior Art Release	Kept as the definition of conformance
Selective quotation listed as a source of error	Prior Art Release	Kept as the qualifier rule and the Disclosure duty
“Immutable upon publication”	Prior Art Release	Reframed: each version is immutable and pinned; the axiom changes only by new versions
Paste facts in quotation marks; cite inline	Technical Report	Kept for verbatim quotation. Converted figures carry a declared transformation instead
Paraphrasing systems are “fundamentally non-intelligent”	Technical Report	Withdrawn as an overclaim. Paraphrasing narrative is allowed; paraphrasing a fact-constant is a Grounding failure
No inference inside the fact layer	Technical Report	Kept: the fact layer retrieves, it does not generate
Fact = Representation	“seshats axiom.pdf”	Kept, limited to declared, reproducible transformations
Fact = Verification	“seshats axiom.pdf”	Becomes the definition of conformance
Fact = Meaning	“seshats axiom.pdf”	Withdrawn. Equality of meaning cannot be verified; translation is not reproducible
Value, scope, provenance, temporal marker	WTF 7.0, as described in the 9.0 synthesis	Kept as required fields of every stored record (Chapter 3)
A human-validated hypothesis is promoted to a Constant	Protocol v7.0	Changed. An inference never becomes a fact-constant. It can be recorded only as an attributed claim: “validated by X on date D”
Constants are “binary (True/False) and immutable”	Protocol v7.0	Changed. A constant is fixed, not thereby true
Minimum of three citations per fact	WTF 10 Framework Engine ^[38]	Moved to the source-validity layer (Chapter 7). One grounded citation satisfies Grounding; corroboration measures belief

WHY ENTROPY WAS THE WRONG TEST

The Prior Art Release required that the Shannon entropy of the information be equal at every operation, and treated any change as a sign of compromise^[33]. Equal entropy is necessary for an exact copy but nowhere near sufficient. Entropy

measures how varied the symbols are, not which symbols appear in which order. Any rearrangement of the same characters has exactly the same entropy.

DERIVED FACT-CONSTANT · COMPUTED FOR THIS BOOK

The strings “2040 BCE” and “2400 BCE” contain the same characters, so both have a character-level Shannon entropy of 2.75 bits per character. Their SHA-256 digests differ completely: they begin `ba3a5f37` and `7a046b23` respectively^[39]. An entropy check would pass a swapped date; a hash check catches it.

The example is not contrived. The two strings differ by the transposition of two digits, which is precisely the kind of error a table of ancient dates can suffer. Version 5 replaces the entropy constraint with hash equality, which detects any change at all.

THE REPORT THAT FAILED ITS OWN TEST

The corrected Technical Report ends with a section headed “Axiom Compliance Verification”, which declares the report self-compliant and states that all its factual claims carry verbatim citations^[35]. Checking its seven attributed quotations against their sources found four that fail: an altered quotation, a sentence attributed to the wrong standard, a quotation that does not appear in the cited document, and two sentences spliced into one. Chapter 11 sets out the full audit. The point here is narrower. The report’s author is the author of this book, and the axiom applies to her too. Recording the failures openly is the axiom working as designed.

HOW THE CHAIN BREAKS

Every failure of a citation chain breaks one of the three equalities or one of the two duties. Naming the failure tells you which clause catches it.

A TAXONOMY FROM THE RECORD

Each row below is a distinct way for a cited fact to go wrong. Each comes with a real case from the sources consulted for this book, so the taxonomy is grounded rather than imagined.

FAILURE	WHAT HAPPENS	CAUGHT BY	CASE
Capture error	The store records a wrong binding	Fidelity	A DOI noted as the axiom paper resolves to the Hempoxy document ^[1, 37]
Corruption	Stored content changes	Fidelity (hash)	Detected by recomputing the digest ^[24]
Drift	A later retrieval returns a different version	Fidelity, version pin	The RAG paper has four arXiv versions ^[22]
Fabrication	The citation points to nothing	Resolvability	Non-existent opinions in <i>Mata v. Avianca</i> ^[18]
Alteration	A quotation is reworded	Grounding	Bender et al. misquoted in the corrected report ^[35, 17]
Splice	Two sentences joined without an ellipsis	Grounding	Lewis et al. abstract ^[35, 22]
Misattribution	A claim is pinned to the wrong record	Grounding	SHA-256 called a SHA-3 variant under FIPS 180-4 ^[35, 27]
Phantom quotation	Quoted words absent from the cited record	Grounding	A sentence attributed to RFC 6920 ^[35, 29]
Frame error	A value is read in the wrong system	Grounding (frame)	Old Kingdom dates labelled Middle Kingdom ^[4, 14]
Decay	A citation stops resolving	Resolvability (snapshot)	Talk:Seshat ^[10] ; legal link rot ^[21]

FAILURE	WHAT HAPPENS	CAUGHT BY	CASE
Wrapper link	A citation resolves to a search page, not the record	Resolvability	Seven of 25 links in the uncorrected report ^[34, 40]
Unlabelled inference	A judgment is presented as fact	Transparency	“Fundamentally non-intelligent” ^[34]
Selective citation	Contrary records are left out	Disclosure	Citing the hemp-leaf reading alone ^[10, 9]

CAPTURE: THE ERROR THAT STARTS IN THE STORE

The author’s saved notes recorded that DOI 10.5281/zenodo.16944339 identified the paper in which Seshat’s Axiom, the Landry Hallucination-Free Protocol and the FC/NV architecture were published^[41, 37]. On 8 October 2026 the DOI resolved to a different work: *Hempoxy – A Master Document on the Science, Synthesis and Strategy of a Sustainable Nanocomposite System*, version v1, published 25 August 2025^[1]. No citation in any later document could repair that, because every later citation would faithfully reproduce the wrong binding. Only a check against the live record at capture time catches a capture error, which is why Fidelity begins at the moment of storage.

ALTERATION, SPLICE AND PHANTOM: THREE GROUNDING FAILURES

Grounding failures are the most common because they happen at the last step, when someone writes. The corrected Technical Report shows three kinds.

Alteration. In its reference list the report quotes Bender et al. as beginning “Language models are a system for...”^[35]. The paper’s wording, as recorded by Quote Investigator, begins with “an LM is a system for”^[17]. The meaning barely changes. The axiom does not care: inside quotation marks, words are fact-constants. The same report’s Key Judgment, a few pages earlier, quotes the passage correctly^[35], which shows how easily a correct record and a drifted copy coexist in one document.

DERIVED FACT-CONSTANT · COMPUTED FOR THIS BOOK

The paper's 14-word excerpt and the report's reworded version produce SHA-256 digests beginning 1ea1ce7a and 194b3444^[39]. A machine detects in microseconds the change a human reader skims over.

Splice. The report quotes Lewis et al. on retrieval-augmented generation as a single sentence^[35]. In the paper's abstract, the quoted words come from two consecutive sentences, with the words between them removed and no ellipsis to show it^[22]. The report also cites the arXiv abstract page without a version, although the record has four versions, the latest dated 12 April 2021^[22].

Phantom quotation. The report places in quotation marks a sentence describing RFC 6920 as defining a URI scheme that identifies information by its content rather than its location^[35]. That sentence does not appear in the RFC's text as retrieved on 8 October 2026. The RFC's abstract says instead that it defines ways to identify a digital object using the output of a hash function^[29]. The paraphrase is not wrong as a description; it is wrong as a quotation.

MISATTRIBUTION

The same report quotes NIST's FIPS 180-4 as describing SHA-256 as a "SHA-3 variant"^[35]. FIPS 180-4 specifies SHA-1 and the SHA-2 family, to which SHA-256 belongs; the SHA-3 family is specified separately, in FIPS 202^[27]. Misattribution is dangerous because it borrows the authority of the cited source for a claim the source never made. A reader who trusts NIST will trust the error.

DECAY AND WRAPPER LINKS

Resolvability can fail without anyone doing anything wrong. The Talk:Seshat page returned by a search on 8 October 2026 could not be fetched minutes later^[10]. The scale of the problem in formal citation is documented: the 2014 Zit-train study found half the links in Supreme Court opinions no longer led to the information cited^[21]. The only defence is to keep a snapshot at capture, which is why Version 5 makes the snapshot part of the stored record rather than an optional extra.

