Sci

MDPI

Article

Quantum Abduction: A New Paradigm for Reasoning Under Uncertainty

Remo Pareschi

Stake Lab, University of Molise, 86100 Campobasso, Italy; remo.pareschi@unimol.it

Abstract

Abductive reasoning—the search for plausible explanations—has long been central to human inquiry, from forensics to medicine and scientific discovery. Yet formal approaches in AI have largely reduced abduction to eliminative search: hypotheses are treated as mutually exclusive, evaluated against consistency constraints or probability updates, and pruned until a single "best" explanation remains. This reductionist framing fails on two critical fronts. First, it overlooks how human reasoners naturally sustain multiple explanatory lines in suspension, navigate contradictions, and generate novel syntheses. Second, when applied to complex investigations in legal or scientific domains, it forces destructive competition between hypotheses that later prove compatible or even synergistic, as demonstrated by historical cases in physics, astronomy, and geology. This paper introduces quantum abduction, a non-classical paradigm that models hypotheses in superposition, allowing them to interfere constructively or destructively, and collapses only when coherence with evidence is reached. Grounded in quantum cognition and implemented with modern NLP embeddings and generative AI, the framework supports dynamic synthesis rather than premature elimination. For immediate decisions, it models expert cognitive processes; for extended investigations, it transforms competition into "co-opetition" where competing hypotheses strengthen each other. Case studies span historical mysteries (Ludwig II of Bavaria, the "Monster of Florence"), literary demonstrations (*Murder on the Orient Express*), medical diagnosis, and scientific theory change. Across these domains, quantum abduction proves more faithful to the constructive and multifaceted nature of human reasoning, while offering a pathway toward expressive and transparent AI reasoning systems.

Keywords: quantum abduction; abductive reasoning; quantum cognition; semantic embeddings; forensic reasoning; scientific discovery

check for updates

Academic Editor: Gerald B. Cleaver

Received: 16 September 2025
Revised: 21 November 2025
Accepted: 4 December 2025
Published: 11 December 2025

Citation: Pareschi, R. Quantum Abduction: A New Paradigm for Reasoning Under Uncertainty. *Sci* **2025**, *7*, 182. https://doi.org/10.3390/sci7040182

Copyright: © 2025 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

1. Introduction

1.1. Motivation

Many critical domains—criminal investigation, clinical diagnostics, scientific discovery—require reasoning under ambiguity, contradiction, and evolving evidence. Classical abduction, as formalized in logic, Bayesian inference, or set-covering models, is typically reduced to eliminative selection: hypotheses compete within a fixed space until one prevails. This reductionist approach fails on two distinct timescales.

First, in immediate decision contexts (emergency medicine, real-time diagnostics), it does not reflect how human experts actually reason, often maintaining competing explanations in play until evidence forces convergence. Second, in extended investigative contexts (complex criminal cases, scientific theory development), the eliminative paradigm creates

counterproductive competition among research teams, investigative units, or theoretical schools, when collaborative "co-opetition"—where competing hypotheses inform and strengthen each other's development—would be more fruitful.

1.2. Dual Role of Quantum Abduction

Quantum abduction addresses both challenges through a unified framework. For immediate decisions, it models the parallel processing observed in expert cognition, where multiple explanatory threads remain active until collapse is necessary. For extended investigations, it transforms competitive hypothesis testing into a coordinated exploration where:

- Different investigative teams pursue distinct but informationally entangled hypotheses
- Evidence gathered for one hypothesis can constructively or destructively interfere with others
- The framework acts as a meta-cognitive coordinator, tracking how partial insights from competing approaches might synthesize into richer explanations
- Resources are allocated based on current amplitude distributions rather than premature winner-takes-all decisions

This "co-opetition" model is particularly valuable in complex criminal investigations (where tunnel vision on a single suspect can blind investigators to alternative leads) and scientific research (where premature paradigm commitment can delay breakthrough syntheses for decades).

1.3. Nature and Originality of This Work

This paper presents a research framework with initial conceptual validation for a novel approach to abductive reasoning. The framework, termed quantum abduction, is entirely original in its specific formulation, though it builds upon established work in quantum cognition [1] and extends our prior research on entangled heuristics [2]. While the latter focuses on strategic concept creation through ontological synthesis, quantum abduction is epistemically grounded—it seeks the best explanation of a single determinate reality rather than creating new conceptual ontologies. The use of quantum formalism as a computational metaphor for cognitive processes, while not new, is applied here in an original way to the specific problem of abductive inference.

From a computational and formal standpoint, we adopt the Hilbert space formalism from quantum cognition to model interference, contextuality, and order effects in explanation; see Section 3.1 for the full rationale and limits.

1.4. Structure of the Paper

The remainder of the paper is organized as follows. Section 2 motivates the use of Hilbert space formalisms for abductive inference and clarifies the scope and limits of the quantum analogy. Section 3 introduces the quantum abduction framework and its computational components. Section 4 illustrates its use across forensic, clinical, literary, and scientific contexts. Section 5 provides an initial empirical assessment using a minimal benchmark. Section 6 contrasts the approach with logic-based, Bayesian, set-covering, and other alternative abductive traditions. Section 7 concludes with implications, limitations, and directions for future development within the broader research program. A glossary of key terms is provided in Appendix A, and formal specifications are consolidated in Appendix B.

2. Background

2.1. Classical Abduction

Abductive reasoning—variously termed abduction, abductive inference, or retroduction—is a form of logical inference that proposes the most plausible explanation for a given set of observations. The concept was first articulated by the American philosopher and logician Charles Sanders Peirce [3] in the late 19th century, who framed abduction alongside deduction and induction as one of the three fundamental modes of reasoning. More recent accounts emphasize its cognitive grounding and role as a general reasoning pattern, both in scientific discovery and in everyday inference [4,5].

Unlike deduction, which guarantees the truth of its conclusions if the premises are true, abduction yields hypotheses that are tentative and provisional. Its conclusions are best understood as plausible guesses or informed conjectures, often described as "best available" or "most likely" explanations. Abduction also differs from induction: whereas induction generalizes from particular cases to broader regularities, abduction remains tied to specific observations, advancing hypotheses without claiming universality.

From the 1980s onward, the concept of abduction has garnered increasing interest in various applied fields, including law, clinical practice, computer science, and artificial intelligence. Diagnostic expert systems, for example, often used abduction to generate candidate explanations for observed symptoms or evidence. However, these implementations revealed both practical and theoretical limits. Practically, they were constrained by the computational resources of the time and by the lack of robust learning mechanisms. Theoretically, they were limited by the eliminative character of classical abduction itself: hypotheses are typically treated as mutually exclusive, and the reasoning process progresses by discarding inconsistent candidates until a single "best" explanation remains.

This eliminativist dynamic, grounded in the principles of non-contradiction and the excluded middle, creates brittleness precisely in situations where human experts exhibit flexibility—namely, in contexts of contradictory, incomplete, or entangled evidence. Classical abduction struggles to sustain multiple explanatory threads in suspension, forcing early closure at the expense of richer, hybrid accounts. These shortcomings provide the rationale for exploring quantum abduction, a framework that models inference in terms of superposition and interference, allowing explanatory alternatives to coexist and interact before eventual resolution.

2.2. Quantum Cognition

Quantum cognition is a theoretical framework that applies mathematical tools from quantum mechanics, not to claim that the brain is a quantum system, but to capture cognitive phenomena that resist classical probabilistic modeling. Introduced in the early 2000s [1,6], the approach has been used to explain puzzling features of human reasoning, such as order effects in decision-making, violations of classical probability axioms, and the coexistence of seemingly incompatible beliefs.

At its core, quantum cognition replaces classical probability spaces with Hilbert spaces, in which mental states are represented as superpositions of potential judgments or beliefs. Observations, questions, or contextual cues act as projection operators, modifying the cognitive state and leading to outcomes that display interference or entanglement effects. For instance, in survey responses, the probability of answering "yes" to two sequential questions often depends on their order—an effect naturally modeled as the collapse of a quantum superposition, but poorly explained by classical Bayesian frameworks.

We use these tools as a *modeling formalism* for non-classical aspects of reasoning (interference, contextuality), not as a claim about quantum hardware or ontological indeterminacy; detailed clarifications appear in Section 3.1.

2.3. Entangled Heuristics and Strategic Inference

The work entitled "Entangled Heuristics for Agent-Augmented Strategic Reasoning" [2] introduces a hybrid deductive reasoning architecture in which strategic heuristics are activated and composed through semantic entanglement. In that context, heuristics extracted from strategic texts (e.g., military treatises, corporate frameworks) are treated not as discrete rules but as semantically interdependent potentials, with transformer models employed to measure and exploit these semantic interdependencies [7]. Conflicting or complementary heuristics are not eliminated but rather fused, leading to the emergence of novel and conceptually richer heuristic syntheses. This process embodies a quasi-ontological creativity: new strategic concepts arise through heuristic entanglement rather than being selected from a fixed menu of pre-existing options.

While the philosophical stance of that project was deliberately constructivist—treating inference as a form of creation—quantum abduction maintains a stronger epistemic orientation. Here, hypotheses represent real but uncertain explanatory elements about a determinate world. The entanglement mechanism remains valuable; however, embedding hypotheses in a Hilbert-style semantic space, modeling their interferences, and enabling synthetic recombination provide the computational machinery needed to capture inference dynamics in uncertain contexts.

Quantum abduction can thus be seen as a specialized adaptation of the entangled heuristics framework:

- It inherits the machinery of semantic superposition and interference.
- But it applies them in service of *explanatory convergence with reality*, rather than conceptual innovation.

This lineage clarifies how formal resources developed for strategic deduction can be repurposed to manage the epistemic complexities of abduction—enabling explanations to interplay, synthesize, or collapse only when evidence warrants it, rather than through premature elimination.

Moreover, the complementarity between two traditions is evident: quantum cognition models context effects, superposition, and interference in human judgment, while generative AI enables compositional and narrative synthesis. Their integration paves the way for quantum-abduction engines capable of generating natural-language explanations of complex entanglements. Such explanations allow ordinary human users to make sense of situations that would otherwise appear inextricably heterogeneous or hopelessly contradictory.

Figure 1 anticipates Section 3 by contrasting eliminative abduction with superposition interference.

