

Computational Epistemology: Foundations

Evidence and Its Justificatory Reach in the Sciences and the Humanities

Gerd Graßhoff ^{ID}

Philosophisches Institut, Forschungsinstitut des Deutschen Museums
gerd.grasshoff@hu-berlin.de

ABSTRACT

This article establishes computational epistemology as the methodological ground for AI-assisted scholarly inquiry. Its central claim is that evidence and justification are structured differently across the sciences of natural phenomena and the humanities of scholarly understanding, yet both rest on a single validity apparatus: argument reconstruction, premise identification, and evidential support. Part I develops the required conceptual foundations — epistemic systems, Wittgenstein’s rule-following, and the sentence/ proposition/argument triad. Part II is the core contribution: a unified theory of evidence that distinguishes the practice of witnessing and causal regularity in the natural sciences from the practice of source-grounded interpretation, hermeneutics, and the reconstruction of intentional action in the humanities. Part III shows why retrieval-augmented generation commits a category error against this backdrop. Part IV describes the computational realisation — argument reconstruction, evidence-directed search, Partitur scores, oracle-verified golden queries as the verification layer, and justified answers over a structured ground, in which every answer to a collection is the conclusion of an argument with addressed premises, declared rules and searched defeaters. Part V draws implications for how scholars and AI platforms should divide epistemic labour, and for how evidential gaps become research agendas.

KEYWORDS computational epistemology · evidence theory · argument reconstruction · hermeneutics · intentionality of action · natural sciences and humanities · epistemic systems · AI scholarly infrastructure

1 Part I: Foundations

1.1 Why Epistemology Now

The platforms that scholars now use to query historical documents, extract claims from correspondence, and reason across thousands of pages share a common design presupposition: that retrieving text is equivalent to finding knowledge. This presupposition is false, and its falsity is not a technical accident. It is a philosophical error — the confusion of a linguistic object with what that object says, and of finding with justifying. Correcting this error requires epistemology.

Epistemology is the discipline that asks how beliefs become knowledge: what justifies the step from holding a proposition to being warranted in asserting it. When we ask whether Euler’s tidal model correctly predicts spring tides, whether a letter attributed to Kepler was actually authored by him, or whether a pattern in correspondence evidence really supports the inference that two scholars shared a research agenda, we are asking epistemological questions. The difficulty is not merely empirical. Each question presupposes a practice of evidential assessment — a set of standards governing what counts as a reason for or against a claim. Those standards differ between physics and philology, between natural history and the history of human intentions. Computational epistemology is the project of making those practices explicit and implementing them as infrastructure for AI-assisted scholarly inquiry.

This article is foundational. It does not build one system; it establishes the conceptual ground on which any honest system must stand. Three companion articles describe concrete realizations: *AI Erasmus* [1] develops the AI Erasmus platform; *AI Scholar* [2] specifies the document infrastructure; *AI Partitur* [3] formalises the score language by which agents are directed. A fourth companion, *Golden Queries in AI Erasmus* [4], shows how oracle-verified test questions make platform trustworthiness measurable. These implementations inherit their epistemological commitments from the present article.

1.2 Epistemic Systems

The concept of an *epistemic system* is the central theoretical contribution of my Habilitation work, developed and applied over three decades in studies of scientific discovery, experimental reasoning, and computational reconstruction [5], [6], [7], [8], [9]. Understanding it is essential before the epistemology of evidence can be stated precisely.

An epistemic system is a structured whole in which one or more agents execute goal-directed actions aimed at acquiring knowledge and avoiding error. What distinguishes an epistemic system from action in general is the character of its goal and its conclusive move. Standard action theory — from Aristotle’s practical syllogism through Davidson’s belief-desire model and Bratman’s account of plans [10] — explains why agents do what they do by pointing to their desires and beliefs. My modification targets a specific class of actions: those whose conclusion is not a generic behavioural outcome but the *formulation, reformulation, or modification of a hypothesis*.

i Note

The Epistemic Action Principle. An agent A, with background knowledge K bearing on hypothesis H, having the goal of advancing H in light of evidence E, and recognising that circumstances C obtain, performs an action whose conclusion is: modify H to H’. That is, the epistemic action produces a specific change in the content of what is believed.

This is not Davidson’s model, which covers all intentional action. It is not a general planning account. It is specifically designed to capture what it means to pursue knowledge: to take an observation, a document, an argument, a correspondence as bearing on a hypothesis, and to revise belief in consequence. The earlier computational reconstruction of the discovery of the urea cycle demonstrated that this framework is not only theoretically motivated but empirically tractable: the sequence of Hans Krebs’s notebook entries and experiments can be reconstructed as a trajectory through hypothesis space, each step explained as an epistemic action [11], [12].

1.2.1 AI Agents as Mechanical Epistemic Actors

Contemporary AI agents — systems that query document collections, extract claims, route evidence, and generate structured outputs — are mechanical actors in epistemic systems. They execute sequences of epistemic actions (querying, retrieving, classifying, inferring) directed toward epistemic goals (answer a question, substantiate a premise, identify a gap). The vocabulary has changed since 1995; the conceptual structure has not. What the literature now calls *agentic AI*, *multi-agent pipelines*, and *orchestration* is nothing other than the division of epistemic labour that epistemic systems have always required: different actors specialising in different phases of inquiry, coordinated toward a shared epistemic goal.

This continuity has a normative implication. Because AI agents are actors in epistemic systems, they are subject to the same normative assessment as other epistemic actors: was the observation careful? was the inference valid? was the evidence weighed correctly? The capacity to plan effectively — which LLM-based agents possess in impressive measure — is distinct from the capacity to assess propositional truth. A system that reliably produces plans may be entirely unreliable as a source of justified belief. Epistemic assessment of AI outputs requires the framework developed here.

The structure of an epistemic system also clarifies who bears responsibility. An agent in an epistemic system is responsible for executing assigned actions according to declared standards. The *epistemic subject* — the one who holds beliefs, makes knowledge claims, and bears the burden of justification — is the human scholar. AI agents contribute to the production of knowledge; they do not *possess* it.

1.3 Wittgensteinian Foundations

The philosophical tradition guiding this article is the analytic tradition, and within it the philosophy of Ludwig Wittgenstein — not chosen arbitrarily but because it provides the most precise tools for the three problems that computational epistemology must solve.

Problem 1: What AI systems process. AI models process tokens. Tokens are marks — syntactic objects. The question of whether processing tokens constitutes understanding, meaning, or knowing is not a technical question. It is a philosophical one, and Wittgenstein’s *Tractatus Logico-Philosophicus* provides the first relevant tool: the distinction between the sentence-sign (*Satzzeichen*), which is a physical mark or sound, and the proposition, which is the sentence-sign *in its projective relation to the world* [13, §3.12]. A system that manipulates sentence-signs without that projective relation is not ipso facto manipulating propositions.

Problem 2: What it means to follow a rule. AI systems are trained to produce outputs that accord with human judgement. Whether this constitutes *understanding* a rule or merely *matching* patterns is the question that Wittgenstein’s rule-following considerations raise directly [14, §§138–242]. His answer is that rule-following is constituted by practice — by participation in a form of life that sustains normative expectations. AI outputs are judged against human practice, not against some abstract rule they internally grasp. This means: the correctness of an AI output is not intrinsic to the output but relational to the practice that assesses it. AI systems can be trained to be reliable; they cannot be the ultimate arbiters of their own reliability.

Problem 3: Where justification ends. Justification is always finite. The demand for an infinite chain of reasons — justify every premise, and every premise used to justify that premise — leads to regress. Wittgenstein’s *On Certainty* resolves this by observing that some propositions function as *hinges*: they do not stand in need of justification because they constitute the framework within which justification operates [15, §341]. In scholarly practice, these hinges include: that documents have physical existence, that paleography as a discipline has developed reliable methods, that physical constants discovered by experiment hold across time. A computational epistemology system must be calibrated against such certainties, not required to prove them.