A subtler failure is the wrapper link. In the uncorrected Technical Report, seven of the 25 listed URLs pointed to a Google search for the source rather than to the source itself^[34, 40]. A search page is not a record. It can return different results tomorrow, so the citation cannot be pinned. The corrected report replaced most of these links^[35], an example of the axiom's remedy applied correctly.

UNLABELLED INFERENCE AND SELECTIVE CITATION

Two failures involve no misquoted fact at all. The Technical Report's judgment that systems which paraphrase facts without verbatim citation are "fundamentally non-intelligent"^[34] is an inference about intelligence, presented with the same force as a fact. Version 5 does not ban the judgment; it requires the label, so the reader knows to weigh it as argument.

Selective citation is harder to see because every individual citation is accurate. Citing only the website that reads Seshat's emblem as a hemp leaf, while leaving out the editor's objection and Thurston's statement that the emblem's meaning is unknown^[10, 9], would pass Grounding and still mislead. The Disclosure duty is the clause that catches it.

INFERENCE

Every case in this chapter was found by one method: open the cited record and compare. None needed expertise in the subject matter. That is the practical promise of the axiom. It turns "is this right?", which needs an expert, into "does this match?", which needs only patience, or a machine.

PART III

THE BOUNDARY

*What the axiom cannot promise, and the strongest objections to
what it does.*

PROVENANCE IS NOT TRUTH

The axiom asks whether you copied the record correctly. It cannot ask whether the record is right. Those are two questions, and a trustworthy output answers both.

A FAITHFUL COPY OF A WRONG DATE

Mark's article on Seshat states that, beginning in the Middle Kingdom, dated there to c. 2613–2181 BCE, she recorded the spoils of war^[4]. The same encyclopedia's timeline for the First Intermediate Period gives c. 2613–2181 BCE as the dates of the Old Kingdom and 2040–1782 BCE as the dates of the Middle Kingdom^[14]. A chapter of a 2020 Cambridge University Press monograph by Lisa K. Sabbahy, citing Shaw (2004), dates the Old Kingdom to roughly 2686–2181 BCE and the Middle Kingdom to 2055–1650 BCE^[30].

Three things follow. First, quoting the article's date exactly, with its citation, satisfies all three clauses: it is a faithful copy of a real record. Second, the record is wrong, by its own publisher's standard. Third, even the correct dates differ between sources, not because either is careless but because Egyptologists use different chronologies.

INFERENCE

The axiom guarantees the first point and is silent on the second. The reference-frame rule handles the third: "Middle Kingdom, 2055–1650 BCE (Shaw chronology)" and "Middle Kingdom, 2040–1782 BCE (World History Encyclopedia)" are both faithful, and naming the frame shows that they are not in contradiction.

THE SECOND QUESTION: SHOULD WE BELIEVE IT?

Intelligence services met this problem long before language models. The NATO grading system, often called the Admiralty Code, rates the reliability of a source from A to F and the credibility of an item of information from 1 to 6, and assesses

the two separately so that a source’s reputation does not inflate the credibility of a particular report^[42]. Under the NATO standard STANAG 2511, information from a usually reliable source judged probably true is graded B2; the same source with information whose truth cannot be judged is graded B6^[43].

The United States intelligence community’s analytic standards, set out in Intelligence Community Directive 203, list nine analytic tradecraft standards^[44]. The first requires that a product properly describe the quality and credibility of its underlying sources, data and methods; the third requires that it properly distinguish underlying intelligence from analysts’ assumptions and judgments^[44]. The standards were established in 2007 and revised in January 2015^[45]. The third standard is the Transparency corollary in another vocabulary.

A TWO-LAYER RECORD

Version 5.1 keeps the two questions apart by giving every fact-constant two separate assessments.

LAYER	QUESTION	SCALE	WHO CAN ANSWER IT
Conformance (the axiom)	Did the claim reach the citation unchanged?	Verified, deferred, failed	Any auditor with the record
Validity (outside the axiom)	How far should the record be believed?	For example, Admiralty reliability and credibility	An assessor with subject knowledge

Keeping the layers apart prevents two errors. A verified citation to an unreliable source is not upgraded by its accuracy, and a credible claim with a broken citation is not excused by its plausibility. The author’s WTF 10 Framework Engine sets a minimum of three citations per fact and an Admiralty vetting floor of B-2^[38]. Under Version 5.1 those are validity policies. They are sensible ones, but they measure belief, not faithfulness, and they belong to the second layer.

PRIMARY AND SECONDARY RECORDS

The W3C’s provenance model, PROV-DM, a Recommendation since 30 April 2013, includes a relation for primary sources: material produced by an agent with direct experience of a topic, at the time, without the benefit of hindsight^[46]. The model notes that secondary sources should reference the primary sources they derive from, so that their reliability can be investigated^[46].

This book cites Wilkinson and David only through Mark's article, because their books were not consulted directly. Appendix D therefore lists those claims as *verified via secondary*: the axiom confirms that Mark quotes them, and the validity layer notes that a primary check remains to be done. That label is the honest price of working from a secondary source.

DISCLOSURE: THE DUTY THAT LIMITS CHERRY-PICKING

Validity also has an adversarial side. A writer can choose sources. The Disclosure duty says that if the declared store holds a record contradicting a cited fact-constant, the output must mention it. For the Middle Kingdom, a writer citing Mark's date must disclose the encyclopedia's own conflicting timeline. For the emblem, a writer citing the hemp-leaf reading must disclose that its meaning is described elsewhere as unknown.

Disclosure is relative to a store, and that is deliberate. A duty to cite everything relevant in the world cannot be audited. A duty to disclose contradictions within a declared store can be: the auditor searches the store and checks. Declaring the store, as Appendix D does for this book, is therefore part of conforming.

OBJECTIONS AND LIMITS

A rule worth adopting has to survive its strongest objections. Here are nine, with the answer to each and the limits that remain.

1. “IT FORBIDS SYNTHESIS.”

A strict identity rule would forbid every summary, translation and argument, because each transforms its sources. Version 5.1 avoids this by scope. It governs fact-constants and leaves narrative free. A summary conforms if every checkable fact inside it is grounded and every conclusion is labelled. The earlier claim that paraphrasing systems are fundamentally non-intelligent^[34] went further than the rule needs, and Chapter 5 records its withdrawal.

2. “TRANSLATION IS UNAVOIDABLE.”

It is, and the axiom does not forbid it. It denies translation the status of a reproducible transformation, because two careful translators produce different text. A translated quotation therefore travels with its original wording, or is labelled as the author’s rendering. For a bilingual author working in English and French, the practical rule is simple: quote the original, then translate.

3. “SOME SOURCES CANNOT BE PUBLIC.”

Investigators protect informants; companies hold confidential data; courts seal records. Resolvability offers two routes. A **named auditor** can hold the record, which converts an anonymous assurance into a checkable claim about a person or institution. Or the author can publish a **hash commitment** now and reveal the record later.

Hash commitments need care. RFC 6920, the IETF standard for naming things with hashes, warns that publishing a hash name does not protect the confidentiality of the object: for a highly formatted object with little variation, an attacker can

guess candidate contents and check whether the digest matches^[29]. A short record, such as a name, a date or a yes-or-no finding, is exactly that kind of object.

RULE ADDED IN VERSION 5.1

A deferred record is committed as the hash of a secret random value (the salt) joined to the record. The salt is revealed together with the record. Until then, guessing the record does not reproduce the digest.

A deferred fact-constant is reported as *deferred*, not verified, until the reveal. That is a smaller promise than verification, and the output must say so.

4. “IT COSTS TOO MUCH.”

Most of what the axiom requires is cheap for a machine. Hashing a document, recording a version and keeping a snapshot are routine operations. The expensive part is human: the discipline of quoting exactly and labelling inferences. The cost of not doing it is visible in Chapter 2. In *Mata v. Avianca*, a single filing of fabricated citations produced an order to show cause, a sanctions hearing and a written opinion^[18].

5. “NOBODY CITES COMMON KNOWLEDGE.”

True, and the axiom does not demand it. A work may declare one version-pinned baseline reference that grounds general facts throughout. What it may not do is exempt facts from grounding altogether, because “common knowledge” is exactly where outdated and false beliefs hide. This book chose the stricter path and declared no baseline.

6. “AUTHORS WILL GAME THE CLASSIFICATION.”

An author who wants to avoid checking can call a fact an interpretation. The lookup test defeats this by classifying claims by what they do, not by what they are called, and the default resolves doubt against the author. A claim that a checker could look up is a fact-constant whatever its label.