3. Computational Premises for Quantum Abduction

Classical abduction, as outlined in Section 2, typically assumes a fixed hypothesis set and an eliminative search process: contradictions are resolved by discarding candidates until a single explanation remains. Rather than reiterating these limitations, we now turn to the computational consequences of moving beyond them.

Quantum abduction shares the classical commitment to a determinate underlying reality—only one sequence of events has actually occurred—but departs from eliminativism in its representation of uncertainty. Hypotheses are maintained in an *epistemic superposition*, with amplitudes that evolve as evidence accumulates. Observations act as projections in a semantic Hilbert space; interactions among hypotheses are modeled through an interference matrix that captures reinforcement, contradiction, or semantic overlap. Collapse marks explanatory convergence, which may result in a dominant or hybrid explanation rather than a single survivor from a fixed menu.

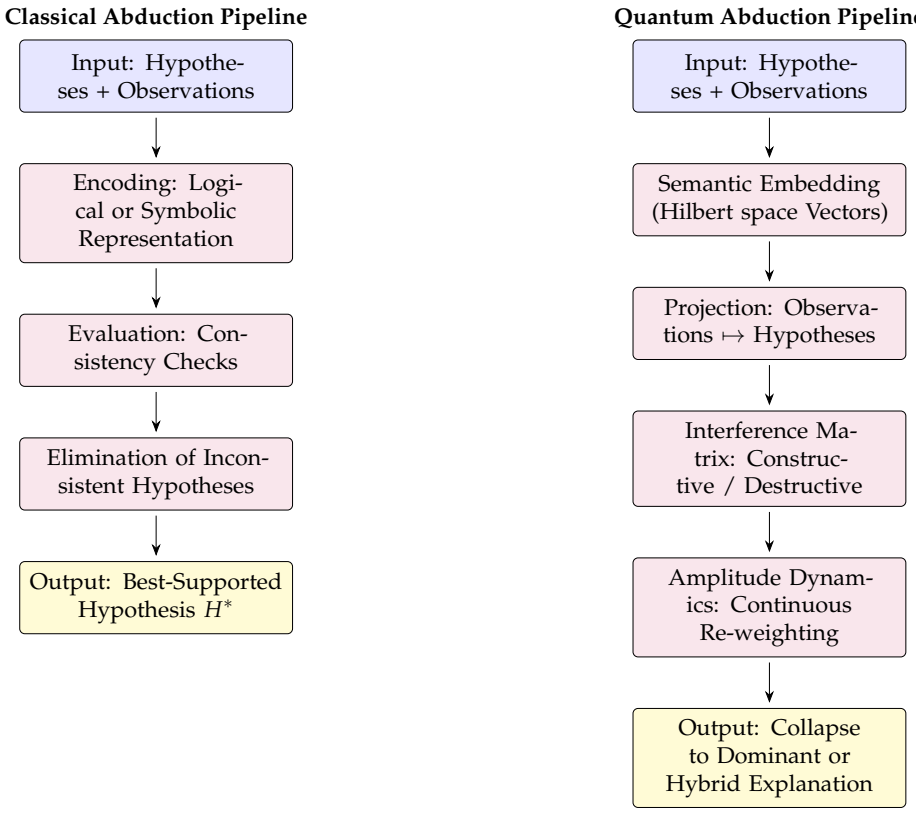


Figure 1. Computational pipelines: Classical abduction prunes the hypothesis space through elimination. Quantum abduction embeds hypotheses in a semantic Hilbert space, models interference, and converges via amplitude dynamics to a dominant or hybrid synthesis.

This section introduces the operational machinery that enables these dynamics, centered on the mathematical summary in Box 1. We treat Hilbert space tools—vectors, projections, interference—not as ontological commitments but as a computational formalism that resists premature elimination and preserves information until evidence warrants decisive convergence. Full derivations appear in Appendix B.

3.1. Why Quantum?—Interpretational Clarifications

As anticipated in the Introduction, "quantum" is used here in an *epistemic and computational* sense rather than an ontological one. Section 2 gave the conceptual motivation; here we state the implications at a more formal level and show how they guide the computational design.

Three Distinct Notions

- **Quantum cognition**: Hilbert space structure to model belief states and their context-sensitive evolution (this paper).
  - *Intuition:* A Hilbert space offers a geometric setting in which each hypothesis is a vector and each observation acts as a projection. Inner products encode explanatory alignment; interference terms capture how alternatives can reinforce or attenuate one another. This yields contextuality and order effects providing the non-classical compositionality needed for abductive reasoning under contradiction.
- **Quantum computation**: algorithms on quantum hardware (not used here).
- **Classical simulation of quantum-like dynamics**: classical hardware implementing interference/superposition mathematics (our regime).

Why the Formalism Matters Here

The formalism captures three features of explanatory reasoning that classical abduction struggles with:

1. *Interference* of explanatory alternatives, yielding non-classical updates (violations of the law of total probability when hypotheses interact);
2. *Contextuality*, where the plausibility of a hypothesis depends on which other hypotheses are co-considered;
3. *effects* in evidence evaluation, as documented in cognitive experiments [6].

Superposition and collapse are *epistemic*: coexistence of competing explanations prior to convergence; collapse as emergence of a coherent synthesis under sufficient constraints.

Relation to Entangled Heuristics

Methodologically, we share machinery with *Entangled Heuristics for Agent-Augmented Strategic Reasoning* [2] (transformers encode semantic interdependence; interference drives composition). The interpretational aim differs:

- In [2], semantic superposition supports *ontological innovation* (new strategic constructs).
- Here, superposition tracks *epistemic uncertainty about a determinate reality*.

Thus, the same tools (semantic embeddings, transformer-mediated similarity, interference dynamics) serve ontological synthesis there, and epistemic inference here.

Quantum Abduction: Core Math Capsule

Abductive state: $\left|\Psi^{(t)}\right\rangle = \sum_i \alpha_i^{(t)} |H_i\rangle, \quad \sum_i |\alpha_i^{(t)}|^2 = 1.$
Projection (evidence): $e_i^{(t)} = \langle H_i \mid O_j \rangle = \cos\left(\mathbf{h}_i, \mathbf{o}_j\right).$
Interference: $I_{ij} = \kappa_{ij} \cdot \mathrm{sim}(H_i, H_j), \ \kappa_{ij} \in [-1, 1].$
Update + normalize: $\tilde{\alpha}_i^{(t+1)} = \alpha_i^{(t)} + \eta\left(e_i^{(t)} + \sum_{k \neq i} I_{ik} \alpha_k^{(t)}\right), \ \alpha^{(t+1)} = \tilde{\alpha}^{(t+1)} / \|\tilde{\alpha}^{(t+1)}\|_2.$
Collapse when $C(\Psi) = \max_i |\alpha_i|^2 > \tau$ (single or hybrid via mix). Formal details: Appendix B, Equations (A1)–(A6).

Transformers and LLMs (Conceptual Primer)

Hypotheses and observations are expressed in natural language; we therefore need a representation that preserves semantic structure. Transformer encoders map text to high-dimensional vectors, where semantic proximity corresponds to geometric proximity. This does not entail that transformers "understand" the hypotheses; rather, they supply a *coordinate system* in which superposition and interference can be computed. LLMs, in turn, *articulate* hybrid explanations in natural language; they do not decide scores.

3.2. Transformer and LLM Integration as Epistemic Infrastructure

Embedding Model

We use the transformer model Sentence-BERT [8] (e.g., all-MiniLM-L6-v2, $d = 384$) to encode hypotheses H_i and observations O_j into a shared semantic space: $\mathbf{h}_i = \mathrm{SBERT}(\mathrm{text}(H_i))$ and $\mathbf{o}_j = \mathrm{SBERT}(\mathrm{text}(O_j))$. Results are largely invariant to the specific SBERT variant up to rotation; higher-dimensional models increase fidelity at modest cost.

Clarifying the Objects Involved

Each hypothesis H_i and observation O_j begins as a text string. These are not numerical objects. To use them computationally, we convert them into fixed-length numerical vectors using the SBERT encoder.

What SBERT Does (Intuitive Explanation)

SBERT is a transformer-based model that maps any piece of text—ranging from a single word to a short paragraph—into a vector in $\mathbb{R}^{384}$. All embeddings have the same dimensionality regardless of the input length. This uniformity is why cosine similarity between $\mathbf{h}_i$ and $\mathbf{o}_j$ is well-defined.

In principle, long hypotheses could be decomposed into clause- or sentence-level embeddings and then pooled (e.g., by weighted averaging or attention). For this initial framework we adopt the simpler and widely used "one SBERT vector per text" strategy, leaving fine-grained composition to future work.

Meaning of H_i, $\mathbf{h}_i$, and e_i

To avoid ambiguity:

- H_i = hypothesis as a linguistic description;
- $\mathbf{h}_i$ = its embedding in $\mathbb{R}^{384}$;
- O_j = observation as text;
- $\mathbf{o}_j$ = its embedding in $\mathbb{R}^{384}$.

The projection score for hypothesis i at time t is

$$e_i^{(t)} = \cos\left(\mathbf{h}_i, \mathbf{o}^{(t)}\right).$$

It has a single index because it measures the activation of one hypothesis, even though its computation involves two embeddings.

Qualitative vs. Quantitative Projection Matrices

The tables in Section 4 use qualitative marks (v = supports, x = contradicts) purely for readability. In the actual computation, each entry is replaced by a cosine similarity between SBERT embeddings, optionally modulated by an interference coefficient κ_{ij} described below.

Projection (Evidence Activation)

Similarity $e_i^{(t)} = \cos(\mathbf{h}_i, \mathbf{o}_j)$ acts as a projection operator, activating or suppressing amplitudes without forcing single-winner selection.

Interference Estimation

We set $I_{ij} = \kappa_{ij} \cdot \mathrm{sim}(H_i, H_j)$ with cosine similarity over SBERT vectors. $\kappa_{ij} \in [-1, 1]$ encodes exclusivity/complementarity and can be initialized (i) heuristically from domain priors, (ii) via weak supervision (optimize ECS/HQI on solved cases), or (iii) by expert elicitation smoothed through metric learning.