These three Wittgensteinian contributions — the sentence/proposition distinction, rule-following as practice, and hinge certainties — run through the entire argument. They are not ornamental. They determine what can be claimed about AI knowledge, AI reliability, and the structure of evidence.

1.4 The Terminological Core: Sentences, Propositions, Arguments

Before evidence can be discussed, the objects that evidence supports must be identified precisely. This requires distinguishing three things that contemporary AI discourse habitually conflates.

A sentence is a syntactically well-formed expression in a natural or formal language. “The Earth orbits the Sun” and “Die Erde umkreist die Sonne” are two distinct sentences. Sentences are linguistic objects; they are marks or sounds.

A proposition is what a declarative sentence expresses — its content, the thought it conveys. The two sentences above express the same proposition: *that the Earth orbits the Sun*. Propositions are the primary truth-bearers: it is propositions, not sentences, that are true or false. The sentence “Snow is white” is neither true nor false; the proposition *that snow is white* is true. This distinction is not pedantry. It is the ground of all subsequent analysis.

A statement is a proposition asserted by a speaker in a particular context. Copernicus’s statement *that the Earth moves* is the proposition *that the Earth moves* as asserted by Copernicus in the *De Revolutionibus* of 1543. The same proposition, asserted by different speakers at different times in different circumstances, constitutes different statements — with different evidential histories and different justificatory demands.

An argument is a structured set of propositions standing in logical relations. Formally: an argument consists of *premises* (propositions offered as reasons) and a *conclusion* (the proposition for which reasons are offered). The notation is: from premises $P_1, P_2, \dots, P_n$, derive conclusion C . An argument is *valid* when there is no possible situation in which all premises are true and the conclusion false — that is, when the premises

guarantee the conclusion. An argument is *sound* when it is valid and its premises are actually true. To know the conclusion, the argument must be cogent: valid, with premises that are not only true but known.

i Note

Why this matters for computational systems. AI systems process sentences (token sequences). The question of whether processing sentences yields access to propositions — to truth-bearers and their logical relations — is precisely what cannot be assumed. A system that retrieves sentences matching a query has found relevant *text*, not verified relevant *propositions*. The sentence/proposition distinction is not optional; it is the hinge on which the critique of retrieval-based reasoning turns.

1.4.1 Inference, Reasoning, and Explanation

Three operations are commonly and damagingly conflated in descriptions of what AI systems do.

Inference is the truth-preserving transition from premises to conclusion. In deductive inference, if the premises are true, the conclusion must be true. Formal logic — propositional and predicate — provides the calculus of such transitions.

Reasoning is the broader activity of constructing, evaluating, and revising such inferences. It includes identifying what premises are needed, whether they are true, whether the inference is valid, and what further inquiry the result demands.

Explanation provides the reasons why a proposition is true. To explain *why* the planets move in ellipses is to exhibit the premises — Newton’s laws, the inverse-square law, initial conditions — from which the elliptical motion follows.

Practical reasoning is a fourth operation that must be distinguished from the three above. Practical reasoning moves from goals and beliefs to intended actions: “I want to achieve G; action X is a means to G; therefore I shall do X.” This is the reasoning AI planning systems perform. It is genuinely reasoning, but it is *action-directed*, not *truth-directed*. The capacity to plan effectively — to chain tool calls, decompose subgoals, write executable steps — is not the capacity to assess whether the propositions consulted in that planning are true. This distinction is fundamental. Chain-of-thought prompting and agentic tool use are practical reasoning. Neither, by itself, constitutes epistemic assessment.

2 Part II: Evidence and What It Can Justify

2.1 The Central Claim

Evidence is whatever an epistemic practice licenses as a reason for or against a proposition. This minimal definition defers nothing — it does not pre-determine what counts as evidence, or how strong it is, or what it can establish. But it carries immediate implications.

First: evidence is practice-relative. What the astronomer counts as evidence for a model, the philologist does not automatically count as evidence for a reading. Each discipline has developed over centuries a set of standards — methods of observation, procedures of source criticism, canons of argumentation — that determine what evidence is within its domain. These standards are not arbitrary; they are responses to specific types of proposition and specific sources of error that different domains of inquiry have encountered.

Second: justification is argument-shaped. A piece of evidence *E* justifies proposition *P* only when there is an argument whose premises include *E* (and perhaps background knowledge *K*) and whose conclusion is *P* (or whose conclusion raises the probability of *P* in a measurable way). Evidence that does not figure in any such argument does not justify anything. Finding a document is not yet evidence for anything; the document justifies a proposition only when it figures as a premise in a reconstructible argument.

Third: the type of proposition determines the type of evidence required. A proposition about a causal law of nature requires different evidence than a proposition about what a historical author meant. This is the insight behind the distinction between natural sciences and humanities — the *Naturwissenschaften* and *Geisteswissenschaften* — that Dilthey articulated at the end of the nineteenth century [16] and that continues to structure scholarly practice. But the distinction is not a distinction in logical form. Both types of proposition are assessed by argument reconstruction; both types of evidence are weighed against standards of quality and independence. The *apparatus* of justification is the same; the *practices* that fill it differ.

2.2 Evidence in the Natural Sciences

2.2.1 Observation under Error Models

The natural sciences — physics, astronomy, chemistry, biology, and their histories — produce propositions about regularities in natural phenomena. Such propositions are justified by observation: controlled, repeatable, and interpreted through the theoretical framework that makes the observation scientifically significant.

Observation is not the passive receipt of sense data. It is an epistemic action: designed to bear on a specific hypothesis, executed under conditions specified by the protocol, and interpreted within a background of theory and prior knowledge. This is what the concept of *epistemic system* captures precisely: the scientist designs observations to bear on the modification of specific hypotheses, and the observation counts as evidence only insofar as it figures in the argument from observation-statement to hypothesis-modification.

What makes natural-science evidence distinctive is the explicit role of *error models*. Every measurement is uncertain. The astronomer measuring the position of a star, the geometer measuring the angle of a sundial groove, the chemist measuring reaction rates — all obtain not a single value but a distribution of values around a true value, with a spread determined by instrument precision, environmental noise, and procedural error. Evidence for a natural-science proposition is never a bare observation; it is an observation-with-uncertainty-estimate, interpreted against a model of what departures from the true value are to be expected.

i Note

Witnessed quantities and independent agreement. The most powerful natural-science evidence comes not from a single measurement but from multiple independent measurements of the same quantity by different methods or instruments. When independent witnesses converge — when optical and radio measurements of the same stellar position agree, when independent archaeological methods date the same stratum to the same century — the probability that all are simultaneously wrong in the same direction drops sharply. Independent convergence is the gold standard of natural-science evidence.

2.2.2 Witnesses and Consensus: The Sundial Case

The reconstruction of ancient Greek conic sundials from three-dimensional laser scans provides a particularly clean illustration of natural-science evidence practice [17]. A conic sundial is a hemisphere-shaped or cone-shaped stone instrument carved with lines that mark the hours of the day. The lines are not directly accessible from the surface: the instrument must be reconstructed by identifying the geometrical parameters — cone axis inclination, latitude for which it was designed, the spacing of hour lines — from the physical traces the stone preserves.

The physical traces available are: the outlines of carved grooves, the position of the gnomon hole, the shape of the carved face, and incidentally preserved alignment marks. Each of these is a *witness* — an independent source of information about the same underlying geometric parameters. The reconstruction method treats the witnesses as providing simultaneous constraints on the parameter values: the cone axis inclination, for instance, is constrained by the gnomon hole position, by the symmetry of the grooves, and by the spacing of the solstitial lines. When all witnesses are satisfied by a single parameter value, that value is taken as the best reconstruction of the ancient craftsman's *intention*.