7. “A FAITHFUL COPY OF A LIE IS STILL A LIE.”

Agreed. The non-claim says so in so many words, and Chapter 7 shows the two-layer record that handles it. The axiom’s contribution is to make the lie traceable to its source, which is where correction has to begin.

8. “WHO AUDITS THE AUDITOR?”

Conformance is transitive, so an auditor’s report is itself an output that can be audited. Named auditors carry reputations that can be checked. Hash commitments need no trusted auditor at all, only a later reveal. What the axiom cannot remove is the need for someone, eventually, to open the record and look.

9. “FAITHFUL MACHINES CAN STILL REASON BADLY.”

Yes. The Transparency corollary labels inferences; it does not validate them. Good reasoning needs methods of its own, such as the Analysis of Competing Hypotheses that Richards Heuer set out for intelligence analysts^[47] and that Chapter 10 builds into the investigative loop. The axiom’s role is narrower and prior: it makes sure the premises are what they claim to be before anyone reasons from them.

THE LIMITS THAT REMAIN

- The axiom cannot verify what was never recorded: an unrecorded conversation, a memory, a lost document.
- It cannot judge truth, only faithfulness to sources.
- The boundary between a checkable fact and an interpretation can be disputed in hard cases; the default rule settles disputes but does not dissolve them.
- Snapshots consume storage, and some sites forbid archiving.
- Deferred records offer a weaker guarantee than public ones, for as long as they stay deferred.

PART IV

PROTOCOLS

*Turning the axiom into procedures that people and machines can
follow, and numbers an auditor can report.*

THE LANDRY HALLUCINATION-FREE PROTOCOL

The axiom says what must be true of an output. The Landry Hallucination-Free Protocol (LHFP) says what to do, step by step, so that it is.

WHERE THE PROTOCOL CAME FROM

The LHFP first appeared in a white paper dated 10 January 2026^[31]. It proposed three mechanisms: a copy mode in which a model, when a fact is required, pulls the string from a verified database instead of generating it; a “surgical” patching interface that edits a single token in place rather than regenerating a whole document; and a symbolic checker that blocks output contradicting a knowledge graph^[31]. Its methodology section verified results by comparing the hash of the AI’s output against the source document^[31].

The copy mode has a published precedent. Pointer-generator networks, described in 2017, copy words from the source text by pointing, which their authors report aids accurate reproduction of information^[23]. The author’s Protocol v7.0 added the refusal rule: if no constant is retrieved, the system returns “Data Missing” rather than generating a value^[32].

WHAT THIS EDITION CHANGES

The white paper contains two figures that the protocol’s own rules cannot carry forward: a 27% hallucination rate for citations and a 15% drift in numerical data^[31]. Neither carries a source in the paper. Under Version 5.1 they are fact-constants without grounding, so this edition removes them pending a source. The paper’s hypothesis of 100% factual integrity with a 99% cost reduction is stated in its own methodology as a hypothesis^[31]; this edition keeps it in that status, as something to test rather than something to report. The WTF 10 engine refers to “LHFP v3”^[38]. The text below is the LHFP as conformed to Seshat’s Axiom 5.1.

THE PROTOCOL IN EIGHT STEPS

STEP	ACTION	CLAUSE SERVED
1. Capture	Fetch the source at a pinned version. Record URL or DOI, version, locator, capture time, frame and qualifiers. Archive a snapshot. Compute SHA-256.	Fidelity, Resolvability
2. Store	Append the record to the store. Never edit a record in place; corrections are new records that point to what they supersede.	Fidelity
3. Retrieve	Fetch by record ID. Recompute the hash and compare. A mismatch halts the pipeline.	Fidelity
4. Copy	Insert fact-constants by copying the stored span, not by generating it. Apply only declared transformations.	Grounding
5. Cite	Attach record ID, version and passage locator to every fact-constant. No record, no fact: return “Data Missing”.	Resolvability
6. Check	Compare every cited span with the stored record: exact match, or match after the declared transformation.	Grounding
7. Label and disclose	Mark inferences; search the declared store for contradicting records and report them.	Transparency, Disclosure
8. Report	Publish the conformance state of every fact-constant: verified, deferred or failed.	Conformance

REFERENCE SKETCH

```

capture(source):
    content = fetch(source.url, version=source.pin)
    record = { content, sha256: H(content), url: source.url,
              version: source.pin, locator, captured_at: now(),
              snapshot: archive(content), frame, qualifiers,
              as_of, tier }
    store.append(record)                # append-only; never edit

retrieve(id):
    r = store.get(id)
    assert H(r.content) == r.sha256    # Fidelity
    return r

cite(claim, id, span, transform=None):
    r = retrieve(id)
    assert resolvable(r)                # Resolvability
    assert matches(claim, r.content[span], transform) # Grounding
    return Citation(id, r.version, span, transform)

answer(question):
    facts = [cite(...) for each fact-constant] or "Data Missing"
    inferences = label(reasoning over facts) # Transparency
    conflicts = store.contradicting(facts)  # Disclosure
    return Output(facts, inferences, conflicts, conformance_report())

```

WHAT THE PROTOCOL DOES NOT DO

The LHFP does not make a model truthful. It makes a model's factual claims checkable, and it stops the model from stating a fact it cannot cite. Its name promises freedom from hallucination; its honest guarantee is narrower. Every fact-constant in a conforming output either traces to a stored record or is flagged. Hallucinated reasoning, as distinct from hallucinated facts, remains possible; the Transparency corollary makes it visible, and Chapter 10's methods are there to test it.

WTF AND THE INVESTIGATIVE LOOP

Investigations move in loops: observe, question, hypothesise, test, conclude, repeat. The axiom fits inside any such loop by attaching one duty to each step.

THE WICKED THINKING FRAMEWORK

The author’s Wicked Thinking Framework (WTF) developed through numbered versions. The synthesis of versions 1.0 to 9.0 describes WTF 6.0, “The New Scientific Method”, as a six-step recursive process: deep observation and problem definition; targeted questioning; hypothesis generation with Analysis of Competing Hypotheses; comprehensive experimentation with Admiralty source vetting; integrated critical analysis; and an actionable conclusion that can trigger the next cycle^[26].

Analysis of Competing Hypotheses comes from Richards J. Heuer Jr., whose *Psychology of Intelligence Analysis*, published by the CIA’s Center for the Study of Intelligence in 1999, devotes its eighth chapter to the method^[47]. Its discipline, evaluating evidence against several hypotheses at once rather than collecting support for a favourite, is the reasoning counterpart to the axiom’s discipline about records.

ONE DUTY PER STEP

WTF 6.0 STEP	WHAT THE AXIOM ADDS
1. Deep observation and problem definition	Capture every observation as a stored record with provenance; separate observed facts from first impressions (Transparency)
2. Targeted questioning	Each question names the records that could answer it, so that answers are resolvable
3. Hypothesis generation and ACH	Hypotheses are labelled inferences; the evidence matrix cites record IDs, not recollections

WTF 6.0 STEP	WHAT THE AXIOM ADDS
4. Experimentation and source vetting	New data are captured, hashed and pinned; Admiralty grades go in the validity layer, apart from conformance
5. Integrated critical analysis	Every fact used in the analysis is checked for Grounding; contradictions in the store are disclosed
6. Actionable conclusion	The report carries its conformance states; the conclusion is labelled as inference and may be recorded only as attributed (“concluded by X on date D”)

THE VERIFICATION AGENTS

The 9.0 synthesis describes WTF 7.0 and 8.0 as adding a Multi-Agent Verification System with four agents: a Historian that tracks citation drift and retractions, a Baseline Scout that hunts for omitted comparisons, a Verification Engine that checks citations against the full text of their sources, and a Calibration Tracker that compares stated confidence with historical accuracy^[26]. Each maps onto a part of Version 5.1.

AGENT	ITS TASK	AXIOM COUNTERPART
Historian	Citation drift, retraction status	Version pins and as-of dates (Fidelity, Grounding)
Baseline Scout	Missing comparisons	Disclosure duty
Verification Engine	Checking citations against full text	Grounding
Calibration Tracker	Confidence against accuracy	Validity layer (outside the axiom)

THE EARLIER INVESTIGATIVE FRAMEWORK

The Seshat Scientific Investigative Framework of September 2024, written for police investigations, sets out an eight-step chain-of-thought workflow: incident response; data collection and preliminary analysis; scene analysis and hypothesis formation; experimentation and evidence testing; critical thinking and systems analysis; data integration and synthesis; report structuring and presentation; and feedback, iteration and continuous learning^[2]. It emphasises a secure chain of custody from the moment evidence is collected^[2]. The axiom can be read as that chain of custody, generalised from physical evidence to every cited fact.