Amplitude Update (Single Step)

```
Input: alpha[1..n], I[1..n,1..n], evidence e[1..n], eta in (0,1]
for i in 1..n:
    alpha_tilde[i] = alpha[i] + eta * ( e[i] + sum_{k != i} I[i,k]
    * alpha[k] )
Normalize: alpha = alpha_tilde / ||alpha_tilde||_2
Collapse if max_i |alpha[i]|^2 > tau; else continue
```

Mix Operator and LLM Articulation

When multiple hypotheses retain high amplitudes and exhibit constructive interference, the system prepares a *candidate hybrid* by combining their embeddings via a weighted

semantic blend. The precise operator, its justification, and its role in case studies are elaborated in Section 3.3; at this point we note only that the LLM's role is strictly limited to articulating the hybrid in natural language, not to scoring or selecting it.

Human-in-the-Loop

Experts adjust κ_{ij} (conflict/compatibility), calibrate collapse thresholds τ, and edit LLM phrasing for clarity and faithfulness. This design follows a *human-in-the-loop* principle [9], in which domain experts operate at a semantic rather than mathematical level. They do not need Hilbert space or amplitude calculus; their inputs are judgments about *complementarity*, *conflict*, or *partial overlap* among hypotheses (i.e., signs and strengths for κ_{ij}, and acceptable collapse thresholds τ). Within a broader *centaurian* paradigm [10], the assistant provides formal control and interpretability over superposition/interference dynamics while preserving human agency over meaning and decision.

Full definitions, equations, and derivations are provided in Appendix B; here we retain only the operational summary.

3.3. The Mix Operator for Hypothesis Synthesis

When several hypotheses retain high amplitudes and exhibit constructive interference, the system applies the mix operator to generate hybrid explanations:

$$\text{mix}(H_i, H_j) = \lambda_{ij}\mathbf{h}_i + (1 - \lambda_{ij})\mathbf{h}_j, \qquad \lambda_{ij} = \frac{|\alpha_i|^2}{|\alpha_i|^2 + |\alpha_j|^2}. \tag{1}$$

The natural-language articulation follows the LLM strategy above. Formal variants and derivations appear in Appendix B.

Why This Matters in Practice

Retention of incompatible evidence (no premature pruning), *dynamic reweighting* as data arrives, and *hybrid resolutions* when interference is constructive are the core computational benefits.

How the Machinery Appears in Our Cases (Preview of Section 4)

Case	**Mechanism in Play**
Ludwig II	e_i from letters/reports boosts H1/H2/H3; positive I_{12} and I_{23} sustain a hybrid "entangled conflict" before collapse.
Mostro di Firenze	Stable e for weapon/M.O. raises H1/H3; vehicle variance drives $I_{12} > 0$ (imitators); new DNA toggles I_{35}.
Bossetti–Gambirasio	Strong e for nuclear DNA on H1 opposed by mtDNA/Y-line (negative I toward H1); superposition resists forced collapse.
Botulism vs.GBS/MFS	Parallel high e on H1/H2; defer collapse and allow treatment-on-superposition until decisive labs arrive.
Drift → Tectonics	Long-lived superposition H1/H2; instruments increase e on H1 and flip I as mechanism emerges, prompting synthesis.

Implementation Note

Off-the-shelf SBERT encoders provide embeddings; I_{ij} is initialized from similarity and tuned by priors or weak supervision; updates follow the capsule above. A short reference sketch and complexity notes appear in Appendix B.

3.4. Temporal Scales and Co-Opetition in Quantum Abduction

Immediate Decision Contexts

In time-constrained scenarios (such as emergency medicine, crisis response, and real-time diagnostics), experts naturally maintain multiple hypotheses in superposition, allowing for parallel consideration until evidence or urgency forces a collapse. Here,

quantum abduction serves both as a *descriptive model* of expert cognition and a *prescriptive tool* for decision support.

Extended Investigative Contexts

In long-duration investigations (complex criminal cases; scientific paradigm shifts), classical eliminativism fosters destructive competition. Quantum abduction coordinates "co-opetition":

- **Hypothesis Allocation**: distribute resources proportionally to $|\alpha_i|^2$;
- **Information Sharing**: evaluate evidence under H_i for projection onto all active hypotheses;
- **Interference Tracking**: monitor cross-effects via I_{ij} and amplitude sensitivities;
- **Synthesis Triggers**: detect constructive interference suggesting breakthrough hybrids.

Organizational Implementation

Possible forms include investigation coordination centers, research consortia with shared amplitude tracking, and multi-unit intelligence analysis with interference monitoring.

From Competition to Co-opetition

The same mathematics (α, I, τ) scales to organizational coordination. A simple value decomposition is

$$\text{Value}_{\text{co-op}} = \sum_i |\alpha_i|^2 V_i + \sum_{i \neq j} I_{ij} \alpha_i \alpha_j S_{ij}, \tag{2}$$

where V_i is the standalone value of H_i and S_{ij} is the synthetic value from interaction. The second term—absent in classical abduction—captures the added value of maintaining superposition.

4. Case Studies

We present case studies drawn from criminal investigation, clinical diagnostics, and scientific inquiry to illustrate how quantum abduction operates in the face of ambiguity, contradiction, and evolving evidence. Together, they span the full temporal range of the framework's application: from real-time decision settings—such as emergency medical management where parallel hypotheses must remain active until action is unavoidable—to long-horizon contexts like scientific research programs or decades-long criminal investigations. Across these domains, the case studies show how quantum abduction sustains multiple explanatory lines in a structured superposition, countering the eliminative bias of classical approaches, and how it enables synthesis or collapse precisely when the evidence warrants it.

4.1. Forensic Reasoning

4.1.1. Ludwig II of Bavaria

The mysterious death of King Ludwig II of Bavaria in 1886, together with his physician Dr. Gudden, is a paradigmatic case where classical abduction struggles and quantum abduction offers a richer interpretive space. The historical record presents conflicting evidence: eyewitness accounts of the bodies in shallow water, reports of blunt trauma on Gudden, Ludwig's documented political deposition and alleged insanity, and the absence of a thorough autopsy. Classically, these contradictions enforce a premature binary: *suicide* versus *murder*.

Classical framing forces "suicide vs. murder"; the superpositional lens lets both narratives coexist until constraints compel synthesis. Concretely, political orchestration and personal despair can be treated as entangled contributors rather than exclusives, so that

evidence about Gudden's trauma, deposition dynamics, and Ludwig's letters interact before any collapse. In our setting, this interaction yields an emergent H_5 (entangled conflict) that is not posited a priori but is constructed by interference among H_1–H_3; Section 5 demonstrates this effect quantitatively using randomized evidence orderings.

Furthermore, at the *hypothesis level*, quantum abduction supports hybrid explanations such as the following:

- H_1: Suicide. Explains resignation but ignores Gudden's injuries and political orchestration.
- H_2: Murder. Explains political context and Gudden's trauma but leaves Ludwig's fatalism underexplored.
- H_3: Struggle. Accounts for joint death and injuries but underplays political necessity.
- H_4: Medical accident. Considers sudden seizure or collapse but neglects political-military involvement.
- H_5: **Entangled conflict** (emergent from interference). This hypothesis is not defined a priori but emerges through the quantum abductive process from constructive interference between H_1, H_2, and H_3. It represents Ludwig resisting removal in a context where suicidal impulses, physical resistance, and political suppression became entangled, projecting onto most observations with high combined explanatory power.

Historical documents support this reframing, as illustrated in Table 1. Ludwig's letters to Richard Wagner and confidants reveal fatalistic tones and longing for escape, aligning with H_1; political records demonstrate ministerial urgency to depose him, aligning with H_2; Gudden's ambiguous role aligns with H_3; reports of seizures align with H_4. Only H_5—the entangled conflict hypothesis—integrates them without contradiction.

Table 1. Hypothesis projection matrix for Ludwig II case. Observations: O_1 = both died together, O_2 = drowning, O_3 = Ludwig did not drown, O_4 = Gudden trauma, O_5 = seizure, O_6 = dubious insanity, O_7 = political pressure. (✓ = supports hypothesis; ✗ = contradicts hypothesis.)

Hypothesis	O_1	O_2	O_3	O_4	O_5	O_6	O_7
H_1 Suicide	✓	✓	✗	✗	✗	✗	✗
H_2 Murder	✓	✗	✓	✓	✗	✓	✓
H_3 Struggle	✓	✗	✓	✓	✓	✗	✗
H_4 Medical	✓	✗	✗	✗	✓	✗	✗
H_5 Entangled conflict	✓	(✓/✗)	✓	✓	✓	✓	✓

Through quantum abduction, the Ludwig II case is reframed not as "who killed whom," but as an entangled historical event in which political orchestration, personal despair, and interpersonal struggle jointly collapsed into the observed outcome. This illustrates the unique explanatory gain of a superpositional abductive framework over classical elimination.

4.1.2. The "Monster of Florence"

The "Mostro di Firenze" killings (1968–1985) are one of the most studied yet unresolved serial crime cases in Europe, involving eight double homicides of couples near Florence. Despite decades of investigation, no definitive suspect has been convicted. This case highlights the limitations of classical abduction and the interpretive power of a quantum abductive framework.

Observations

Eight double homicides (1968–1985) with a consistent weapon/M.O. and post-mortem mutilations; heterogeneous vehicles reported near scenes; 2024 DNA update introduces an unknown external profile.

Competing Hypotheses

H_1 single serial killer; H_2 multiple actors/imitators; H_3 transient actor with wide mobility; H_4 ritual/cult network; H_5 external predator.

Quantum Reframing

Forensic consistency (O_2) boosts H_1/H_3; vehicular diversity activates H_2; the 2024 DNA boosts H_5. Constructive interference yields hybrids: $H_{1,2}$ (core killer + imitators), $H_{3,4}$ (transient acts appropriated into cultic narratives), $H_{3,5}$ (external predator intersecting local patterns).

Quantum Advantage

(i) No premature elimination of inconsistent cues; (ii) contradictions become interferential structure, not noise; (iii) composite explanations are first-class outcomes rather than afterthoughts.

Reframed Question

The quantum abductive stance shifts from "Which hypothesis explains all?" to the following:

> *"Which entangled combination of agency, imitation, mobility, and symbolic behavior coherently explains the mixture of forensic consistency and observational variability in the Florence homicides?"*

Illustrative Synthesis

A plausible abductive synthesis might state the following:

> *A dominant transient actor with access to varied vehicles carried out the core attacks, whose stylistic and symbolic violence later inspired local imitators and speculative cultic interpretations. The consistency of weapon and mutilations suggests a single origin, while the observed vehicular variance and the emergence of new DNA indicate auxiliary actors or replication effects. The "Mostro" phenomenon thus emerges less as a single killer and more as an entangled field of violence, imitation, and unresolved identities.*

Quantum Advantage

Here, quantum abduction achieves the following:

- Avoids premature exclusion of inconsistent evidence.
- Integrates contradictions as interference effects rather than eliminations.
- Supports composite explanations (killer + imitators + external unknown).