Crucially, outliers are not averaged away. An isolated groove that departs from the predicted position at fourteen standard deviations is not an argument for a revised model; it is evidence of localised artisan error — a slip of the chisel that left a trace inconsistent with the design. The distinction between *systematic* departures (which revise the model) and *localised* departures (which are attributed to execution error) requires an explicit error model and an explicit account of what counts as independent evidence for the design parameters. Without the error model, all measurements look equivalent; with it, the thirteen-sigma outlier is recognised as noise.

This example exhibits the general structure of natural-science evidence:

1. Identify the proposition to be established (the design parameters of the instrument).
2. Determine what independent witnesses bear on that proposition (groove positions, gnomon hole, face geometry).
3. Specify an error model (expected measurement uncertainty for each witness type).
4. Find the parameter values for which all witnesses are jointly satisfied within their error bounds.
5. Report outliers explicitly and attribute them to their causes (design deviation, artisan error, instrument damage).

2.2.3 Causal Regularities

Beyond individual observations, the natural sciences establish propositions about *causal regularities*: laws that hold across instances. The evidence for a causal regularity is not a single observation but a systematic pattern of observations interpreted within a theoretical framework.

The theory of causal regularities that my work with Michael May developed — what we call *minimal theories* — provides the logical structure [18], [19]. A *minimal sufficient condition* (MSC) for an effect E is a set of factors $\{F_1, F_2, \dots, F_k\}$ such that whenever all factors in the set are present, E occurs, and no proper subset of the factors suffices. A *minimal theory* for E is the disjunction of all MSCs: either this cluster or that cluster of conditions is sufficient to produce E .

The evidence for a causal regularity in this sense is a *regularity across instances*: repeated occurrence of E whenever the relevant cluster of conditions is present, and absence of E whenever the cluster is absent. Crucially, what counts as evidence for one MSC can be evidence *against* another: if factor F is absent in a case where E occurs, then no MSC that includes F is a sufficient condition for E . The logical structure of the evidence space is rich — the evidential relations are not simply accumulative — and reconstructing them requires the argument-reconstruction apparatus of Part I.

This account of causal evidence is explicitly distinct from Pearl’s interventionist or structural causal model framework [20]. The minimal-theory framework is a regularity-theoretic account: causal claims are about what *reliably* co-occurs and co-varies across instances, not about hypothetical interventions in counterfactual models. The distinction matters for computation: evidence relevant to a regularity-theoretic causal claim is found in observed instances, not in intervention studies that may not be available for historical data.

2.3 Evidence in the Humanities

2.3.1 The Page as Unit of Evidence

The humanities — history, philology, the history of science, the history of ideas — produce propositions about *texts*, about *persons*, and about *human action*. The evidence for such propositions is documentary: documents that have survived, that have been transmitted, copied, edited, translated, and interpreted across centuries.

The *page* is the primary unit of evidence for the humanities, understood as the image-and-transcription pair: a facsimile of the original surface, together with a scholarly transcription that makes its content computationally processable [2]. The page in this sense is *Grounding Evidence* — the direct textual or observational basis for historical claims, as distinct from secondary literature that analyses, paraphrases, or summarises those bases. When a claim is disputed, the dispute moves to a specific page: “show me where in the document this is asserted” is the humanities equivalent of “show me the raw measurement.”

The evidential hierarchy in the humanities is anchored in this distinction:

- *Primary sources*: original documents, manuscripts, letters, instruments — the page is the ultimate ground
- *Critical editions*: scholarly reconstructions of primary texts, with apparatus that documents variants and editorial decisions
- *Secondary literature*: scholarly analysis and interpretation of primary sources
- *Tertiary accounts*: summaries, encyclopaedia articles, textbook synopses

This hierarchy is not a quality ranking of authors. It is a ranking of *epistemic proximity to the fact*. A claim about what Euler wrote in E57 is best grounded in the facsimile of E57, better grounded in a critical edition, less well grounded in a scholarly monograph about Euler, and least well grounded in a general reference work. Each step away from the primary source introduces an additional layer of interpretation that must be checked against the original.

2.3.2 Source Criticism and Provenance

Document evidence carries conditions that natural-science observation evidence does not: it has been *transmitted*, which means it may have been copied with errors, interpolated, altered, or forged. Source criticism — the systematic assessment of a document’s authenticity, integrity, and transmission history — is the humanities analogue of instrument calibration and error modelling in the natural sciences.

Provenance is the chain of custody: from the original composition of a document through its manuscript tradition, its editions, and its current physical location, to the scan or microfilm from which a researcher works. A computational epistemology system must track provenance as part of evidential metadata. The question is not only “what does this sentence say?” but “what is the evidential status of this sentence, given the document’s transmission history?” A sentence copied from a corrupted manuscript through three unreliable intermediaries warrants much weaker belief than a sentence from the author’s autograph. Provenance does not determine truth — corrupt transmission can preserve true content, and a pristine autograph can contain falsehoods — but it determines the strength of the evidential connection between the sentence and what the author wrote.

2.3.3 Hermeneutics as Rule-Governed Practice

Interpretation is the act of determining what a text *means* — what proposition or set of propositions it expresses. In the humanities, interpretation is never automatic: texts written in earlier periods, in languages now less familiar, in genres with conventions no longer in use, require skilled reading that draws on philological training, knowledge of the period, and familiarity with the author’s other work.

A common mischaracterisation of humanistic interpretation treats it as the intuitive *divination* of an author’s private mental state: reading behind the words to recover the subjective intention that produced them. This mischaracterisation has generated the so-called hermeneutic circle as a logical problem, and various romantic accounts of *Einfühlung* (empathetic understanding) as its resolution.

The Wittgensteinian tradition offers a different and more productive account. Understanding a text is not divining a private mental state; it is applying *shared interpretive practices* that constitute the meaning of the text for a reading community. These practices are rule-governed in Wittgenstein’s sense: they can be taught, learned, corrected, and extended. They include:

- Grammar and syntax: rules for parsing the words and sentences of the language
- Genre conventions: rules for reading a letter, a treatise, a scientific report
- Period conventions: rules governing what expressions were used to convey what ideas in a given time and place
- Context rules: rules for reading a passage in light of the surrounding text and the author’s known positions

Meaning, on this account, is not a hidden inner state but a pattern of use within a practice. The LWgrammar project [21] operationalises exactly this Wittgensteinian insight: rather than asserting what a term means, it surveys how the term is used across a declared corpus (a set of primary sources within a declared time,

author, and discipline scope), and compiles usage rules from cited instances. No sense is asserted without a page behind it. The dictionary entry is a *grammar of use*, not a metaphysical statement about conceptual content.

i Note

Interpretation as argument. An interpretation of a text passage can be expressed as an argument: the premise is the text (with its provenance and context), background knowledge (grammar, period conventions, the author’s other work), and the interpretive principle (a rule of the relevant practice); the conclusion is the proposition that the passage means such-and-such. The interpretation is justified to the degree that the premises are established and the interpretive principle is well-grounded in shared scholarly practice. Interpretive disagreements are disagreements about premises or principles — and they can, in principle, be adjudicated by evidence.

This means that textual interpretation, while requiring forms of background knowledge not required for natural-science measurement, is not epistemically different in kind. It is argument-shaped. It is open to criticism and revision. It is grounded in evidence (the text and its context). And it admits of degrees of epistemic quality: interpretations grounded in better evidence, employing more reliable interpretive principles, and commanding wider agreement among informed readers are more warranted than interpretations that rest on single readings, contested principles, or unverified assumptions.