A NOTE ON THE 10-STEP LOOP

The author has also written an open research-method framework, the Scientific Investigative 10-Step Loop, licensed CC BY 4.0^[41]. Its full text was not among the records consulted for this edition, so this chapter does not describe its steps. Applying the duties above to its ten steps is left to the next edition, once the text is in the declared store.

MEASURING CONFORMANCE

An axiom you cannot measure is a slogan. This chapter defines the measurements, gives the audit procedure, and applies it to a real document.

STATES, NOT SCORES

The unit of measurement is the fact-constant, and each one ends in a state: verified, deferred or failed. Two further labels describe the auditor’s situation rather than the claim: *verified via secondary*, when the record was checked only through another source that quotes it, and *not checked*, when the auditor could not reach the record. Reporting states rather than a single score keeps the information that matters: which claims failed, and why.

FIVE MEASURES

MEASURE	DEFINITION	TARGET
Failed count	Number of fact-constants in the failed state	Zero for a strict pass
Grounding rate	Fact-constants matching their record exactly or under a declared transformation, divided by all fact-constants checked	1.0
Resolvability rate	Citations an independent auditor could follow to the exact passage, divided by all citations	1.0
Primary share	Fact-constants verified against a primary record, divided by all verified	Report it; no fixed target
Decay rate	Citations that resolved at publication but no longer resolve, measured at a stated later date	Report it; snapshots should hold it at zero

The grounding rate adapts the Fact Integrity Rate of the FC/NV architecture, defined there as facts preserved exactly divided by facts referenced, with a required value of 1.0^[25]. Version 5.1 counts a declared, reproducible transformation as preservation.

THE AUDIT PROCEDURE

- 1. Declare the store: list every record the output may cite.
- 2. Extract every fact-constant using the lookup test, including attributed and derived ones.
- 3. For each, follow the citation. Record whether it resolves to the exact passage.
- 4. Compare the claim with the passage: exact, declared transformation, or mismatch.
- 5. Check that qualifiers, frame and as-of date travelled with the claim.
- 6. Recompute every derived fact-constant from its stated method.
- 7. Search the declared store for contradicting records and check that they were disclosed.
- 8. Report each fact-constant’s state, the five measures, and what could not be checked.

A WORKED AUDIT

The corrected Technical Report on Seshat’s Axiom contains seven quotations attributed to outside sources^[35]. The table reports each one’s state after the checks made for this book on 8 October 2026.

QUOTATION ATTRIBUTED TO	FINDING	STATE
Bender et al. (2021), in the Key Judgment	Wording matches the paper as recorded by Quote Investigator; the quotation stops at a clause end without an ellipsis ^[17]	Verified via secondary
Bender et al. (2021), in the reference list	Opening reworded from “an LM is” to “Language models are” ^[17]	Failed
NIST FIPS 180-4	Calls SHA-256 a SHA-3 variant; FIPS 180-4 specifies the SHA-2 family, and SHA-3 is in FIPS 202 ^[27]	Failed
IETF RFC 6920	Quoted sentence not found in the RFC text; the abstract’s wording differs ^[29]	Failed
Lewis et al. (2020), definition of RAG	Two abstract sentences joined without an ellipsis; no version pinned ^[22]	Failed
Lewis et al. (2020), human evaluation figures	The body of the paper could not be opened during this check	Not checked

QUOTATION ATTRIBUTED TO	FINDING	STATE
Wittgenstein, <i>Tractatus</i>	The cited Project Gutenberg edition could not be opened during this check ^[48]	Not checked

Result: one of seven verified (via a secondary source), four failed, two not checked. The grounding rate over the five quotations that could be checked is 1 in 5, or 0.2. The report does not pass, and its section declaring self-compliance^[35] is itself an unsupported claim.

Both versions of the report state that they were generated with Gemini 2.5 Flash Preview, and that the model acted as a deterministic scribe under the 1=1=1 rule^[34, 35].

INFERENCE

The failed quotations read like reconstructions of their sources rather than copies of them, which is the behaviour the axiom exists to prevent. They appeared in a document produced by a model that had been told to act as a deterministic scribe. An instruction is not a mechanism. A model asked to copy can still compose, and only a procedure that copies from the stored record and then checks the result, as Chapter 9 sets out, turns the request into a guarantee.

THIS BOOK’S OWN AUDIT

Appendix D applies the same states to every source this book cites. The declared store holds 59 records. 37 were verified directly or recomputed; 6 were verified only through a secondary source; 2 are used only as attributed claims; 12 are held in the author’s restricted archive and were read in full; 1 could not be checked; and 1 failed and appears only as a case of failure. Web sources are pinned by version where one exists and otherwise by access date; snapshots are to be archived with the Zenodo deposit. No source in the failed state is used to support a claim; failed sources appear only as cases of failure.

PART V

APPLICATIONS

*The axiom at work in machines, investigations, publishing and
law.*

AI AND RETRIEVAL SYSTEMS

Retrieval gives a model the right documents. The axiom makes it prove, claim by claim, that it used them faithfully.

A CONFORMING RETRIEVAL PIPELINE

Retrieval-augmented generation already has the shape the axiom needs: a store, a retriever and a generator^[22]. What it lacks by default is the bookkeeping that would let a stranger check its output. Six additions supply it.

1. **Hash every chunk at indexing time** and keep its source, version and locator beside the vector. The index becomes a store in the axiom's sense.
2. **Pin the corpus.** An answer cites corpus version N; re-indexing creates version N+1 and leaves N retrievable.
3. **Return records, not just text.** The retriever passes record IDs and hashes to the generator, and the generator's output carries them through.
4. **Copy, then check.** Fact-constants are inserted by copying spans; a post-generation checker compares every cited span with its record and fails the answer on any mismatch.
5. **Refuse when empty.** With no supporting record, the system states that the data are missing rather than generating a value^[32].
6. **Emit a conformance report** with every answer: each fact-constant, its record, its state.

WHAT AUDITORS AND REGULATORS WILL ASK FOR

Regulation is moving toward the same records. Article 12 of the EU Artificial Intelligence Act requires that high-risk AI systems technically allow the automatic recording of events, or logs, over the system's lifetime^[49]. As of 8 October 2026, the AI Act Explorer reports that the article applies from 2 December 2027 for high-risk systems listed in Annex III and from 2 August 2028 for those under An-

nex I^[49]. In the United States, NIST’s AI Risk Management Framework 1.0, published in January 2023, sets out characteristics of trustworthy AI; it treats “valid and reliable” as the base on which the others rest, and “accountable and transparent” as relating to all of them^[50].

INFERENCE
A conformance report is a log of exactly the kind these frameworks value: for every factual output, the record it came from and whether it matched. A system built to the axiom generates evidence of validity and transparency as a side effect of normal operation, rather than reconstructing it after an incident. This is an observation about fit, not a statement that conformance satisfies any legal requirement.

A SYSTEM CARD FOR CONFORMANCE

FIELD	EXAMPLE ENTRY
Declared store	Corpus name, version, record count, hash of the manifest
Capture policy	Snapshot service, hash algorithm (SHA-256), version pinning rule
Allowed transformations	Units, rounding rules, date formats
Refusal rule	Behaviour when no record supports a requested fact
Grounding checker	Matching method; behaviour on mismatch
Disclosure search	How contradicting records are found and reported
Measured performance	The five measures of Chapter 11 on a stated test set and date

OSINT CHAIN OF CUSTODY

Open-source investigators already treat a screenshot as evidence only if they can show where it came from and that it has not changed. The axiom generalises that discipline to every fact in the final report.

THE BERKELEY PROTOCOL

The Berkeley Protocol on Digital Open Source Investigations, published by the UN Office of the High Commissioner for Human Rights and the Human Rights Center at the University of California, Berkeley, sets international standards for online research into alleged violations of international criminal, human rights and humanitarian law, including methods for gathering, analysing and preserving digital information^[51]. It was launched on 1 December 2020, and it addresses the identification, collection, preservation, verification and analysis of digital open-source information^[52].