Thus, the framework captures the narrative multiplicity of the case, yielding explanatory depth beyond the eliminativist paradigm.

4.1.3. Bossetti and the Yara Gambirasio Case

The murder of Yara Gambirasio, a 13-year-old girl found dead in 2011 near Brembate di Sopra in the province of Bergamo, remains one of Italy's most controversial criminal cases. Massimo Giuseppe Bossetti was convicted in 2016 based on DNA evidence; this outcome of the criminal proceedings remains legally solid, despite numerous anomalies in the forensic chain and contradictory interpretations that continue to raise doubts about both the reliability of the evidence and the coherence of the investigative reasoning. In this case, a typical reductionist shift toward classical *eliminativist* abduction—driven by the judicial necessity of closure—was required to uphold Bossetti's conviction. Yet, as with any conviction based purely on circumstantial or interpretative evidence, shadows remain. A

quantum abductive approach, instead of enforcing premature closure, would keep these shadows in play, offering a fuller picture of the inferential landscape.

While the Bossetti case itself was not formally described in abductive terms within judicial proceedings, its evidential dynamics—mutually conflicting hypotheses, partial support, and iterative narrative reconstruction—are archetypally abductive. In this sense, our analysis is not a historical reconstruction but an abductive re-framing aimed at revealing how eliminative reasoning constrained the court's interpretive space.

Observational Contradictions

The prosecution rested on the matching of Bossetti's nuclear DNA to the so-called *Ignoto 1* profile recovered from the victim's clothing. However, several anomalies emerged:

- The **Y-chromosome haplotype** of Ignoto 1 did not match that of Bossetti's known male relatives.
- The **mitochondrial DNA** found on Yara's clothes did not match Bossetti's maternal line.
- The amount of **nuclear DNA** was unusually high, despite its well-known tendency to degrade rapidly in environmental exposure.
- Conversely, **mitochondrial DNA**, typically far more stable and abundant in degraded samples, was found only in trace quantities.

This paradox constitutes an inversion of forensic expectations: nuclear DNA, normally fragile, persisted strongly, while mitochondrial DNA, normally resilient, was nearly absent. Classical reasoning treats such anomalies as marginal noise; a quantum abductive model instead elevates them to central dimensions of explanatory tension.

Limits of Classical Abduction

Classical abductive reasoning proceeds as follows:

1. Observation: Bossetti's nuclear DNA matches the crime scene sample.
2. Inference Rule: If someone's DNA matches a crime scene sample, then they were likely present.
3. Abductive Conclusion: Bossetti was at the crime scene and likely committed the murder.

This model cannot accommodate the contradictions above. To preserve logical consistency, it must eliminate or downplay conflicting evidence—mitochondrial mismatch, degradation anomalies, or kinship inconsistencies—resulting in a form of *explanatory closure* that obscures deeper explanatory potentials.

Quantum Abductive Reframing

Quantum abduction offers a framework in which conflicting observations need not be collapsed or discarded. Instead, each hypothesis is represented as a semantic vector in a conceptual Hilbert space, and each observation as a projection operator that activates or interferes with subsets of these hypotheses.

The hypothesis space may include the following:

- H_1: Bossetti is the biological donor and the murderer.
- H_2: The donor is a relative or unregistered sibling.
- H_3: The DNA match results from laboratory contamination or error.
- H_4: Bossetti's DNA was planted or transferred indirectly.
- H_5: Degradation or environmental effects inverted nuclear vs. mitochondrial persistence.
- H_6: Amplification or measurement bias distorted the detection process.

These hypotheses form an entangled field rather than discrete alternatives. For instance, H_1 and H_2 share semantic overlap (nuclear match) but diverge on kinship

assumptions; H_1 is weakened by mitochondrial mismatch, which instead reinforces H_2 or H_4. The final explanatory synthesis emerges not from exclusion but from interference-informed composition.

Projection Matrix

To visualize this entangled reasoning, in Table 2 we map hypotheses (H_i) against key observations (O_j):

Table 2. Hypothesis projection matrix for the Bossetti–Gambirasio case. Observations: O_1 = nuclear DNA match, O_2 = Y-haplotype mismatch, O_3 = mitochondrial mismatch, O_4 = nuclear abundance anomaly, O_5 = mitochondrial scarcity, O_6 = contamination possibility, O_7 = court ruling affirming nuclear DNA integrity. (✓ = supports hypothesis; × = contradicts hypothesis.)

Hypothesis	**O_1**	**O_2**	**O_3**	**O_4**	**O_5**	**O_6**	**O_7**
H_1 Bossetti guilty	✓	×	×	×	×	×	✓
H_2 Relative involvement	✓	✓	✓	×	×	×	×
H_3 Contamination	✓	×	×	×	×	✓	×
H_4 Secondary transfer	✓	×	✓	×	✓	×	×
H_5 Degradation anomaly	×	×	✓	✓	✓	×	×
H_6 Measurement bias	✓	×	✓	✓	✓	✓	×

This matrix shows that no single hypothesis aligns with all observations. Instead, they form overlapping projections, sustaining multiple explanatory amplitudes.

Forensic Reinforcement and Recalibration

On appeal, the Court of Cassation reaffirmed the integrity of nuclear DNA, which, in our model, acts as a strong projection, increasing $|\alpha_{H_1}|^2$, while applying destructive pressure to alternatives (H_2, H_3). As Figure 2 shows, orthogonal constraints (mitochondrial and Y-line mismatches) prevent full collapse, explicitly separating *judicial closure* from *epistemic closure*.

Thus, a quantum abductive assistant might generate the following:

> *The nuclear DNA match strongly implicates Bossetti or a close relative, but mitochondrial and Y-chromosome mismatches suggest genealogical or methodological gaps. The unusual abundance of nuclear markers, compared to the scarce mitochondrial evidence, indicates either secondary transfer, selective preservation, or amplification bias. While the Cassation ruling reinforces Bossetti's responsibility, quantum abduction models this as a high-amplitude projection that suppresses but does not eliminate alternative hypotheses. The explanatory field remains entangled: identity, mechanism, and chain-of-custody processes co-determine the inferential landscape.*

This reveals the epistemic innovation of quantum abduction: it accounts not only for forensic contradictions but also for the institutional dynamics of closure, showing how alternative amplitudes persist beyond verdicts. In this way, quantum abduction spans both modern forensic science and historical mysteries (as in the Ludwig II case), demonstrating its scope across contexts.

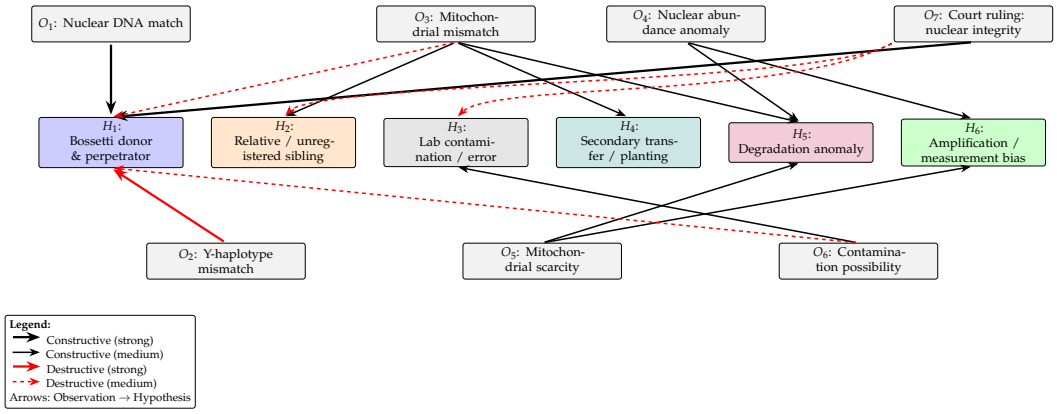


Figure 2. Bossetti–Gambirasio: split-layout interference map. Hypotheses occupy the central band; observations sit above/below. Solid black arrows indicate constructive activation; solid/dashed red arrows indicate destructive influence (strong/medium). Line thickness encodes salience. Strong projections from O_1 and O_7 amplify H_1, while O_2, O_3, and O_6 apply destructive pressure to H_1 and activate alternatives (H_2, H_3, H_5, H_6).

4.2. Literary Demonstration: Murder on the Orient Express

Before turning to scientific and clinical applications, we round up the forensic section above with a literary case—Agatha Christie's *Murder on the Orient Express*—itself inspired by the real kidnapping and murder associated with the Armstrong/Lindbergh tragedy. The novel is not merely a puzzle; it is a precise dramatization of abductive structure. Detective Hercule Poirot recognizes that the binary framing "*which one passenger is guilty?*" collapses the space of explanations too early. Rather than eliminate suspects one by one, he keeps all viable stories in play until the evidence compels a synthesis: the disjunction of suspects resolves into a coordinated conjunction. In other words, Poirot behaves as a *quantum investigator*: hypotheses remain in superposition, interact via thematic constraints, and finally collapse to a collective explanation.

4.2.1. Plot in Brief

The train is snowbound; an American traveling under an alias (Ratchett) is found stabbed. The compartment contains a scatter of contradictory clues (a monogrammed handkerchief, a pipe cleaner, a burnt note, a red kimono, a spare conductor's uniform), and several alibis partially corroborate one another. Poirot infers that Ratchett is in fact Cassetti, the gangster behind the Armstrong child murder. Every principal passenger (and the conductor) turns out to be directly connected to that crime. Two explanations are then staged: (i) a lone intruder escaped in the night; (ii) the passengers, acting as a self-constituted jury, each delivered a stab, distributing guilt and masking individual responsibility. Poirot presents both "solutions", but only the second integrates all clues *without contradiction*.