2.3.4 The Intentionality of Human Action

The hardest case for a unified epistemology of evidence is the attribution of *intentions* to historical actors. When we say that Kepler intended to demonstrate the harmony of the heavens, that Pirckheimer’s correspondence reveals a strategic management of humanist networks, or that Euler’s tide treatise was motivated by rivalry with d’Alembert — we are making claims about mental states: about what agents were aiming at, what they believed, what they were trying to achieve.

Propositions about intentions seem to lie outside the scope of the evidence practices described above. We cannot directly observe what Kepler intended; we cannot measure Pirckheimer’s network strategy; we cannot verify Euler’s motivations as we verify the physical parameters of a sundial. The empiricist tradition has often concluded that intention attributions are either unverifiable (and therefore not genuine propositions) or reducible to behavioural patterns (and therefore not genuinely mental). Neither conclusion is satisfying to the practising historian.

The account I propose draws on Anscombe’s analysis of intentional action [22] and on the epistemic action principle introduced in §1.2. An action is intentional under a description when the agent would acknowledge that description as answering the question “Why did you do this?” That is, an intentional action is one that the agent explains by citing reasons: the beliefs and goals that made the action intelligible as a means to an end. The evidence for an intention attribution is therefore *the agent’s expressed reasons*, together with the pattern of actions that those reasons render intelligible.

In the humanities, the documentary record — letters, manuscripts, annotations, laboratory notebooks, administrative records — provides evidence for intention attributions precisely because it preserves the expressed reasons of agents: their stated goals, their acknowledged beliefs, their self-presented motivations. The *Actor–Agenda–Act* framework, developed in the context of epistemic systems analysis [5], [6], captures this structure formally:

- An *Actor* is an agent with goals, beliefs, and resources
- An *Agenda* is a structured goal hierarchy that the Actor pursues over time, decomposed into sub-agendas, tasks, and communicative actions
- A *Communicative Act* is an action that advances, modifies, or signals the state of an agenda in a way accessible to others

This structure is the humanities analogue of the causal model. Just as causal evidence supports a claim of the form “C caused E because C regularly co-occurs with E under background conditions K”, intention evidence supports a claim of the form “A intended G because the pattern of A’s communicative acts is best explained by the hypothesis that A was pursuing agenda $\bar{A}$ toward goal G.”

The evidential logic is the same: convergence of independent witnesses (different letters confirming the same stated goal), coherence of the inferred agenda with the observed pattern of actions, and absence of alternative explanations that better fit the evidence. The difference from causal evidence is not logical but material: the *type* of premises that figure in the argument (expressed reasons rather than measured quantities), and the *background knowledge* required to evaluate them (knowledge of period conventions for how goals were communicated, rather than knowledge of instrument calibration and error models).

This account has an important consequence for AI-mediated humanities scholarship. The claim that an AI system can “understand” a historical agent’s intentions is to be assessed not by asking whether the AI has inner states resembling intentions, but by asking whether the AI’s outputs participate correctly in the shared interpretive practice that licenses intention attributions. An AI agent that extracts stated reasons from letters, maps them to agenda structures, identifies inconsistencies, and flags gaps in the evidentiary basis is a genuine participant in the epistemic practice — not because it has intentions of its own, but because it executes the rule-governed operations that the practice requires.

■ 2.4 The Unified Account: One Apparatus, Two Practices

The evidence practices of the natural sciences and the humanities are structurally parallel despite their surface differences:

Dimension	Natural Sciences	Humanities
Primary unit of evidence	Measurement / observation	Page / document
Error model	Instrument uncertainty	Transmission and reading uncertainty
Independent witnesses	Measurements by different methods	Independent sources testifying to the same fact
Causal claim	Regular co-variation under conditions	Agenda structure rendered intelligible by expressed reasons
Quality hierarchy	Direct observation → experiment → secondary analysis	Autograph → critical edition → secondary literature
Interpretation layer	Physical theory (reading the instrument)	Genre convention, grammar, period knowledge

In both domains, evidence justifies propositions through argument. In both domains, the quality of evidence depends on independence, proximity, and reliability. In both domains, interpretation is required at every level — even reading a measurement instrument presupposes theory. And in both domains, disagreements about evidence are in principle resolvable by appeal to shared standards, even when those standards are themselves contestable and historically variable.

The synthesis is this: computational epistemology provides a single *validity apparatus* — argument reconstruction, premise identification, evidential support assessment — that applies across both families of practice. What varies is the ontology of evidence and the background knowledge required to assess it. The apparatus is domain-neutral; the filling is domain-specific.

This is the sense in which computational epistemology is the *methodological ground* for AI-assisted scholarly inquiry across the sciences and the humanities. It does not flatten the differences between disciplines;

it identifies the common logical structure that all justified belief-formation must exhibit, and makes that structure computationally tractable.

■ 3 Part III: Why Retrieval Cannot Establish Truth

■ 3.1 What Retrieval-Augmented Generation Does

Retrieval-augmented generation (RAG) is the currently dominant paradigm for “grounding” AI language-model outputs in document collections [23]. Its operation can be described precisely:

1. A query sentence is encoded as a vector in a high-dimensional embedding space
2. Documents or document segments whose embeddings are near the query embedding are retrieved
3. A generative language model produces a response conditioned on the retrieved segments

At no point in this sequence is propositional content identified, logical structure represented, or evidential relations assessed. The system operates on sentences (or segments of text) as vectors, and the retrieval relation is *similarity in embedding space* — not logical entailment, not evidential support, not semantic equivalence.

This is not a failure of implementation. It is a feature of the architecture. And when one understands what RAG does, it becomes clear what it cannot do: it cannot establish truth, because establishing truth requires justification, and justification is an operation on propositions in argumentative structures, not on sentences as vectors.

■ 3.2 The Category Error

Ryle’s concept of a category mistake identifies the error of treating an entity of one logical kind as if it belonged to another [24]. A category mistake is not merely an empirical error; it is a conceptual error, one that cannot be corrected by more data or better training.

RAG systems commit a category error when they treat the retrieval of similar sentences as a species of justification. The error is threefold:

Similarity is not entailment. Two sentences that are semantically similar — whose embeddings are near each other in the representation space — may express logically independent or even contradictory propositions. The sentences “Combustion requires oxygen” and “Combustion releases phlogiston” would have high embedding similarity, since they concern the same domain. But the propositions they express are not merely logically independent; they are incompatible. Similarity does not carry logical relations.

Relevance is not truth. A highly relevant retrieved segment — one that addresses precisely the topic of the query — may express a false proposition. The phlogiston sentence is maximally relevant to a query about combustion; it is also false. Relevance and truth are orthogonal dimensions. RAG optimises for relevance; truth is simply not in its objective function.

Finding is not justifying. Justification is a normative relationship: E justifies P iff E provides a reason — within an appropriate epistemic practice — for holding P. Finding a sentence E that expresses a proposition similar to P is not a normative relationship. It is a distributional one. The move from “E is similar to P” to “E justifies P” is not licensed by any argument; it is a category confusion.

These failures cannot be patched by scaling, by better embeddings, or by larger corpora. A system that retrieves sentences cannot assess propositional content by retrieving more or better sentences. The limitation is structural.