BERKELEY PROTOCOL PROCESS	AXIOM COUNTERPART
Identification	Declaring the store and the locators
Collection	Capture with provenance (Fidelity)
Preservation	Hash and snapshot (Fidelity, Resolvability)
Verification	Source validity (outside the axiom) and Grounding
Analysis	Labelled inference (Transparency) and Disclosure

HASH AT THE MOMENT OF CAPTURE

The single most valuable habit is to hash a capture the moment it is made. SHA-256, from the SHA-2 family specified in FIPS 180-4^[27], is the default. RFC 6920 standardises a way to name an object by its hash, the `ni` URI, and makes SHA-256 mandatory to implement^[29]. A hash-named capture can be shared, mirrored and re-hosted while remaining checkable: whoever holds the bytes can recompute the name. The example in Chapter 4, the Named Information form of

“Hello World!”, was recomputed for this book and matches the RFC’s printed example character for character^[28, 29].

KEEPING ANALYSIS APART FROM EVIDENCE

ICD 203’s third tradecraft standard requires analysts to distinguish underlying intelligence from their own assumptions and judgments^[44]. The Admiralty Code grades each source and each item separately^[42]. Together with the axiom, they give an OSINT report three layers that must never blur: what the record says (conformance), how far to believe it (validity) and what the analyst concludes (labelled inference).

PROTECTING SOURCES WITHOUT LOSING CHECKABILITY

An investigator cannot publish an informant’s message, but can publish a salted commitment to it (Chapter 8). The report then carries the fact-constant in the deferred state, with a promise that a named court, editor or ombudsman can be shown the record and salt. The reader learns exactly how much to trust the claim: not verified, but fixed in advance and open to later check.

CASE · THE PAGE THAT VANISHED

On 8 October 2026, a search for this book returned the Wikipedia talk page for Seshat, including a discussion of the hemp-leaf reading of her emblem^[10]. Minutes later, a direct request for the page failed^[10]. Nothing sinister happened; a page simply became unavailable to the tools in use. But an investigator who had cited it without a snapshot would now hold a citation that nobody could follow. Under Version 5.1, the capture would have included the snapshot and the hash, and the citation would still resolve.

OPEN SCIENCE AND PUBLISHING

*Scholarly publishing already has most of the machinery the axiom needs.
The work is in using it precisely.*

IDENTIFIERS THAT PERSIST

The Digital Object Identifier system is an international standard, ISO 26324, which the DOI Foundation's fact sheet dates to 1 May 2012; it is governed by the International DOI Foundation, which the same fact sheet says was launched in 1998, and it resolves through the Handle System^[53]. The Foundation's own handbook gives different dates: it says the Foundation was created in 1997 and the standard approved in 2010^[54]. Approval and publication can fall in different years, so the two records may both be right; this book reports both rather than choose, as the Disclosure duty requires. When a resource moves, its DOI continues to resolve to the new location^[55]. The ORCID iD does the same for people: it is a persistent digital identifier that distinguishes one researcher from every other^[56].

CITE THE VERSION, NOT THE CONCEPT

Zenodo's documentation explains that a new version of a record is a completely new record, with its own metadata, files and persistent identifier, linked to all previous and future versions; citing a specific version ensures that its files have not changed^[57]. For the axiom this is decisive. A citation to the version DOI satisfies Resolvability and protects Fidelity. A citation that does not name a version leaves open which files the reader will find, and invites drift.

METADATA ARE FACT-CONSTANTS

Repository metadata can contradict themselves, and when they do the axiom treats each field as a claim to reconcile. The Zenodo record for the Hempoxy master document describes the work as released into the public domain as prior art,

while its licence field reads Creative Commons Attribution 4.0 International^[1]. Public-domain dedication and CC BY are different grants. One of the two fields is wrong, and a reader cannot tell which from the record alone.

The same record lists an MD5 checksum for its file^[1]. A published checksum is exactly what Fidelity needs. MD5, however, is the kind of weak hash function that RFC 6920 tells its registry experts not to accept because of collision weakness^[29]. Depositing a SHA-256 digest alongside the repository's checksum costs nothing and closes the gap.

ONE AUTHOR, SEVERAL NAMES

This book is published under the name Marie-Soleil Seshat Landry. The Hempoxy record on Zenodo lists its creator as Landry, Marie Seshat, linked to ORCID 0009-0008-5027-3337^[1], and the 2024 investigative framework gives the author as Marie Seshat Landry^[2]. Names vary; the ORCID iD does not. Under the axiom, the identifier is the fact-constant and the name is its rendering.

Identifier bindings need the same care. A research aggregator lists a related record by the author, *Towards AGI – Organic Deterministic Intelligence*, under DOI 10.5281/zenodo.19015971^[58]. That binding was not checked against Zenodo itself for this edition, so it is reported here as an attributed claim, and Chapter 6 shows why the check matters.

PRIOR ART AND THE DEPOSIT PLAN

The Prior Art Release of 26 March 2026 declared the axiom into the public domain as open-source prior art and stated that it established priority as of that date^[33]. A priority claim is only as strong as the evidence of the date and of the exact text. A file in a private drive establishes neither to an outsider. A timestamped repository deposit, with a published SHA-256 digest of the file, establishes both.

DECISIONS FOR THIS EDITION

- Deposit this book on Zenodo under CC BY 4.0, with a SHA-256 digest of the PDF in the description.
- Deposit the archive records listed in Chapter 5, unchanged and errors included, as a companion dataset titled *Sesbat's Axiom: The Version Record*, so that every restricted citation in Appendix D becomes public.
- Link the two deposits with related-identifier fields, and cite version DOIs throughout.
- Correct the Hempoxy record's licence statement so that description and licence field agree.

LAW AND COMPLIANCE

Courts and regulators already demand that cited authority exist and say what it is cited for. The axiom states that demand precisely enough to automate it.

GATEKEEPERS

In sanctioning the lawyers in *Mata v. Avianca*, Judge Castel wrote that existing rules “impose a gatekeeping role on attorneys to ensure the accuracy of their filings”^[59]. The sanctions turned on more than the first false citations: the opinion found that the lawyers continued to stand by the fake opinions after judicial orders called their existence into question^[18]. The axiom adds nothing to that duty. It describes how to discharge it: resolve every citation to a pinned record, check every quotation against the record, and when a citation is challenged, open the record before answering.

THE RECORD MUST OUTLIVE THE FILING

The 2014 study of link rot found half the URLs in Supreme Court opinions no longer led to the information cited, and its authors proposed long-term preservation of link contents by libraries^[21]. That proposal is the snapshot rule applied to the law. A legal citation to a web page without a preserved copy fails Resolvability on a delay.

COMPLIANCE ENGINES

A compliance engine checks a submission, such as a building-permit application, against a body of rules. The author’s DICE project, a Deterministic AI Compliance Engine for building-permit compliance, is one such design^[41]. The axiom suggests four requirements for any engine of this kind.

1. **Rules are stored records.** Each rule is captured with its code, edition, section number and effective date, and hashed.

2. **Findings cite clauses.** Every finding of compliance or non-compliance names the rule record, edition and section it rests on.
3. **Amendments are disclosed.** If the store holds a later amendment or a conflicting local bylaw, the finding says so.
4. **The non-claim is stated.** A finding certifies conformity with the cited edition of the cited rules. It does not certify that the building is safe, and the report must say so.

LOGS AS EVIDENCE

Where the EU AI Act's record-keeping article applies, providers of high-risk systems must make automatic logging technically possible^[49]. A compliance engine built to the axiom produces, for every finding, the record it relied on and whether it matched. Whether such a log satisfies any particular legal requirement is a question for counsel in the relevant jurisdiction; this chapter claims only that the axiom produces the raw material.

Nothing in this chapter is legal advice.

CONCLUSION

A STANDARD FOR VERIFIABLE MACHINES

The axiom promises one thing and refuses to promise another. Both halves matter.

WHAT A CONFORMING OUTPUT GUARANTEES

- Every fact-constant traces to a specific version of a specific record, at an exact passage.
- The record retrieved is bit-for-bit the record stored, and the record stored is the record published.
- Every fact-constant matches its record exactly, or under a transformation the output names and anyone can repeat.
- Qualifiers, frames and as-of dates travel with the facts they qualify.
- Every inference is marked as one, and rests on grounded facts.
- Contradicting records in the declared store are disclosed.

WHAT IT DOES NOT GUARANTEE

- That the sources are true.
- That the reasoning is sound.
- That the declared store contains every relevant record in the world.