4.2.2. Classical vs. Quantum-Abductive Reading

Classical abduction would select a single culprit (*OR* over suspects), pruning away incompatible threads. Yet the evidence is purpose-built to defeat eliminativism: each clue simultaneously implicates and misdirects. Poirot's procedure maintains a superposition of individually plausible narratives that, under further constraints (shared motive, synchronized opportunity, mutually supporting alibis), *interfere constructively* into a collective act (*A*ND). The final "collapse" is a hybrid: distributed agency coherently explaining contradictory traces (multiple wound types, planted items, cross-checked alibis). Thus, the case exemplifies the core move of quantum abduction: from selection to composition. The two readings are represented and contrasted in Figure 3.

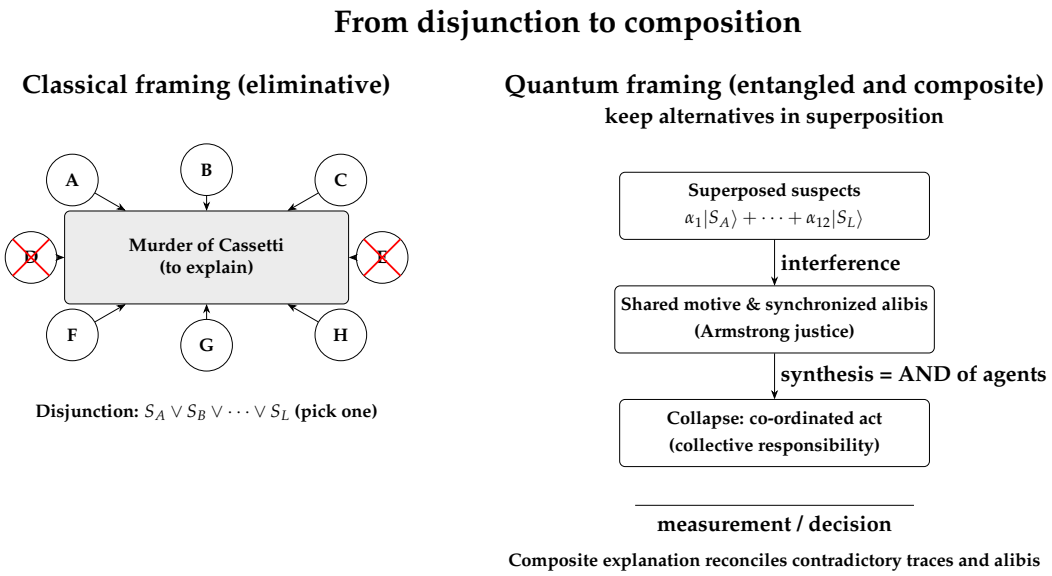


Figure 3. From classical eliminative framing (**left**), where suspects are treated as a disjunction and progressively discarded, to quantum framing (**right**), where suspects remain in superposition, interfere via shared motive and synchronized alibis, and collapse to a coordinated act that reconciles contradictory traces and alibis. (Red crosses indicate suspects eliminated in the classical process.)

4.3. Medical Diagnostics: Botulism vs. GBS/MFS

Having demonstrated quantum abduction in action in a literary case involving multiple culprits, we now turn to the clinical and scientific domains, where the integration of alternative explanatory states can be crucial to saving lives or upending scientific paradigms. A particularly illustrative case is that of a "suspended" diagnosis already discussed in [11], where the challenge of balancing competing hypotheses under uncertainty foregrounds the abductive dimension of reasoning.

Provenance. The clinical vignette is adapted from a real diagnostic narrative to illustrate the abductive process; it does not report a live deployment of the assistant. The reconstruction foregrounds how superposition, order effects, and deferred collapse align with expert practice in time-constrained care.

The episode involved a young man admitted to the emergency room with rapidly progressing paralysis. The clinical team (and the AI assistant replicating their reasoning) had to decide between two life-threatening but distinct conditions: *botulism* and *Guillain–Barré syndrome* (GBS) or its *Miller–Fisher* variant (MFS). Each diagnosis implied radically different therapeutic strategies: antitoxin administration for botulism versus immunoglobulin or plasma exchange for GBS/MFS. Early test results, cranial nerve involvement, and autonomic instability produced a contradictory picture, precluding definitive commitment to either path.

4.3.1. Classical Abductive Limitation

Under a classical abductive frame, clinicians would be pressed to select the most plausible hypothesis (H_1: Botulism or H_2: GBS/MFS) as "the explanation" for the observed syndrome. This eliminativist move, however, risks therapeutic error: a premature collapse toward the wrong diagnosis may fatally delay the correct treatment.

4.3.2. Quantum Abductive Reframing

Quantum abduction instead maintains both hypotheses in superposition:

$$\Psi = \alpha|H_1\rangle + \beta|H_2\rangle,$$

where amplitudes evolve as new evidence projects onto the diagnostic state. Contradictory test results are not discarded but interfere with the relative weight of competing explanations. This allows clinicians to act on the entangled state itself: a "parallel treatment" strategy can be pursued—administering both antitoxin and IVIG/plasma until further observations collapse the superposition toward a dominant diagnosis.

4.3.3. Clinical Alignment

In the real case [11], this is exactly what occurred: as represented in Figure 4, the medical team suspended binary decision-making, treating for both conditions in parallel until the confirmatory test results clarified the diagnosis. The AI assistant, reasoning abductively, mirrored this strategy in real time, demonstrating both the plausibility and utility of quantum abductive framing in clinical settings.

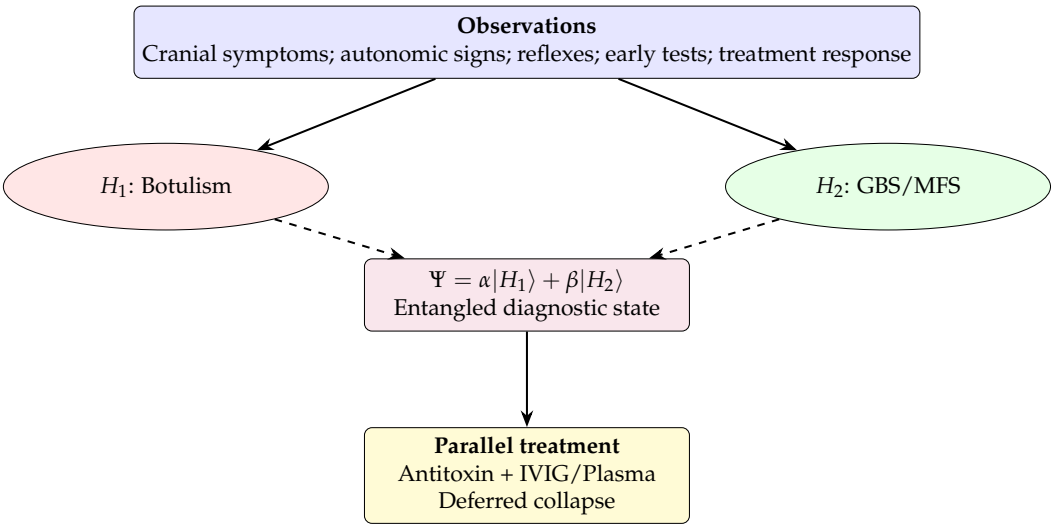


Figure 4. Quantum abductive diagnosis sustains co-activation and permits parallel, low-regret therapy before evidential collapse. (Solid arrows: evidence projection; dashed arrows: contribution to the entangled state.)

This medical scenario exemplifies how quantum abduction manages uncertainty: not by suppressing contradictions but by harnessing them to sustain a safe and generative diagnostic trajectory until reality itself enforces collapse.

4.4. Scientific Theory Change

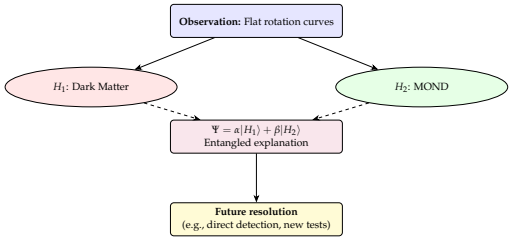
Scientific inquiry is a fertile domain for quantum abduction. Unlike courtrooms or clinics, science rarely resolves hypotheses quickly: competing explanatory frameworks can coexist for decades, each with partial explanatory successes, until new instruments or data shift the balance. In this sense, scientific progress often exemplifies abductive "superposition," with explanatory collapse delayed until decisive evidence emerges. Two emblematic cases are astrophysical dark matter versus MOND, and Wegener's continental drift versus geological fixism.

4.4.1. Astrophysics: Dark Matter vs. MOND

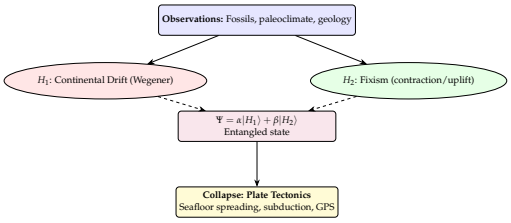
One of the most visible examples of suspended explanatory alternatives in modern science is the ongoing debate between *dark matter* and *modified Newtonian dynamics (MOND)*. Both are attempts to explain why galactic rotation curves remain flat rather than decaying at large radii.

- H_1: Invisible non-baryonic matter pervades galaxies, altering their gravitational profiles.
- H_2: Newton's law of gravitation requires modification at very low accelerations.

For decades, neither model has been definitively falsified. Dark matter has accumulated supporting evidence from cosmology (CMB anisotropies, lensing), but direct detection experiments remain null [12,13]. MOND captures empirical regularities in galactic dynamics but resists integration into relativistic cosmology. Both thus persist in a suspended explanatory state, as represented in Figure 5a, their amplitudes oscillating with new experimental results.



(**a**) Astrophysics: Dark Matter vs. MOND



(**b**) Geology: Drift vs. Fixism

Figure 5. Scientific superposition in two domains. (**a**) Astrophysics: dark matter and MOND coexist in an entangled explanatory state until decisive evidence induces collapse. (**b**) Geology: Wegener's continental drift and fixist models remain in superposition until the mechanism of plate tectonics emerges, yielding a richer synthesis. (Solid arrows: evidence projection; dashed arrows: contribution to the entangled state).

4.4.2. Geology: Continental Drift → Plate Tectonics

Another striking example of quantum abductive dynamics is the history of *continental drift*. In 1912, meteorologist Alfred Wegener proposed that continents once formed a single landmass (Pangaea) and had since drifted apart. His evidence included fossil correlations across oceans, paleoclimatic traces (e.g., glacial striations in now-tropical zones), and geological continuities between Africa and South America.