■ 3.3 What Cannot Be Replaced by Retrieval

The following epistemic operations all require the apparatus of Part I — argument reconstruction, premise identification, and evidential assessment — and none of them is achievable by retrieval:

- *Validity assessment*: determining whether an argument's conclusion follows from its premises requires identifying the propositions expressed, their logical relations, and the inference pattern — not finding similar sentences.
- *Soundness assessment*: determining whether premises are true requires evidence-directed search — finding evidence that bears on each premise's truth conditions — not similarity-based retrieval.
- *Evidential gap identification*: determining that a premise in a reconstructed argument lacks adequate evidential support requires knowing what would count as evidence for that premise and confirming its absence — not querying for similar content.
- *Interpretation quality*: determining whether a textual interpretation is well-grounded requires applying shared interpretive practices to the source text and assessing the quality of the evidence and principles employed — not retrieving interpretations that resemble the current one.

A system grounded in evidence practice must *invert* the retrieval paradigm: rather than asking “what text is similar to my query?”, it asks “what evidence is required to justify proposition P, and where does such evidence exist or fail to exist?” This inversion is the distinctive contribution of computational epistemology.

■ 4 Part IV: The Computational Realisation

■ 4.1 Two-Phase Architecture

The conceptual framework of Parts I–II can be computationally realised through a two-phase architecture.

Phase 1 — Reconstruction transforms natural language text into explicit argument structures. Its inputs are source documents in a processable format (the PDAI page format, described in the AI Scholar companion [2]). Its outputs are schema-validated argument representations: identified premises, identified conclusions, logical relation types, and notes on implicit (enthymematic) premises that the argument presupposes but does not state.

Reconstruction is guided by the principle of charity: argue for the strongest defensible reading of each argument before criticising it. It is also guided by the principle of literal grounding: premises are attributed to specific passages in specific documents, with page references, not to the text corpus in general.

Phase 2 — Assessment evaluates reconstructed arguments. For each premise, it asks: what is the evidential support? Is the premise grounded in primary sources? Are the sources independent? Does the background knowledge required for the interpretive step have adequate justification? The output is not a verdict but a graded report: which premises are well-evidenced, which are weakly supported, which lack evidence entirely.

The gap between what is needed and what is available is the *evidential gap*. Gap identification is one of the most valuable outputs of computational epistemology: it transforms the scholarly question “what should I look for next?” from an unanswerable question about the general corpus into a specific question about which premises require which types of evidence.

■ 4.2 Argument Representation

Arguments are represented in a schema-validated format that captures the logical structure:

```
<argument id="arg-euler-tides-01">
  <claim>Euler's tidal model predicts spring-tide amplitude correctly.</claim>
  <premises>
    <premise id="p1" source="euler-1740" page="42">
      The gravitational attractions of Sun and Moon add when in alignment.
    </premise>
    <premise id="p2" source="euler-1740" page="47">
```

```

    The amplitude of the combined attraction is proportional to
    the sum of the individual amplitudes.
  </premise>
  <premise id="p3" status="implicit">
    Modern tidal observations are comparable in method to
    eighteenth-century prediction standards.
  </premise>
</premises>
<conclusion>
  Spring tides are approximately 2.5 times the amplitude of neap tides.
</conclusion>
<evidential-assessment>
  <premise-ref id="p1" support="strong" source-quality="primary"/>
  <premise-ref id="p2" support="moderate" source-quality="primary"
    note="Linearisation assumption not explicitly justified"/>
  <premise-ref id="p3" support="weak" status="gap"
    note="No eighteenth-century observational series located"/>
</evidential-assessment>
</argument>

```

The schema makes explicit what remains implicit in prose arguments: which claims are grounded where, which steps are inferential and which observational, and which premises lack the evidential support required for the argument to be cogent.

4.3 Evidence-Directed Search

Evidence-directed search is the operational inversion of RAG. Given a proposition *P* whose evidential support is to be assessed, the system does not ask “what text is similar to *P*?” It asks:

1. What type of proposition is *P*? (causal regularity, interpretive claim, intention attribution, ...)
2. What would count as evidence for *P* in the relevant epistemic practice?
3. Where in the document collection might such evidence be located?
4. What is the quality of the evidence found, and what is the quality of what is missing?

For a causal regularity claim, the evidence type is a pattern of co-variation across instances, and the search targets instances meeting the relevant background conditions. For an interpretive claim about a text passage, the evidence type is usage patterns across the relevant corpus (which the LWgrammar index provides), plus parallel passages in the author’s other work. For an intention attribution, the evidence type is expressed reasons in the actor’s correspondence and the coherence of the inferred agenda with the observed pattern of actions.

This means: the same document collection supports radically different evidence searches depending on the type of claim being assessed. An index of causal co-occurrence is needed for natural-science evidence; a lexical usage index is needed for interpretive evidence; a correspondence and agenda index is needed for intention evidence. The AI Erasmus platform provides all three, built over the same PDAI document substrate [1], [25].

4.4 The Partitur: Inspectable Method

One of the distinctive features of the computational realisation described here is that the method by which evidence is assembled is itself a scholarly artifact, not a black box. This is what the *Partitur* (score) paradigm achieves [3].

A Partitur is a YAML document that specifies, in a human-readable and scholar-editable form, the sequence of epistemic operations an AI orchestration performs, the standards that outputs must meet at each stage, the failure modes that invalidate a stage, and the learning targets that define what improvement means. The

Partitur is not implementation code; it is the score that a Conductor executes. Like a musical score, it can be cited, criticised, revised, and compared with alternative scores.

The epistemic importance of the Partitur is that it makes *declared method* a condition of evidential validity. A conclusion produced by a pipeline whose Partitur has not been declared — whose operations are not inspectable — cannot be cited as scholarly evidence, because the method by which the conclusion was produced cannot be scrutinised or replicated. A declared Partitur is the humanities equivalent of a published experimental protocol: it makes the conclusion criticisable.

i Note

Epistemic disconnect. When AI systems produce outputs whose derivation cannot be traced, scholars face an epistemic disconnect: they hold conclusions without access to the reasoning that produced them. The Partitur closes this disconnect not by making AI infallible but by making AI operations accountable. The scholar can follow the Partitur, assess each stage, and identify where the output deserves confidence and where it requires further scrutiny.

4.5 Golden Queries and Oracle Verification

The trustworthiness of an epistemic infrastructure is not established by its designers' assertion but by its measurable performance on questions with known, independently verified answers. This is the role of *golden queries* in the AI Erasmus platform [4].

A golden query is a scholarly question for which an *oracle* can be independently specified: a count, a set of documents, a pattern of relationships, that can be verified directly against primary sources without relying on the system being tested. The oracle is not produced by the system; it is the independently verified ground truth against which the system's output is measured. The question "How many letters in the Kepler correspondence mention the term *harmonia*?" has an oracle answer: a human scholar counting in the critical edition. When the system produces the same answer, the answer is trusted; when it diverges, the divergence is evidence for a specific failure — in retrieval, in lemmatisation, in scope specification — that can be diagnosed and corrected.

The governance structure for golden queries is explicit:

- *Harvest*: the fraction of goldens whose oracle is satisfied in a given run
- *Monitor log*: the inspectable record of each run, stage by stage
- *Learning loop*: a structured process by which diagnosed failures lead to Partitur patches, tested under an A/B gate against the golden suite before promotion

This is the computational realisation of what epistemic systems have always required: a test of whether the knowledge-acquisition machinery is working, measured against independently verified truths.

4.6 Factual Validation of Sentences

Golden queries test the machinery on questions. A scholarly text poses the converse problem: it consists of sentences, each of which may express propositions about sources, objects and measurements, and the author wants to know, sentence by sentence, which of these propositions the registered evidence justifies. The Famulus lector (ai-lektorat) answers this question for every sentence of a paper. Its design follows directly from the terminological core of Part I (Section 1.4): a sentence is not true or false; the propositions it expresses are. *Validating a sentence* therefore means establishing that every factual proposition the sentence expresses is justified by registered evidence and that none is contradicted by it.