THE RULE IN ONE LINE

WHAT IS CITED IS WHAT WAS RETRIEVED;
WHAT WAS RETRIEVED IS WHAT WAS STORED.

The rule is small enough to state on a cover and strict enough to audit by machine. It does not replace judgment about what is true. It clears the ground for

that judgment, by making sure that when people argue about a fact, they are arguing about the same fact.

This book has tried to obey the rule it states. Appendix D reports where it succeeded and where it could not. Readers are invited to audit it, to find what the audit missed, and to publish what they find. Under the axiom, a correction is not an embarrassment. It is a new version, pinned and kept, with the old one still on the record.

• • •

THE CANONICAL TEXT, VERSION 5.1

This is the text to cite. Cite it as: Landry, M.-S. S. (2026). Seshat's Axiom, Version 5.1, with the version DOI of this edition.

SESHAT'S AXIOM · VERSION 5.1 · 8 OCTOBER 2026

WHAT IS CITED IS WHAT WAS RETRIEVED; WHAT WAS RETRIEVED IS WHAT WAS STORED. (I = I = I)

1. **Fidelity (Store = Retrieve).** A record is captured with its provenance: source, exact locator, version, capture time, a content hash and an archived snapshot. Anything later retrieved is identical to what was stored.
2. **Grounding (Retrieve = Cite).** Every fact-constant in an output matches its cited record exactly, or under a declared, reproducible transformation within a declared reference frame. The record's qualifiers and its "as of" date travel with it.
3. **Resolvability (precondition).** Every citation is version-pinned and locates the exact passage. It is public, held by a named auditor, or committed by a salted hash for later disclosure.

Corollary, Transparency. Every inference is labelled as inference, and every chain of inference ends in grounded fact-constants.

Duty, Disclosure. If the declared store holds a record that contradicts a cited fact-constant, the output says so.

Non-claim. The axiom guarantees faithfulness to sources, not their truth.

DEFINITIONS

FACT-CONSTANT

An assertion a checker can confirm by looking it up in the record it cites. When in doubt, an assertion is a fact-constant; its class is set by the lookup test, not by the author's label.

DERIVED FACT-CONSTANT

A result of reproducible computation from grounded fact-constants, with the method shown.

ATTRIBUTED CLAIM

An assertion of the form "Source S states P". It is a fact-constant about S, not about P.

INFERENCE

A conclusion, interpretation or causal claim of the author. It is never promoted to a fact-constant; it may be recorded only as attributed ("concluded by X on date D").

NARRATIVE-VARIABLE

Wording, order, tone and structure that assert nothing checkable.

DECLARED TRANSFORMATION

A change that two careful people, applying the stated rule, would carry out identically: unit conversion, stated rounding, formatting. Translation and paraphrase are not declared transformations.

REFERENCE FRAME

The system within which a value has meaning: chronology, unit system, calendar, coordinate datum, currency and base year.

DECLARED STORE

The list of records an output may cite and against which Disclosure is checked.

SALTED HASH COMMITMENT

The hash of a secret random value joined to a record, published before the record is revealed.

CONFORMANCE

An output conforms when an independent auditor, following its citations, can verify every fact-constant. Each fact-constant is verified, deferred or failed. Conformance is transitive: citing a conforming output grounds the fact; citing a non-conforming one grounds only the attributed claim. A citation cycle grounds nothing. One declared, version-pinned baseline reference may ground general knowledge for a whole work.

CHANGELOG

VERSION	DATE	CHANGE
5.1	8 Oct. 2026	Hash commitments must be salted (after RFC 6920, section 10). “1=1=1=1” retired as a name. Record fields from FC/NV v2.0 and WTF 7.0 adopted as required fields.
5.0	8 Oct. 2026	Three clauses, corollary, Disclosure duty, non-claim; reference frames; mandatory snapshots; metadata as fact-constants.
Earlier	Sept. 2024 – Mar. 2026	See Chapter 5. All earlier versions remain on the record.

CONFORMANCE CHECKLISTS

Two lists: one for the person writing, one for the person checking.

FOR AUTHORS AND SYSTEMS

- ☐ I declared the store: every record this output may cite.
- ☐ Every record was captured with source, version, locator, capture time, hash and snapshot.
- ☐ Every fact-constant was copied from its record, not written from memory.
- ☐ Every quotation is exact, and every omission inside it is marked with an ellipsis.
- ☐ Every changed value names its transformation; every date and number names its frame.
- ☐ Qualifiers and as-of dates travel with their facts.
- ☐ Every citation names a version and an exact passage.
- ☐ Restricted sources name an auditor or carry a salted commitment.
- ☐ Every derived figure shows its method.
- ☐ Every inference is labelled, and rests on grounded facts.
- ☐ I searched the store for contradicting records and disclosed them.
- ☐ I published the conformance state of every fact-constant.

FOR AUDITORS

- ☐ Extract every fact-constant with the lookup test, regardless of how it is labelled.
- ☐ Follow each citation; record whether it reaches the exact passage at the pinned version.
- ☐ Recompute hashes where the record provides them.
- ☐ Compare claim and passage: exact, declared transformation, or mismatch.
- ☐ Check frames, qualifiers and as-of dates.
- ☐ Recompute derived figures.
- ☐ Search the declared store for undisclosed contradictions.
- ☐ Report states, the five measures of Chapter 11, and everything you could not check.

GLOSSARY

ADMIRALTY CODE

The NATO system that grades source reliability A–F and information credibility 1–6, separately.

ACH (ANALYSIS OF COMPETING HYPOTHESES)

Heuer’s method of testing evidence against several hypotheses at once.

CONCEPT AND VERSION IDENTIFIERS

On Zenodo, each version of a record has its own persistent identifier; citing a version pins its files.

CONFORMANCE

The state of an output whose every fact-constant an independent auditor can verify through its citations.

DECAY

Loss of resolvability over time, as when a cited page disappears.

DOI

Digital Object Identifier, a persistent identifier standardised as ISO 26324.

FC/NV

Fact-Constant/Narrative-Variable: the author’s architecture separating immutable facts from fallible narrative.

FIDELITY, GROUNDING, RESOLVABILITY

The three clauses of the axiom (Chapter 3).

HASH (SHA-256)

A fixed-length digest of content; any change to the content changes the digest.

LHFP

Landry Hallucination-Free Protocol: the step-by-step procedure that implements the axiom (Chapter 9).

NI URI

A Named Information identifier from RFC 6920, naming an object by its hash.

ORCID ID

A persistent identifier for a researcher.

PRIMARY SOURCE

In PROV-DM, material produced by an agent with direct experience of a topic, at the time.

PROVENANCE

In PROV-DM, a record of the people, institutions, entities and activities involved in producing, influencing or delivering data.

SNAPSHOT

An archived copy of a source taken at capture time.

VALIDITY LAYER

The assessment of how far a record should be believed, kept separate from conformance (Chapter 7).

WTF

The author's Wicked Thinking Framework, a family of truth-finding methods developed through numbered versions.

SOURCES AND CONFORMANCE REGISTER

Every record this book cites, with its version pin, access tier and conformance status as of 8 October 2026. This list is the book’s declared store.

The declared store holds 59 records. 37 were verified directly or recomputed; 6 were verified only through a secondary source; 2 are used only as attributed claims; 12 are held in the author’s restricted archive and were read in full; 1 could not be checked; and 1 failed and appears only as a case of failure. Web sources are pinned by version where one exists and otherwise by access date; snapshots are to be archived with the Zenodo deposit.