Yet geology at the time was dominated by *fixist* paradigms: mountains were thought to arise from the Earth's contraction as it cooled, and continents were assumed immobile. As a result, Wegener's hypothesis was widely ridiculed, mainly because he lacked a plausible mechanism to explain how continents could move.

For half a century, the scientific community lived in a quantum-like suspension:

- H_1: Continents drift (Wegener's view).
- H_2: Continents are fixed; apparent similarities arise from chance or bridges (fixism).

The wave function did not collapse until the 1960s, when new instruments revealed seafloor spreading, subduction zones, and the dynamics of the lithosphere. With plate tectonics [14], the two alternatives merged into a richer synthesis, represented in Figure 5b: Wegener was right about drift, but wrong about the mechanism. His bold intuition thus collapsed into modern geology's most unifying paradigm (this example illustrates a central feature of quantum abduction: instead of enforcing early elimination of hypotheses—as classical abductive frameworks often do—the model allows suspended coexistence until decisive evidence becomes available. Historically, continental drift was rejected not because it failed to explain observations, but because it lacked a mechanism; quantum abduction formalizes precisely this long-lived state of justified non-elimination).

Wegener's own fate underlines the personal stakes of abductive suspension. A passionate polar explorer, he died in 1930 on the Greenland ice sheet during a resupply expedition. His body, buried under drifting snow, was found the next year. In a sense, his theory too lay buried for decades, preserved in superposition, until new data revived and vindicated it.

4.5. Classical vs. Quantum Abduction in Scientific Reasoning

The three cases above—clinical diagnosis, astrophysical theory choice, and geological paradigm change—illustrate how classical abduction and quantum abduction diverge in practice. Classical approaches push toward premature elimination, whereas quantum abduction sustains superposition until decisive evidence or new mechanisms emerge. To highlight these contrasts across domains, Table 3 summarizes the key differences in abductive dynamics.

Table 3. Classical vs. Quantum Abduction in Scientific Reasoning

Case Study	Classical Abduction	Quantum Abduction
Medical Diagnosis (Botulism vs. GBS)	Seek the single best diagnosis; early resolution; risk of excluding valid alternatives.	Maintain co-activation; parallel low-regret therapy; collapse on decisive evidence.
Astrophysics (Dark Matter vs. MOND)	Commit to the dominant paradigm; treat the other as fringe.	Keep both entangled; design tests targeting interference zones; defer collapse.
Geology (Drift vs. Fixism)	Reject Drift for lack of mechanism; persist with Fixism.	Sustain superposition until the mechanism appears (tectonics); then collapse.

4.6. Key Takeaway: The Usefulness of Pre-Collapse

In high-stakes settings, the *pre-collapse state* is not a temporary inconvenience but a valuable analytical product. While the case studies in the previous section focus on law, medicine, and scientific theory change, the same dynamics extend naturally to domains such as intelligence analysis and emergency management. In every setting, interference patterns reveal where evidence fails to discriminate, and amplitude trends indicate when to act, what to collect next, and how to document uncertainty with transparency.

4.6.1. Law (Reasonable Doubt)

A distribution such as $|\alpha_{H_1}|^2 = 0.55$ (guilt), $|\alpha_{H_2}|^2 = 0.35$ (alternative), and an emergent $H_3 = 0.10$ identifies evidential gaps, clarifies burdens, supports proportional or qualified verdicts, and provides a structured trajectory for appellate review (cf. Bossetti, Section 4.1).

4.6.2. Intelligence (Collection Planning)

Amplitudes over H_1 (centralized), H_2 (autonomous), and H_3 (hybrid) guide resource allocation proportionally to $|\alpha_i|^2$. Positive I_{12} highlights the need for discriminating evidence (e.g., communication topology), while growth in α_{H_3} suggests preparing hybrid-network models.

4.6.3. Medicine (Parallel Treatment)

When botulism and GBS/MFS maintain high but mutually destructive amplitudes, pre-collapse supports low-regret dual therapy, provides an ethical and administrative rationale for resource use, and indicates when commitment becomes safe as amplitudes diverge (see Figure 4).

4.6.4. Science (Avoid Premature Closure)

Amplitude tracking for drift vs. fixism (historically) or dark matter vs. MOND (today) highlights which evidence classes are decisive and which are merely reinforcing, converting rivalry into targeted exploration rather than premature paradigm commitment.

4.6.5. Crisis Response (Public Safety)

Under ambiguous causes, amplitudes justify graduated protocols—e.g., simultaneous HAZMAT deployment, security lockdown, and meteorological monitoring—and provide a clear after-action trail explaining branch prioritization.

4.6.6. Institutional Value

Pre-collapse visibility supports cognitive load sharing, expert training, procedural auditing, and public communication by making uncertainty explicit, structured, and operational rather than hidden or improvised.

5. Benchmarking on the Ludwig II Case

The purpose of this section is not to provide large-scale benchmarking, which is outside the scope of the present paper, but to demonstrate that quantum abduction (QA) behaves competitively and distinctively when compared to classical abduction methods on a concrete historical case characterized by contradictory evidence and competing explanatory lines. We use the death of King Ludwig II of Bavaria (1886) as a compact abductive problem: the evidence is sparse, fragmentary, and contested, and multiple explanations remain historically defensible.

5.1. Case Setup

We consider four hypotheses:

- **H1**—Suicide (self-drowning).
- **H2**—Homicide (killing by guards or political actors).
- **H3**—Struggle or manslaughter (accidental drowning following altercation).
- **H4**—Medical episode leading to drowning.

We use seven evidence items $\mathcal{O} = \{O_1, \ldots, O_7\}$, expressed as natural-language snippets rather than categorical codes:

- O_1: Autopsy reports consistent with drowning.
- O_2: Bruising or struggle marks reported in some documents.
- O_3: Letters and remarks interpreted as suicidal ideation.
- O_4: Context of political removal and active opposition.
- O_5: Witness statements are inconsistent or contradictory.
- O_6: Prior episodes suggestive of mental health instability.
- O_7: Conflicting timelines regarding guards' proximity and reaction.

Hypotheses and evidence are encoded using a Sentence-BERT model [8], producing vectors $\mathbf{h}_i$ and $\mathbf{o}_j$ in a shared semantic space.

5.2. Methods Compared

We compare three approaches.

Quantum Abduction (QA)

We initialize a uniform abductive state $\left|\Psi^{(0)}\right\rangle = \frac{1}{\sqrt{4}} \sum_i |H_i\rangle$ and update amplitudes using the evidence projection and interference update rule described in Section 3.2. Interference signs κ_{ij} encode known incompatibilities or partial compatibilities (e.g., H_1 vs. H_2 is strongly exclusive; H_2 and H_3 share structural features and allow constructive interference). Collapse occurs when $\max_i |\alpha_i|^2 > \tau$ (we use $\tau = 0.8$).

Logic-Based Abduction (L-ABD)

We construct a small set of transparent rules:

$$\{O_3, O_6\} \Rightarrow H_1, \quad \{O_4, O_2\} \Rightarrow H_2, \quad \{O_2, O_1\} \Rightarrow H_3, \quad \{O_1\} \Rightarrow H_4 \text{ (fallback)}.$$

Hypotheses are scored by rule-support count, with ties broken by parsimony (least assumptions) and causal directness.

Bayesian Abduction (B-ABD)

We define a small Bayesian network with priors $p(H_1) = 0.25, p(H_2) = 0.30$, $p(H_3) = 0.25$, $p(H_4) = 0.20$ and hand-elicited likelihoods $p(O_j \mid H_i)$ reflecting historically plausible support patterns. Posteriors are computed by Bayes' rule, assuming conditional independence of evidence ($p(O_{1:m} \mid H) = \prod_j p(O_j \mid H)$).

5.3. Evaluation Protocol

We study two regimes:

R1: Full-Evidence Evaluation. We supply all 7 evidence items, but in *random order*. For each method, we compute the top-ranked hypothesis across 50 random evidence permutations. This probes order sensitivity. Classical abduction and Bayesian inference are order-invariant; QA is not, by design, since order encodes reasoning context.

R2: Perturbation Test. We remove O_2 (struggle indications) and repeat the experiment. This tests robustness under contradictory or missing evidence.

We report the following:

- **Top-1 distribution**: proportion of runs in which each hypothesis is selected.
- **ECS (QA only)**: Explanatory Coherence Score ECS $= \bar{e} - \left|\overline{\min(I, 0)}\right|$, measuring net support vs. destructive interference.
- **Calibration proxy (QA)**: $\max_i |\alpha_i|^2$ at collapse.
- **Brier score (Bayes)**: posterior concentration vs. its own top prediction.

5.4. Results

The results of the benchmark are summarized in Table 4. The full benchmark implementation and reproducibility instructions are available at https://github.com/Aribertus/qa-poc (accessed on 3 December 2025). All experiments were run on a local Intel i7-8700 CPU machine with 16 GB RAM and no GPU acceleration. We report results from 50 randomized evidence orderings in two regimes: R1 (full evidence) and R2 (removal of O_2, which encodes struggle indication). Unless otherwise noted, embeddings use Sentence-BERT (`all-MiniLM-L6-v2`) and κ_{ij} is initialized via sign-based exclusivity priors and smoothed by metric similarity.

Table 4. Ludwig II mini-benchmark. QA = quantum abduction (ours); L-ABD = rule-based abductive reasoning; B-ABD = Bayesian abduction. Top-1 values are fractions over 50 random evidence orders.

	QA (Ours)	L-ABD	B-ABD
Top-1 (R1)	29/50 (0.58)	0/50 (0.00)	50/50 (1.00)
Top-1 (R2)	43/50 (0.86)	0/50 (0.00)	50/50 (1.00)
ECS (QA, R1)	0.176	–	–
Conf (QA, R1)	0.374	–	–
Brier (Bayes, R1)	–	–	0.083

Interpretation.

Under full evidence (R1), QA splits between H_2 and H_3 depending on order (*table reports only the H_2 Top-1 count:* 29/50), with moderate confidence ($\max_i |\alpha_i|^2 = 0.374$ on

average) and positive coherence (ECS = 0.176). L-ABD does not select H_2 in any run (table: 0/50), while the Bayesian baseline (B-ABD) consistently selects H_2 (50/50), reflecting likelihood dominance.