Decomposition. A language model decomposes the sentence into a conjunctive normal form: a set of clauses, each a disjunction of literals over a registered vocabulary of predicates (*Bestand(x, collection)*, *Wert(x, feature, value)*, *StehtIn(x, feature, source)*, ...). Every clause carries a role — *core*, *identification*, *source statement*, *connection* or *presupposition* — with a weight (1.0, 0.8, 0.6, 0.4, 0.2). Parts of the

sentence that assert no fact — goals, norms, stipulations, evaluations, requests, references of the text to itself — are recorded as non-factual and never enter the evidence test. A sentence with no factual clause receives the verdict *no factual statement*.

The evidence base. For sentence s , the clauses to be tested are resolved against

$$W(s) = F(s) \cup H|_s,$$

where $F(s)$ are the facts that registered sources supply for the terms of s (manifests, object files, measurement runs, bibliographies), each with its address, and $H|_s$ is the part of the background rules H — identity equations, part-whole relations, closed-record rules — that the terms of s touch. A clause is *justified* if it follows from $W(s)$, *refuted* if its negation follows, and *unsupported* otherwise. The degree of justification $\delta(s)$ is the weight of justified clauses divided by the total weight.

The verdict. A sentence is *sachlich validiert* (factually validated) if and only if five conditions hold: (1) $W(s)$ together with the normal form is consistent; (2) every core clause is justified; (3) $\delta(s) \geq \theta$, with $\theta = 0.6$; (4) three checks on the decomposition pass — every term of the sentence appears as an argument of a literal or inside a classified non-factual part (term fidelity), no literal introduces a term foreign to the sentence, and a second, independent decomposition agrees on the load-bearing clauses and on what is non-factual; (5) the sentence is compatible with the earlier sentences of the text and with the registers. Otherwise it is *nicht validiert*. There is no third verdict “open”: *not validated* is the contradictory negation of *validated*, and it does not mean *false* — it means the registered evidence does not suffice. The design goal is asymmetric: no sentence whose propositions are false may be validated, at the price that some true sentences remain unvalidated while a source is missing. Such a sentence marks an evidential gap in the sense of Section 5.2.

Condition 5: compatibility. Consistency with earlier sentences is tested without matching surface strings. Every accepted clause is written as a ground fact into a per-paper DuckDB store; incompatibility patterns — two values of a functional feature, disjoint collections, an event before a birth or after a death, a count that contradicts an enumeration, a spelling that departs from the reference form of a name, a use of a term that a binding usage rule excludes — are negative binary Horn clauses, evaluated as indexed SQL joins. Names are linked by key equality, then by edit-distance bounds with a fixed initial, then by a lookup in a reference work (Wikidata, restricted to humans with life dates); an embedding similarity is used only as a hint, never as a ground. A finding blocks the verdict only under a *confirmation principle*: every fact it relies on must also be carried by the independent second decomposition. A defect of one model reading therefore produces a hint, not a verdict.

Resolution. Because the evidence base is almost entirely Horn, the decision procedure is fast. Unit propagation with two watched literals per clause [26] decides Horn sets in linear time [27]; only the small non-Horn remainder goes to a DPLL search [28], and the residue to saturation by resolution [29]. Every derived literal keeps a pointer to the clauses that justified it, so the reasons shown to the author are the derivation itself, read backwards to the registered facts. On the Famulus platform paper (557 sentences) the local tests take a median of 72 ms per sentence; the time a sentence costs is dominated by the model calls of the two decompositions, not by resolution.

Terminology. Technical terms are checked once per paper, in the background: for every term, a dictionary entry in the style of the LWgrammar institution [21] records its definition and the roles in which the paper may use it (sign, truth-bearer, part of the world, action, agent ...), bound to active LWgrammar senses and blocking usage rules where they exist. Later checks match against this entry. This is where the terminological core becomes operational: under the binding rule that consequence and contradiction hold between propositions, not sentences, the sentence “The number was right; the sentence contradicted it” uses *sentence* as a truth-bearer and is not validated.

i Note

A worked case. The platform paper stated that the sundial ObjID126 “from the Waepler corpus” carries its latitude $\varphi = 38.01^\circ$ as the consensus of five independent measurement series. The decomposition yielded, among others, `Bestand(objid:126, waepler)`, `Wert(objid:126, breite, 38.01°)` and `Wert(objid:126, zeugen, fünf)`. The first was refuted by the Berlin Sundial Collaboration records, which place the object in the collection *Ancient Sundials* (Delos, Inv. B3652); the second was justified by the joint estimate of the reconstruction (Section 2.2.2); the third was refuted, because that estimate rests on four witnesses. After the author corrected collection and count, both clauses are justified. The sentence nevertheless remains unvalidated: its remaining core clauses, “as a consensus by the witness principle” and “not interpolated from a model”, have no registered evidence. That is the correct outcome under the asymmetric design.

Meta-validation. The validator is itself subject to the discipline of Section 4.5. A reader who holds a verdict to be wrong reports it; the report is assigned to exactly one error site — decomposition gap, role error, rule error in H , register error, source error or source gap — and yields candidate changes to rules, registers or the decomposition instruction. A candidate is accepted only by the owner and only after a gate: the reported sentence obtains the expected verdict, the golden sentences keep zero false validations and no higher rate of false non-validations, and a mutation test — sentences whose numbers, names and relations are deliberately falsified — validates none of them. On the platform paper, 7 sentences are validated, 474 not validated and 76 carry no factual statement; none of 28 mutants of the validated sentences is validated. That the calibration of weights and threshold barely moves these figures shows where the bottleneck lies: in missing sources, not in the setting of θ . The verdicts, with their reasons, are open to the same communal scrutiny that Section 5.3 assigns to every scholarly claim.

■ 4.7 Justified Answers over a Structured Ground

Factual validation asks of a given sentence whether the propositions it expresses are justified. A question put to a collection poses the converse task: there is no sentence yet, only a propositional form with an open place — *Kármán corresponded between 1930 and 1940 with x persons* — and the system is to supply the value together with its justification. The Famulus Notebook realises this for the PDAI collections, whose letters, pages, sentences, mentions, topics and speech acts are held in a DuckDB sidecar per collection [30].

The realisation inverts the order of retrieval (Section 3.2). A question is assigned to one of 56 declared query forms; the form’s single SQL statement proposes an answer and names its witnesses (the letters, acts or sentences that would carry it). For every witness the addressed facts are read — the register’s sender, recipient and date, the sentence a mention stands in, the model’s reading of a request with the precision measured for such readings — and the rules of the form’s *argument scheme* are grounded over them. The same ground-resolution prover that decides factual validity (Section 4.6) decides whether the instance follows; its leaves are the evidence. The scheme’s *critical questions* [31] are the collection’s error model written as queries over the same data: letters without a resolved party or date, register entries that share a name form, duplicates, a relayed request, an evidence sentence that does not carry the reading’s words, an unreliable language label, period boundaries that are only proposed. They are searched before any sentence is written, and each has a declared effect: it turns a count into a lower bound, merges what is shown to be the same, downgrades or — when two independent witnesses agree — upgrades an instance, or refutes the conclusion [32].

The search for this justification is an expansion of a goal, and the turn from search to justification is where the two tasks meet. The goal is to justify the proposition the form’s query proposes; it is expanded by the rules whose head it matches, an alternative per rule, with the atoms of a rule’s body as new goals. A branch may end only in a ground: an addressed record, a reading of measured precision, a declared stipulation or a computed statement. A predicate with no ground and no rule that reaches one has no stock, and its branch closes at once as a gap. Where every open atom is closed, the expansion is a refutation of the negated

conclusion; reducing it by deletion to the clauses without which the conclusion no longer follows at the same grade leaves a core in which every node contributes. The same tree is drawn for a sentence of a paper, from the verdict of factual validation (Section 4.6): its leaves are again registered grounds or marked gaps [30].