STATUS	MEANING
Verified	Opened and checked directly on 8 October 2026, or recomputed for this book
Via secondary	Checked only through another source that quotes or describes it
Attributed	Used only for what the source says, not for the truth of what it says
Restricted	Author’s private archive, read directly; named holder: the author; to be deposited publicly
Not checked	Could not be opened; cited only to report that fact
Failed	Did not resolve or did not match; cited only as a case of failure

NO.	RECORD	PIN · TIER · HOW CHECKED	STATUS
1	Landry, Marie Seshat. “Hemphox – A Master Document on the Science, Synthesis and Strategy of a Sustainable Nanocomposite System.” Zenodo, 25 Aug. 2025. DOI 10.5281/zenodo.16944339. https://zenodo.org/records/16944339	Version v1; 151 views and 90 downloads as of 8 Oct. 2026 Public · Opened in full	Verified
2	Landry, Marie Seshat. Seshat Scientific Investigative Framework (SSIF). Sept. 2024. Google Drive file 1ga81InzVMSbur8w2LMRA3Kn7xmG1oLJl	Dated Sept. 2024 Restricted · Read	Restricted

NO.	RECORD	PIN · TIER · HOW CHECKED	STATUS
3	Encyclopaedia Britannica. "Seshat." Britannica. https://www.britannica.com/topic/Seshat	No revision date shown; Accessed 8 Oct. 2026 Public · Search excerpt	Verified
4	Mark, Joshua J. "Seshat." World History Encyclopedia, 18 Nov. 2016. https://www.worldhistory.org/Seshat/	Published 18 Nov. 2016; Accessed 8 Oct. 2026 Public · Opened in full <i>Used only for what the article states; contains a mislabelled Middle Kingdom date.</i>	Attributed
5	Wilkinson, Richard H. Symbol & Magic in Egyptian Art. Thames & Hudson, 1994, p. 174.	Page 174 as quoted in Mark 2016 Public (print) · Through Mark 2016 <i>Primary copy not consulted.</i>	Via secondary
6	Wikipedia contributors. "Seshat." English Wikipedia. https://en.wikipedia.org/wiki/Seshat	Revision ID not recorded; Accessed 8 Oct. 2026 Public · Search excerpt <i>Pin to a revision ID at deposit.</i>	Verified
7	David, Rosalie. Religion and Magic in Ancient Egypt. Penguin, 2003, p. 411.	Page 411 as quoted in Mark 2016 Public (print) · Through Mark 2016 <i>Primary copy not consulted.</i>	Via secondary
8	Wilkinson, Richard H. The Complete Temples of Ancient Egypt. Thames & Hudson, 2000, p. 38.	Page 38 as quoted in Mark 2016 Public (print) · Through Mark 2016 <i>Primary copy not consulted.</i>	Via secondary
9	Thurston, Jacqueline. "Sacred Deities of Ancient Egypt." ARAS (Archive for Research in Archetypal Symbolism). https://aras.org/node/632125	Accessed 8 Oct. 2026 Public · Search excerpt	Verified
10	Wikipedia contributors. "Talk:Seshat." English Wikipedia. https://en.wikipedia.org/wiki/Talk:Seshat	Seen in a search result on 8 Oct. 2026; direct request failed minutes later Public · Search excerpt only <i>Decayed for the tools in use; no snapshot. Cited only as a case.</i>	Failed
11	Seshat: Global History Databank. "Seshat Data" (downloads page). https://seshat-db.com/downloads_page	Accessed 8 Oct. 2026 Public · Search excerpt	Verified
12	"Seshat: A tool for managing and verifying annotation campaigns of audio data." arXiv:2003.01472. https://arxiv.org/pdf/2003.01472	Version not recorded; Accessed 8 Oct. 2026 Public · Search excerpt	Verified

NO.	RECORD	PIN · TIER · HOW CHECKED	STATUS
13	Kim, Jinhan, et al. “Predictive Mutation Analysis via Natural Language Channel in Source Code.” arXiv: 2104.10865. https://arxiv.org/pdf/2104.10865	Version not recorded; Accessed 8 Oct. 2026 Public · Search excerpt	Verified
14	World History Encyclopedia. “First Intermediate Period of Egypt” (timeline). https://worldhistory.org/timeline/First_Intermediate_Period_of_Egypt	Undated page; Accessed 8 Oct. 2026 Public · Search excerpt <i>Chronology: World History Encyclopedia.</i>	Verified
15	Ji, Ziwei, et al. “Survey of Hallucination in Natural Language Generation.” DOI 10.1145/3571730. arXiv: 2202.03629. https://arxiv.org/abs/2202.03629	arXiv v7, 14 Jul. 2024 Public · Opened (abstract page)	Verified
16	Bender, Emily M., Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. “On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?” FAccT ’21, 2021. DOI 10.1145/3442188.3445922. https://doi.org/10.1145/3442188.3445922	2021 proceedings Public · Wording via Quote Investigator; DOI= title binding seen in search listings <i>Primary PDF not opened.</i>	Via secondary
17	Quote Investigator. “Stitching Together Sequences of Linguistic Forms ... A Stochastic Parrot.” 7 Oct. 2022. https://quoteinvestigator.com/2022/10/07/word-parrot/	Published 7 Oct. 2022; Accessed 8 Oct. 2026 Public · Search excerpt	Verified
18	Mata v. Avianca, Inc., 678 F. Supp. 3d 443 (S.D.N.Y. 2023), Opinion and Order on Sanctions, signed 22 June 2023 (Castel, J.). https://www.courts.nj.gov/PDF/Mata%20v.%20Avianca%20Inc.pdf	Signed 22 June 2023; Accessed 8 Oct. 2026 Public · Search excerpt of the opinion	Verified
19	Maruf, Ramishah. “Lawyer apologizes for fake court citations from ChatGPT.” CNN (syndicated by KRDO). https://krdo.com/?p=972183	Accessed 8 Oct. 2026 Public · Search excerpt	Verified
20	Louisiana Law Blog. Report on the Mata v. Avianca sanctions opinion. https://www.louisianalawblog.com/?p=7151	Accessed 8 Oct. 2026 Public · Search excerpt	Verified
21	Zittrain, Jonathan, Kendra Albert, and Lawrence Lessig. “Perma: Scoping and Addressing the Problem of Link and Reference Rot in Legal Citations.” 127 Harvard Law Review Forum 176 (2014). https://hls.harvard.edu/bibliography/perma-scoping-and-addressing-the-problem-of-link-and-reference-rot-in-legal-citations	2014; Accessed 8 Oct. 2026 Public · Search excerpt (abstract)	Verified
22	Lewis, Patrick, et al. “Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks.” arXiv:2005.11401. https://arxiv.org/abs/2005.11401	arXiv v4, 12 Apr. 2021 Public · Opened (abstract page) <i>Body text (human evaluation) not opened.</i>	Verified

NO.	RECORD	PIN · TIER · HOW CHECKED	STATUS
23	See, Abigail, Peter J. Liu, and Christopher D. Manning. “Get To The Point: Summarization with Pointer-Generator Networks.” arXiv:1704.04368. https://arxiv.org/abs/1704.04368	arXiv v2, 25 Apr. 2017 Public · Opened (abstract page)	Verified
24	National Institute of Standards and Technology. FIPS 180-4: Secure Hash Standard (SHS). Aug. 2015. DOI 10.6028/NIST.FIPS.180-4. https://csrc.nist.gov/pubs/fips/180-4/upd1/final	Aug. 2015 edition Public · Opened (publication page)	Verified
25	Landry, Marie-Soleil Seshat. Fact-Constant / Narrative-Variable Architecture, Version 2.0. Jan. 2026. Google Drive file 1eSLBAA4W9NPLF4I0cLWwa7nPXNIEGgh-	Last updated Jan. 2026 Restricted · Read	Restricted
26	Landry, Marie-Soleil Seshat. Synthesizing WTF Frameworks into 9.0. Google Drive file 1MTgUC2ITzA_mSuqdeus5Q2zrzwpqELv1	Undated Restricted · Read	Restricted
27	National Institute of Standards and Technology. FIPS 202: SHA-3 Standard. Aug. 2015. DOI 10.6028/NIST.FIPS.202. https://csrc.nist.gov/pubs/fips/202/final	Aug. 2015 edition Public · Search excerpt	Verified
28	Computation: SHA-256 of the 12-byte ASCII string “Hello World!”, encoded base64url without padding (Python 3, hashlib). Result: f40xZX_x_F05LcGBSKHWXfwtSx-jlnc0St3SABJtkGk	Recomputed 8 Oct. 2026; matches RFC 6920 §8.1 Derived · Recomputed	Verified
29	Farrell, S., et al. RFC 6920: Naming Things with Hashes. IETF, Apr. 2013. Proposed Standard. https://datatracker.ietf.org/doc/html/rfc6920	RFC 6920 (immutable) Public · Opened in full	Verified
30	Sabbahy, Lisa K. “Background to Ancient Egypt.” In <i>Kingship, Power, and Legitimacy in Ancient Egypt</i> , Cambridge University Press, 2020, pp. 4–28 (citing Shaw 2004, p. 184). https://www.cambridge.org/core/books/kingship-power-and-legitimacy-in-ancient-egypt/background-to-ancient-egypt/D58485329A866F2C110007FF07E0A337	Published online 26 Nov. 2020; Accessed 8 Oct. 2026 Public (summary); full text paywalled · Search excerpt <i>Chronology: Shaw 2004.</i>	Verified
31	Landry, Marie-Soleil Seshat. White Paper: The Landry Hallucination-Free Protocol (LHFP). 10 Jan. 2026. Google Drive file 1Nm9uLM6G0ByYf1936dvy9_9pVvuHwpuM	Dated 10 Jan. 2026 Restricted · Read	Restricted
32	Landry, Marie-Soleil Seshat. Deterministic Retrieval and Constrained Inference, Protocol Version 7.0. Google Drive file 1uEmr9f9zIUYQZr4AsXw6EVBPJLmckNng	Undated; archived 27 Jan. 2026 Restricted · Read	Restricted