Under perturbation (R2), where O_2 is removed, QA shifts toward H_2 (43/50 in the H_2 column), while L-ABD still does not select H_2 and B-ABD remains at H_2 (50/50). Overall, QA manages uncertainty gracefully (maintaining superposition until constraint accumulates), whereas the classical baselines either hard-select or revert to priors.

5.5. Summary

This compact benchmark shows that quantum abduction is as follows:

- Competitive with classical abductive methods under fully informative conditions;
- More robust when evidence is contradictory or incomplete;
- Uniquely capable of modeling order/context effects and hybrid explanatory outcomes.

A full-scale benchmark suite (multiple domains, larger hypothesis sets, robustness to noise, expert-elicited interference structures, and human evaluation of hybrid explanations) is left for future work, as discussed in Section 7.

6. Related Work

Abduction has been studied in artificial intelligence and cognitive science for decades, producing distinct schools of thought that differ in their formalisms and computational commitments. Broadly, three main traditions can be identified: *logic-based*, *Bayesian*, and *set-covering* approaches. Each has made important contributions, from rigorous proof systems to probabilistic reasoning and diagnostic optimization. Yet, as will be shown, they share a common eliminativist orientation: explanations are generated, tested, and ultimately reduced to a single "best" survivor. This contrasts with the framework of *quantum abduction*, which emphasizes coexistence, interference, and synthesis of hypotheses. Below, we review these traditions before highlighting how quantum abduction generalizes beyond their limitations.

6.1. Logic-Based Approaches

Within artificial intelligence and computational logic, abduction has traditionally been cast as the task of extending a background theory with explanatory assumptions such that an observation is entailed while consistency is preserved. Proof-theoretic methods such as sequent calculus and semantic tableaux provided sound and complete procedures for first-order logic [15,16]. These methods gave abduction a rigorous logical footing, but enforced eliminativist dynamics: candidate explanations were pruned until only one minimal, consistent explanation remained.

Abductive Logic Programming (ALP) further extended this tradition by embedding abduction into logic programming with abducible predicates and integrity constraints [17,18]. ALP offered practical implementations and found application in diagnostics and reasoning under incomplete information. Yet, it retained the same structural limits: explanations were selected from a fixed hypothesis space, contradictions were excluded rather than explored, and novelty was absent. In this sense, logic-based approaches, while formally elegant, struggled to accommodate the richness and ambiguity of real-world explanatory contexts.

6.2. Bayesian Approaches

Probabilistic extensions reframed abduction in terms of likelihood estimation. Ishihata and Sato [19] introduced *statistical abduction* using Markov Chain Monte Carlo to evaluate posterior probabilities of explanations in probabilistic logic frameworks. Raghavan [20] proposed BALP, combining Bayesian Logic Programs with abductive reasoning, enabling

structured logical representations of explanations to coexist with Bayesian inference. These methods elegantly handled uncertainty and provided robust inferential machinery for domains such as plan recognition and natural language understanding.

Yet Bayesian approaches treat hypotheses as probabilistically weighted *alternatives* to be ranked and selected. They remain tied to fixed hypothesis spaces and ultimately aim at a Most Probable Explanation (MPE). In contrast, quantum abduction treats hypotheses as interfering amplitudes, allowing contradictory or partial explanations to co-shape one another and even fuse into hybrids that transcend the original hypothesis set. Thus, while Bayesian methods broaden abduction with uncertainty modeling, they remain eliminative in orientation.

6.3. Set-Covering Approaches

Set-covering abduction, pioneered by Reggia and colleagues [21–23], frames explanation as a combinatorial optimization task: find minimal sets of causes that "cover" observed effects. Widely applied in medical diagnosis and fault detection, this approach is computationally tractable and intuitively appealing in domains with clear causal mappings. Variants such as Parsimonious Covering Theory further emphasized minimality as a criterion of explanatory adequacy.

However, set-covering remains committed to binary inclusion/exclusion of hypotheses and relies on predefined cause-effect mappings. It cannot represent ambiguous or overlapping causation, nor can it preserve contradictory evidence for later synthesis. By enforcing parsimony, it often discards elements that—within a quantum abductive framework—would be retained as contributors to interference patterns, allowing richer composite explanations to emerge.

6.4. Integrative Assessment

Across these traditions, abduction is implemented in structurally different ways. To enable their direct comparison as in Table 5, we assess them along four shared dimensions:

- **Hypothesis representation:** whether hypotheses are treated as discrete alternatives or as jointly coexisting explanatory states.
- **Conflict handling:** whether contradictions are eliminated, suspended, or allowed to interact.
- **Update dynamics:** how new evidence changes explanatory states.
- **Outcome structure:** whether reasoning terminates in selection, ranking, or synthesis.

Table 5. Unified comparison across abductive paradigms. Classical approaches select among hypotheses, while quantum abduction allows coexistence and interaction prior to collapse.

Dimension	Logic-Based	Bayesian	Set-Covering	Quantum Abduction (Ours)
Hypothesis representation	Discrete alternatives	Weighted alternatives	Minimal covering sets	Coexisting amplitudes (superposition)
Conflict handling	Eliminated by consistency	Resolved by posterior dominance	Removed via minimality constraints	Retained with constructive/destructive interference
Update dynamics	Rule application and pruning	Likelihood reweighting	Constraint-/cost-based elimination	Amplitude rotation and phase interaction

Table 5. *Cont.*

Dimension	Logic-Based	Bayesian	Set-Covering	Quantum Abduction (Ours)
Outcome structure	Single minimal explanation	Most probable explanation (MPE)	Parsimonious cover	Collapse to dominant or *hybrid* explanation

6.5. Alternative and Hybrid Approaches

Beyond the three main paradigms reviewed above—logic-based, Bayesian, and set-covering—other research lines approximate abductive reasoning through mechanisms that relax bivalence, incorporate argumentative dynamics, or combine neural and symbolic representations. Although these approaches introduce valuable flexibility, they remain limited in their ability to model superposition, interference, and synthesis within a unified formal framework.

Fuzzy Logic Systems.

Fuzzy logic provides a graded alternative to classical deduction by replacing Boolean truth values with degrees of membership [24,25]. In abductive contexts, fuzzy inference allows for partial matching between observations and hypotheses, accommodating uncertainty in linguistic and perceptual descriptions. However, fuzzy systems operate on scalar degrees of truth and lack the vectorial and phase-sensitive interactions that characterize quantum interference. As a result, while fuzzy logic captures *vagueness*, quantum abduction captures the *interaction* among competing explanatory tendencies.

Argumentation Frameworks.

Another major non-classical approach involves formal argumentation and defeasible reasoning. Dung's abstract argumentation framework [26] and its numerous descendants model the interplay of conflicting arguments through attack and defense relations, enabling structured deliberation and justification. These systems effectively represent contradiction and defeat but typically resolve conflict through the selection of stable or preferred extensions. Quantum abduction, by contrast, maintains superposed explanatory states in which conflicting hypotheses may coexist and influence one another until contextual constraints induce collapse, thus supporting a non-eliminative form of reasoning.

Neural-Symbolic Integration.

A third trajectory aims to bridge symbolic and sub-symbolic inference through neural-symbolic computation [27,28]. These frameworks embed logical rules within neural architectures, enabling differentiable reasoning under uncertainty. While this integration has proven powerful for knowledge representation and learning, it remains primarily a *fusion architecture*: logic guides neural adaptation, but the symbolic layer itself does not exhibit quantum-like contextuality or interference. Quantum abduction complements neural-symbolic integration by providing a principled mathematical model—rooted in Hilbert space semantics—for managing contradictory symbolic information within continuous neural embeddings.

Summary.

In summary, these alternative approaches enrich the abductive landscape by addressing uncertainty, conflict, and hybrid reasoning. Yet none explicitly capture the non-classical compositionality central to human explanatory cognition—the coexistence and dynamic blending of partial hypotheses. Quantum abduction provides a coherent formalism that encompasses these features within a single, interference-enabled representational framework.

7. Conclusions, Outlook, and Future Work

This paper has presented *quantum abduction* as a coherent extension of abductive reasoning beyond the eliminative logics that dominate classical, Bayesian, and set-covering paradigms. Where those traditions converge on ranking and discarding hypotheses, quantum abduction preserves them in a dynamic superposition that allows constructive and destructive interference, culminating in collapse only when contextual evidence supports a coherent, possibly hybrid, synthesis.

The framework thus reinterprets explanation as a process of *synthesis through interaction* rather than *selection through exclusion*. It accounts for how human reasoning sustains ambiguity, navigates contradiction, and gradually achieves insight through the tension of partially competing interpretations. At the computational level, this paradigm integrates semantic embeddings, projection operators, and interference matrices into an amplitude-based dynamics capable of representing explanatory coexistence.

Conceptually and architecturally, this work stands as a step within a broader *research program* that seeks to overcome eliminativism in reasoning—across logic, cognition, and artificial intelligence. Quantum abduction offers not a speculative draft but an operational framework that is already formalized at the representational level and computationally implemented in proof-of-concept systems. Its orientation is *centaurian* [9,10]: human experts remain within the loop, guiding semantic alignment and interpretive synthesis, while the computational substrate maintains the full superpositional structure inaccessible to unaided intuition. In this sense, the framework realizes a concrete model of human–AI complementarity, aligning with emerging paradigms of agentic and hybrid intelligence.

7.1. Current Scope and Limitations

The evaluation presented here is illustrative rather than exhaustive. Benchmarks such as the Ludwig II study demonstrate the feasibility of the approach but do not yet provide statistical generalization. Interference coefficients κ_{ij} are set heuristically; large-scale learning from empirical data remains to be achieved. Likewise, while prototype implementations exist, no controlled human studies have yet quantified improvements in interpretive accuracy or decision quality. Finally, computational cost grows with the dimensionality of the hypothesis space, motivating research into sparse and approximate formulations.

7.2. Future Directions

Future work will extend the current framework along the following three convergent axes:

1. **Formal and Logical Development.** We will elaborate the proof-theoretic and model-theoretic foundations of quantum abduction, establishing soundness and completeness properties for amplitude-based reasoning. This involves importing techniques from quantum logic and category theory to represent composition, interference, and collapse within a unified semantic calculus.
2. **Computational Expansion.** A scalable software library and benchmark suite will be released as part of the ongoing open-source program. Optimizing κ_{ij} through learning, expanding to higher-dimensional embeddings, and integrating retrieval-augmented generation (RAG) pipelines will support domain adaptation across forensic, clinical, and scientific settings.
3. **Human-in-the-Loop Validation.** Controlled studies with experts will test how the quantum abductive assistant influences reasoning transparency, confidence calibration, and interpretive synthesis. These experiments will advance hybrid decision frameworks where collapse is a collaborative, not unilateral, outcome.