Two consequences follow from the account of Part II. First, a count is itself an inference that needs a closure premise: *exactly* n follows only if the register is declared complete for the range of the question. For a register without declared closure the justified conclusion is *at least* n , and absence does not follow at all. The quantifier of the answer is therefore not a caveat added to it but what its grounds entail; in the pilot, the 491 counterparts the query returns for Kármán's correspondence of 1930–1940 become, after duplicates and register entries for one person are merged, *at least 443 persons (1,390 letters)*. Second, the weight of a model-read premise is its measured precision, not a confidence the model assigns itself: a reading counts as a medium-grade premise only if the lower bound of the 95 % interval of its precision on a sample reaches 0.7. Conclusions that rest on readings below that grade, or on conventions not yet confirmed by the owner, are *not validated*, and the argument names the weakest premise; in the pilot, 63 of 74 catalog questions were validated before any learning, and every one of the remaining ten named a specific task — confirm the war boundaries, measure the act reading apart from its locus, add a second witness to a surname mention.

The answer is written from the validated argument: its first sentence is the conclusion, and the model may add detail only through placeholders for row values, so that no number reaches the answer that the argument does not carry. A retrieval baseline on the same 74 questions, answering from twenty BM25 passages with the same model, cited passages in 63 answers; of the 62 questions whose argument establishes a count, it answered 28 with a count, 24 of them with numbers that neither the query nor the argument establishes, and none of the 61 lower-bound questions with *at least*. The category error of Section 3.2 is measurable.

The argument is stored as one file per answer and shown in the notebook as an interactive tangled tree with an explanation box per node, down to the register entry and the facsimile page. A reader who finds it unconvincing annotates a node; the annotation starts a learning episode in which a lesson — a change to the declarative schemes, never to one case — is diagnosed, written, and gated against the annotated case, every affected case and a mutation test with invented and displaced witnesses, before it is promoted; every promoted lesson remains revocable by the owner. The discipline of Section 4.5 thus extends from the machinery's answers to the machinery's reasons.

■ 5 Part V: Implications

■ 5.1 AI as Famulus, Not Oracle

The title of the AI Erasmus companion derives from Erasmus's Basel workshop of the 1520s [1]: a household in which trained assistants (*famuli*) copied, collated, fetched, and checked — so that Erasmus's judgement could be spent where only judgement serves. The familiar in the humanist workshop did not settle arguments; he prepared the material on which arguments were settled.

This image captures the correct epistemic role for AI in scholarly inquiry. An AI system can:

- Process large document collections at scales impossible for human readers
- Extract candidate argument structures for human refinement
- Identify evidential gaps — premises with inadequate support — that direct inquiry
- Compile usage statistics and cross-reference patterns across corpora
- Maintain provenance chains from assertion to primary source

An AI system cannot:

- Assert truth (it produces outputs, not truth-bearers)
- Know propositions (it executes operations, not epistemic stances)
- Bear epistemic responsibility (it is a tool, not a knowing subject)

- Replace the interpretive judgement that requires engagement with the material, the period, and the scholarly community

The appropriate relationship is complementary, not competitive: AI contributes speed and scale; the scholar contributes judgement and accountability. The danger is not that AI will replace scholarly judgement but that scholars will abdicate it — treating AI outputs as authoritative without subjecting them to the same evidential scrutiny they apply to any other source. Computational epistemology provides the framework that makes such scrutiny possible.

■ 5.2 Evidential Gaps as Research Agendas

When argument reconstruction reveals a premise with inadequate evidential support, it identifies a research gap — a specific question that the existing evidence base cannot yet answer. Evidential gaps, made explicit by computational means, are more actionable than the general sense that “more research is needed”:

- A gap in the support for premise P_3 in argument A says: *this specific claim, playing this role in this argument, lacks grounding.*
- The type of gap — missing primary source, contested interpretation, unverified causal mechanism — determines where to look: archives, critical editions, further experiments.
- The *criticality* of the gap within the argument determines priority: a gap in a non-essential premise is less urgent than a gap in a premise on which the conclusion directly depends.

This transforms gap identification from a limitation into a research instrument. The computational system does not merely tell the scholar what it found; it tells the scholar, in a structured way, what it did not find and why that matters.

■ 5.3 The Scholarly Community as Ultimate Arbiter

No individual — human or artificial — is the final judge of scholarly knowledge. Knowledge is a social achievement: produced through public articulation of claims, critical response, iterative refinement, and eventual convergence toward warranted consensus. This is what Longino calls the *fate of knowledge* — that it depends on the social organisation of critical discourse [33] — and what Shapin traces in the history of truth as social institution [34].

Computational epistemology provides infrastructure for this social process, not a replacement for it. When argument structures are explicitly represented, scholars can scrutinise them — add evidence, challenge premises, propose alternative reconstructions. When Partitur scores are public and citable, the method is open to community assessment. When evidential gaps are made explicit, the scholarly community can collectively decide which gaps are most urgent and which archival, experimental, or interpretive efforts should address them.

This infrastructural role is not new; it is continuous with the history of scholarly communication. The critical edition, the scholarly journal, the peer-review process — all are institutional arrangements that make knowledge claims publicly accessible and open to communal assessment. The *Republic of Letters* functioned as an epistemic community precisely because correspondence made claims public and invited criticism [35]. AI Erasmus and its companions are the contemporary continuation of that tradition: making scholarly evidence publicly accountable at the scale that current document collections require.

■ 5.4 The Proper Scope of Computational Epistemology

Computational epistemology is not a theory of all knowledge; it is a methodology for scholarship — for the disciplined, evidence-based inquiry into historical, scientific, and textual matters that constitutes the core of university research. It does not claim that all justified belief must take the form described here. It claims that scholarly inquiry — the kind of inquiry that aims at public, criticisable, evidence-grounded propositions — has this structure, and that making this structure computationally tractable is a genuine service to scholarship.

The scope is also limited by the condition of interpretability. A computational system that produces conclusions without traceable reasoning is not a participant in scholarship, however sophisticated its outputs. Only

when reasoning is explicit — when Partiturs are declared, provenance is tracked, and golden query performance is measured — can AI outputs be incorporated into the scholarly evidentiary chain. This limitation is not a bug; it is the condition of trustworthiness.

6 Conclusion: The Methodological Ground

Computational epistemology is the methodological ground for AI-assisted scholarly inquiry because it identifies the logical structure that evidence and justification must exhibit regardless of the domain in which inquiry proceeds. That structure is argumentative: evidence justifies propositions by figuring as premises in arguments whose conclusions are the propositions to be established. The quality of justification depends on the quality of evidence, the independence of witnesses, the proximity to primary sources, and the reliability of the interpretive steps. These conditions apply in physics as in philology, in astronomy as in the history of humanist correspondence.

The key contributions of this article are:

- A foundational account of epistemic systems as the conceptual ground for understanding both human and AI epistemic actors
- A Wittgensteinian framework that precisely characterises what AI systems process (sentence-signs), what normative assessment requires (agreement with practice), and where justification terminates (hinge certainties)
- A unified theory of evidence that articulates the structural parallels and material differences between evidence in the natural sciences and evidence in the humanities, including the distinctive logic of hermeneutic interpretation and intention attribution
- A diagnosis of retrieval-augmented generation as a category error — treating sentence similarity as propositional justification — that cannot be remedied by scaling
- A description of the computational realisation: two-phase argument reconstruction and assessment, evidence-directed search, Partitur-governed method, and oracle-verified golden queries as the trustworthiness measure

What is not claimed is that computational epistemology makes AI infallible or eliminates the need for scholarly judgement. On the contrary: by making the operations of AI systems explicit, the framework exposes exactly where human judgement is required — in assessing evidential quality, in interpreting the significance of gaps, in deciding what counts as adequate support in the given context of inquiry. The famulus prepares the material; the scholar does the work.