NO.	RECORD	PIN · TIER · HOW CHECKED	STATUS
33	Landry, Marie-Soleil Seshat. Seshat's Axiom (1=1=1) Prior Art Release. 26 Mar. 2026. Google Drive file 15TxvGRx1GzeKZheTKGdqid5RJhGfp87A	Dated 26 Mar. 2026 Restricted · Read	Restricted
34	Landry, Marie-Soleil Seshat. Technical Report: Seshat's Axiom (uncorrected). Google Drive file 1_HxYLpaYHG16guRV1iWCGufE4oJAVD7m	Undated Restricted · Read	Restricted
35	Landry, Marie-Soleil Seshat. Technical Report: Seshat's Axiom (corrected). Mar. 2026. Google Drive file 1hHSXYViRi2ZKhe4en2P5KjSsegMxmyR6	Published Mar. 2026 Restricted · Read	Restricted
36	Landry, Marie-Soleil Seshat. "seshats axiom.pdf" (1 = 1 = 1 Law of Verified Truth). Google Drive file 19rAJozWuZdzbB_-A9axje8jtEzCwBUZh	Undated Restricted · Read	Restricted
37	Landry, Marie-Soleil Seshat. Seshat's Axiom — Book Foundation (working document, Version 5). 8 Oct. 2026. Claude Docs artifact fd7f4eb6-0f50-424c-aef8-6be58429f379	8 Oct. 2026 Restricted · Read	Restricted
38	Landry, Marie-Soleil Seshat. WTF 10 Framework Engine. Google Drive file 13SZ0FbSw1GATvqMZQh07oJNGr14qZW1F	Undated Restricted · Read	Restricted
39	Computation: character-level Shannon entropy and SHA-256 digests of "2040 BCE", "2400 BCE", and of the Bender et al. excerpt in its original and reworded forms (Python 3). Results: 2.75 bits/char for both dates; digests ba3a5f37... vs 7a046b23...; 1ea1ce7a... vs 194b3444...	Recomputed 8 Oct. 2026 Derived · Recomputed	Verified
40	Computation: count of URLs in the uncorrected Technical Report that point to a Google search page. Result: 7 of 25 listed URLs (22 unique)	Recomputed 8 Oct. 2026 Derived · Recomputed	Verified
41	Landry, Marie-Soleil Seshat. Author's notes on her frameworks and projects (10-Step Loop; DICE). Author's private notes	As held 8 Oct. 2026 Restricted · Read <i>Author's own description.</i>	Restricted
42	Wikipedia contributors. "Admiralty code." English Wikipedia. https://en.wikipedia.org/wiki/Admiralty_code	Revision ID not recorded; Accessed 8 Oct. 2026 Public · Search excerpt	Verified
43	Dragos, Valentina. Definitions of reliability and credibility according to STANAG 2511 (Jan. 2003). ETURWG, George Mason University. https://eturwg.c4i.gmu.edu/?q=node/128	Accessed 8 Oct. 2026 Public · Search excerpt <i>Reproduces STANAG 2511 definitions.</i>	Via secondary

NO.	RECORD	PIN · TIER · HOW CHECKED	STATUS
44	Office of the Director of National Intelligence. Intelligence Community Directive 203: Analytic Standards. 2015. https://irp.fas.org/dni/icd/icd-203.pdf	2015 revision Public · Search excerpt of the PDF	Verified
45	Intelligence.gov. “Objectivity” (IC Analytic Standards). https://www.intelligence.gov/mission/our-values/objectivity	Accessed 8 Oct. 2026 Public · Search excerpt	Verified
46	Moreau, Luc, and Paolo Missier, eds. PROV-DM: The PROV Data Model. W3C Recommendation, 30 Apr. 2013. https://www.w3.org/TR/2013/REC-prov-dm-20130430/	REC-prov-dm-20130430 Public · Opened in full	Verified
47	Heuer, Richards J., Jr. Psychology of Intelligence Analysis. Center for the Study of Intelligence, CIA, 1999. Ch. 8, p. 95. https://www.cia.gov/resources/csi/static/Pyschology-of-Intelligence-Analysis.pdf	1999 edition Public · Search excerpt (contents)	Verified
48	Wittgenstein, Ludwig. Tractatus Logico-Philosophicus. Project Gutenberg eBook 5740. https://www.gutenberg.org/ebooks/5740	Not opened Public · Request not completed <i>Cited only in the audit of Chapter 11.</i>	Not checked
49	EU AI Act Explorer. “Article 12: Record-Keeping,” EU Artificial Intelligence Act. https://artificialintelligenceact.eu/article/12/	Application dates as displayed on 8 Oct. 2026 Public · Opened in full <i>Secondary presentation of the official text; dates may change.</i>	Verified
50	National Institute of Standards and Technology. NIST AI 100-1: Artificial Intelligence Risk Management Framework (AI RMF 1.0). Jan. 2023. DOI 10.6028/NIST.AI.100-1. https://nvlpubs.nist.gov/nistpubs/ai/nist.ai.100-1.pdf	Jan. 2023 edition Public · Search excerpt of the PDF	Verified
51	UN OHCHR and Human Rights Center, UC Berkeley. Berkeley Protocol on Digital Open Source Investigations. HR/PUB/20/2, 2022. https://digitallibrary.un.org/record/3973652	HR/PUB/20/2 (2022) Public · Search excerpt (catalogue record)	Verified
52	American Society of International Law. “Strengthening Atrocity Crime Cases with Digital Evidence: New International Standards on Digital Open Source Investigations.” https://asil.org/?p=15645	Accessed 8 Oct. 2026 Public · Search excerpt	Verified
53	DOI Foundation. “Key Facts on Digital Object Identifier System.” https://www.doi.org/factsheets/DOIKeyFacts.pdf	Accessed 8 Oct. 2026 Public · Search excerpt	Verified

NO.	RECORD	PIN · TIER · HOW CHECKED	STATUS
54	International DOI Foundation. DOI Handbook, §1.5 “Introduction (ISO 26324).” Last updated 17 Oct. 2015. https://albany.doi.org/doi_handbook/1_Introduction.html	Last updated 17 Oct. 2015 Public · Search excerpt <i>Dates differ from the fact sheet (record 1997/2010 vs 1998/2012).</i>	Verified
55	International DOI Foundation. “The DOI System ISO 26324” (home page). https://albany.doi.org	Accessed 8 Oct. 2026 Public · Search excerpt	Verified
56	University of Illinois Chicago Library. “ORCID & Author IDs: About ORCID” (quoting ORCID.org). https://researchguides.uic.edu/orcid	Accessed 8 Oct. 2026 Public · Search excerpt <i>Describes ORCID as stated by ORCID.org.</i>	Via secondary
57	Zenodo. “Manage versions.” Zenodo Documentation. https://help.zenodo.org/docs/deposit/manage-versions/	Accessed 8 Oct. 2026 Public · Opened in full	Verified
58	Synapse. Listing for “Towards AGI – Organic Deterministic Intelligence,” DOI 10.5281/zenodo.19015971. https://www.synapsesocial.com/es/papers/69b79fc18166e15b153ac5a7	Accessed 8 Oct. 2026 Public · Search excerpt <i>Aggregator listing; Zenodo record not opened.</i>	Attributed
59	LawSites (LawNext). “Court Imposes Sanctions on Lawyers Who Filed Bogus Cases After Relying on ChatGPT.” June 2023. https://www.lawnext.com/2023/06/court-imposes-sanctions-on-lawyers-who-filed-bogus-cases-after-relying-on-chatgpt-for-legal-research.html	June 2023; Accessed 8 Oct. 2026 Public · Search excerpt <i>Quotes the opinion.</i>	Verified