The next steps are specific:

1. Argument reconstruction agents that extract premise/conclusion structures from PDAI documents, constrained by argument.dtd schema validation
2. Evidential quality assessors that classify evidence by type (primary/secondary, independent/dependent, measurement/interpretation) and assign quality scores
3. Gap identification reports that rank evidential absences by their criticality within reconstructed arguments
4. Extension of the golden-query suite to cover the full range of evidence types — causal regularity claims, interpretive claims, intention attributions — with type-specific oracles
5. Partitur integration: each evidence type has its own evidence-assessment sub-score, declared and revisable

These are not aspirational goals; they are the next steps in a programme whose foundational methodology this article establishes, whose platform infrastructure AI Erasmus and AI Scholar provide, and whose trustworthiness mechanism the golden-queries companion specifies.

Acknowledgements

Earlier work on epistemic systems, conducted with Michael May and Kärin Nickelsen in the context of studies on the urea-cycle discovery and on experimental methodology in the natural sciences, provided the conceptual foundation for this article. The AI Erasmus platform team, and the ongoing conversations about evidence practice and scholarly infrastructure that have shaped the platform's design, have shaped every argument here.

References

Bibliography

- [1] G. Graßhoff, "AI Erasmus — A Platform for a New Era in Digital Humanities," 2026.
- [2] G. Graßhoff, "PDAI: A Certified Document Infrastructure for Computational Historical Scholarship," 2026.
- [3] G. Graßhoff, "The AI Partitur — A Score Language for the Humanities," 2026.
- [4] G. Graßhoff, "Golden Queries in AI Erasmus: A Testbed for Scholarly Inquiry," 2026.
- [5] G. Graßhoff, *Die Kunst wissenschaftlicher Entdeckungen: Grundzüge einer Theorie epistemischer Systeme*. Hamburg: Habilitationsschrift, Universität Hamburg, 1995.
- [6] G. Graßhoff, R. Casties, and K. Nickelsen, *Zur Theorie des Experiments: Untersuchungen am Beispiel der Entdeckung des Harnstoffzyklus*. Bern: Bern Studies in the History, Philosophy of Science, 2001.
- [7] G. Graßhoff, "The Discovery of the Urea Cycle: Computer Models of Scientific Discovery," in *Computer Simulations in Science and Technology Studies*, P. Ahrweiler and N. Gilbert, Eds., Berlin: Springer, 1998, pp. 71–90.
- [8] K. Nickelsen and G. Graßhoff, "Concepts from the Bench: Krebs and the Urea Cycle," in *Going Amiss in Experimental Research*, G. Hon, J. Schickore, and F. Steinle, Eds., Boston: Springer, 2009, pp. 91–117.
- [9] K. Nickelsen and G. Graßhoff, "In Pursuit of Formaldehyde: Causally Explanatory Models and Falsification," *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, vol. 42, no. 3, pp. 297–305, 2011.
- [10] M. E. Bratman, *Intention, Plans, and Practical Reason*. Cambridge, MA: Harvard University Press, 1987.
- [11] G. Graßhoff and M. Schneegans, "Experimentation and Methodology: Applied to Cold Fusion Research," *Philosophia Naturalis*, vol. 32, no. 1, pp. 47–70, 1995.
- [12] G. Graßhoff and M. May, "Hans Krebs' and Kurt Henseleit's Laboratory Notebooks and their Discovery of the Urea Cycle: Reconstructed with Computer Models," in *Reworking the Bench: Research Notebooks in the History of Science*, vol. 7, F. L. Holmes, H.-J. Rheinberger, and J. Renn, Eds., in *Archimedes*, vol. 7., Dordrecht: Kluwer Academic Publishers, 2003, pp. 269–294.
- [13] L. Wittgenstein, *Tractatus Logico-Philosophicus*. London: Kegan Paul, Trench, Trubner, 1922.
- [14] L. Wittgenstein, *Philosophical Investigations*, 1st ed. Oxford: Blackwell, 1953.
- [15] L. Wittgenstein, *On Certainty*. Oxford: Blackwell, 1969.
- [16] W. Dilthey, *Die Entstehung der Hermeneutik*. Tübingen: Mohr, 1900, pp. 185–202.
- [17] G. Graßhoff, "Construction Geometry of the Conic Sundial with Horizontal Gnomon," *arXiv preprint*, 2026.

- [18] G. Graßhoff and M. May, “Causal Regularities,” in *Current Issues in Causation*, W. Spohn, M. Ledwig, and M. Esfeld, Eds., Münster: Mentis Verlag, 2001, pp. 85–114.
- [19] M. Baumgartner and G. Graßhoff, *Kausalität und kausales Schliessen*. Bern: Bern Studies in the History, Philosophy of Science, 2003.
- [20] J. Pearl, *Causality: Models, Reasoning, and Inference*, 2nd ed. Cambridge: Cambridge University Press, 2009.
- [21] G. Graßhoff, “Wittgensteinian Grammar and OED Dictionaries for AI: A Foundation of Semantic Understanding of Scientific Texts,” *arXiv preprint*, 2026.
- [22] G. E. M. Anscombe, *Intention*. Oxford: Blackwell, 1957.
- [23] P. Lewis *et al.*, “Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 9459–9474, 2020.
- [24] G. Ryle, *The Concept of Mind*. London: Hutchinson, 1949.
- [25] G. Graßhoff, “A Multi-Layer Index and Retrieval Architecture for Historical Document Collections,” 2026.
- [26] M. W. Moskewicz, C. F. Madigan, Y. Zhao, L. Zhang, and S. Malik, “Chaff: Engineering an Efficient SAT Solver,” in *Proceedings of the 38th Design Automation Conference*, Jun. 2001, pp. 530–535. doi: 10.1109/dac.2001.935565.
- [27] W. F. Dowling and J. H. Gallier, “Linear-Time Algorithms for Testing the Satisfiability of Propositional Horn Formulae,” *The Journal of Logic Programming*, vol. 1, no. 3, pp. 267–284, 1984, doi: 10.1016/0743-1066(84)90014-1.
- [28] M. Davis, G. Logemann, and D. Loveland, “A Machine Program for Theorem-Proving,” *Communications of the ACM*, vol. 5, no. 7, pp. 394–397, 1962, doi: 10.1145/368273.368557.
- [29] J. A. Robinson, “A Machine-Oriented Logic Based on the Resolution Principle,” *Journal of the ACM*, vol. 12, no. 1, pp. 23–41, 1965, doi: 10.1145/321250.321253.
- [30] G. Graßhoff, “Famulus Reasoning: Evidence, Reasoning Chains and Validation for Questions to Scholarly Collections,” 2026.
- [31] D. Walton, *Argumentation Schemes for Presumptive Reasoning*. Mahwah, NJ: Lawrence Erlbaum Associates, 1996.
- [32] J. L. Pollock, “Defeasible Reasoning,” *Cognitive Science*, vol. 11, no. 4, pp. 481–518, 1987, doi: 10.1207/s15516709cog1104_4.
- [33] H. E. Longino, *The Fate of Knowledge*. Princeton: Princeton University Press, 2002.
- [34] S. Shapin, *A Social History of Truth: Civility and Science in Seventeenth-Century England*. Chicago: University of Chicago Press, 1994.
- [35] L. Daston, “The Ideal and Reality of the Republic of Letters in the Enlightenment,” *Science in Context*, vol. 4, no. 2, pp. 367–386, 1991, doi: 10.1017/S0269889700001010.