7.3. Programmatic Outlook

The quantum abductive approach contributes to the ongoing effort to construct reasoning systems that do not merely simulate deductive or statistical inference, but actively support the generative and exploratory dimensions of human thought. Its integration of symbolic clarity, probabilistic nuance, and superpositional synthesis marks a shift from the competition of explanations to their *entanglement*. In this sense, quantum abduction represents both a conceptual unification and a methodological bridge—one that connects the epistemic pluralism of human inquiry with the algorithmic precision of modern AI.

Funding: Remo Pareschi has been funded by the European Union–NextGenerationEU under the Italian Ministry of University and Research (MUR) National Innovation Ecosystem grant ECS00000041-VITALITY–CUP E13C22001060006.

Data Availability Statement: The code and data for the proof-of-concept implementation (Ludwig II benchmark) described in this manuscript are available in the public repository https://github.com/Aribertus/qa-poc (accessed on 3 December 2025).

Acknowledgments: The author is grateful to Hervé Gallaire and the reviewers for comments and suggestions which have been of great help in improving this article. While preparing this manuscript, the author used ChatGpt 5 and DeepSeek to guide him in LaTeX formatting, particularly in using the TikZ package for creating graphs, as well as in the polishing of the text, aiming to make it as fluent and communicable as possible. The author reviewed and edited the resulting output and takes full responsibility for the content of this publication.

Conflicts of Interest: The author declares no conflicts of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

Appendix A. Glossary of Key Terms

Table A1. Key terms and concepts in quantum abduction

Term	Definition
Transformer	A neural architecture for encoding sequences into context-sensitive vector representations. Used here to embed hypotheses and observations in a shared semantic space.
LLM (Large Language Model)	A generative model producing natural-language text. In quantum abduction, used only for *articulating* hybrid hypotheses, not for scoring or selecting them.
Superposition	The coexistence of multiple hypotheses as weighted amplitudes in the explanatory state, rather than treating them as mutually exclusive alternatives.
Interference	The phenomenon whereby hypotheses can reinforce (constructive) or diminish (destructive) each other's explanatory power through semantic interaction.
Amplitude (α_i)	The complex coefficient associated with hypothesis H_i in the superposition, with $\|\alpha_i\|^2$ representing its relative weight.
Collapse	The convergence of the superposed state to a dominant explanation or synthesized hybrid when coherence exceeds a threshold.
Projection	The semantic alignment between an observation and a hypothesis, computed as cosine similarity in the embedding space.

Table A1. *Cont.*

Term	Definition
Entanglement	In our context (distinct from physics), the semantic interdependence between hypotheses where their meanings and explanatory power shift based on co-activation.
Mix Operator	The mathematical function that combines high-amplitude hypotheses into hybrid explanations through weighted semantic blending.
Coherence	A measure of how concentrated the explanatory state is, typically $C(\Psi) = \max_i \lvert\alpha_i\rvert^2$.
Epistemic vs. Ontological	Our superposition is epistemic (about our knowledge/uncertainty), not ontological (about reality creating multiple worlds).
κ_{ij}	Interference coefficient encoding domain-specific interaction strength between hypotheses i and j.

Appendix B. Formal Framework and Derivations

Appendix B.1. State, Projection, and Interference

State Representation

$$\left|\Psi^{(t)}\right\rangle = \sum_{i=1}^{n} \alpha_i^{(t)} |H_i\rangle, \qquad \sum_{i=1}^{n} |\alpha_i^{(t)}|^2 = 1 \tag{A1}$$

Semantic Embedding

$$\mathbf{h}_i, \mathbf{o}_j \in \mathbb{R}^d, \qquad \mathbf{h}_i = \text{SBERT}(\text{text}(H_i)), \quad \mathbf{o}_j = \text{SBERT}(\text{text}(O_j)) \tag{A2}$$

Projection Operator

$$\langle H_i | O_j \rangle = \cos(\mathbf{h}_i, \mathbf{o}_j) = \frac{\mathbf{h}_i \cdot \mathbf{o}_j}{\|\mathbf{h}_i\| \|\mathbf{o}_j\|} \tag{A3}$$

Interference Matrix

$$I_{ij} = \kappa_{ij} \cdot \text{sim}(H_i, H_j), \qquad \kappa_{ij} \in [-1, 1] \tag{A4}$$

Appendix B.2. Amplitude Dynamics and Coherence

Update Rule and Normalization

$$\tilde{\alpha}_i^{(t+1)} = \alpha_i^{(t)} + \eta \left[e_i^{(t)} + \sum_{k \neq i} I_{ik} \alpha_k^{(t)} \right], \qquad \alpha^{(t+1)} = \frac{\tilde{\alpha}^{(t+1)}}{\|\tilde{\alpha}^{(t+1)}\|_2} \tag{A5}$$

Coherence and Collapse

$$C(\Psi) = \max_i |\alpha_i|^2, \qquad \text{collapse if } C(\Psi) > \tau \text{ or when a decision is required.} \tag{A6}$$

Appendix B.3. Mix Operator and Synthesis

Mix Operator

$$\text{mix}(H_i, H_j) = \lambda_{ij} \mathbf{h}_i + (1 - \lambda_{ij}) \mathbf{h}_j, \qquad \lambda_{ij} = \frac{|\alpha_i|^2}{|\alpha_i|^2 + |\alpha_j|^2} \tag{A7}$$

For natural language articulation, we utilize template-based synthesis for simple merges and LLM-based critique to refine complex hybrids.

Appendix B.4. Implementation Sketch and Complexity

Reference Sketch (pseudo-code).

```
Input: alpha[1..n], I[1..n,1..n], evidence e[1..n], eta in (0,1]
for i in 1..n:
    alpha_tilde[i] = alpha[i] + eta * ( e[i] + sum_{k != i} I[i,k] *
    alpha[k] )
Normalize: alpha = alpha_tilde / ||alpha_tilde||_2
Collapse if max_i |alpha[i]|^2 > tau; else continue
```

Estimating I_{ij}.

Initialize from cosine similarity; set κ_{ij} via domain priors, weak supervision (optimize ECS/HQI on solved cases), or expert elicitation.

Complexity.

Per step $O(n^2)$ for interference + $O(n)$ for projection aggregation; with m observations, $O(n^2 m)$.

Appendix C. Computational Sketch and Implementation Notes

Appendix C.1. State, Evidence, Interference

Let $\mathcal{H} = \{H_i\}_{i=1}^n$ and $\mathcal{O} = \{O_j\}_{j=1}^m$. Maintain

$$\left|\Psi^{(t)}\right\rangle = \sum_{i=1}^{n} \alpha_i^{(t)} |H_i\rangle, \qquad \sum_i |\alpha_i^{(t)}|^2 = 1.$$

Embed hypotheses and observations with a sentence encoder (e.g., Sentence-BERT [8]) to obtain vectors $\mathbf{h}_i, \mathbf{o}_j \in \mathbb{R}^d$. Use cosine similarity as a proxy for $\langle H_i \mid O_j \rangle$.

Evidence projection for a new batch is as follows $\mathcal{O}^{(t)}$:

$$e_i^{(t)} = \max_{O_j \in \mathcal{O}^{(t)}} \frac{\mathbf{h}_i \cdot \mathbf{o}_j}{\|\mathbf{h}_i\| \, \|\mathbf{o}_j\|}.$$

Interference matrix $I \in \mathbb{R}^{n \times n}$ models hypothesis interaction. A simple, effective choice is the following:

$$I_{ii} = 1, \qquad I_{ij} = \text{cos_sim}(\mathbf{h}_i, \mathbf{h}_j) \cdot \kappa_{ij}, \;\; i \neq j,$$

with $\kappa_{ij} \in [-1, 1]$ permitting constructive (>0) or destructive (<0) coupling. In data-poor settings set $\kappa_{ij} = 1$ (pure similarity). When expert priors are available, hand-tune signs to encode known exclusivities or complementarities.

Appendix C.2. Amplitude Update and Collapse

With learning rate $\eta \in (0, 1)$

$$\tilde{\alpha}_i^{(t+1)} = \alpha_i^{(t)} + \eta \left(e_i^{(t)} + \sum_{k \neq i} I_{ik} \, \alpha_k^{(t)} \right), \qquad \alpha_i^{(t+1)} = \frac{\tilde{\alpha}_i^{(t+1)}}{\sqrt{\sum_\ell |\tilde{\alpha}_\ell^{(t+1)}|^2}}.$$

Define a coherence functional that grows when one (or a consistent subset) dominates:

$$\mathcal{C}^{(t)} = \max_i |\alpha_i^{(t)}|^2 \quad \text{or} \quad \mathcal{C}^{(t)} = \sum_i |\alpha_i^{(t)}|^4.$$

Collapse when $\mathcal{C}^{(t)} \geq \tau$ or when a decision deadline arrives. If several $|\alpha_i|$ remain high and $I_{ij} > 0$, return a *hybrid synthesis* $H_{i \wedge j}$.

Appendix C.3. Synthesis operator (hybrids)

Given two high-mass hypotheses H_i, H_j, form a content-level blend

$$\mathrm{mix}(\mathbf{h}_i, \mathbf{h}_j) = \lambda\, \mathbf{h}_i + (1-\lambda)\, \mathbf{h}_j, \quad \lambda = \frac{|\alpha_i|^2}{|\alpha_i|^2 + |\alpha_j|^2},$$

and instantiate a natural-language synthesis using a controlled LLM prompt that lists (i) active hypotheses and amplitudes, (ii) salient observations and projections, (iii) entries of I_{ij} that justify merger or exclusion. This yields faithful, auditable text (see [11] for a clinical instance and [2] for a strategic instance).

Appendix C.4. Implementation Note

A lightweight pipeline suffices the following:

1. Sentence-BERT embeddings for $\mathbf{h}_i, \mathbf{o}_j$ (384–768d).
2. Evidence scores e_i via cosine similarity.
3. I_{ij} from similarity; optionally sign κ_{ij} with expert priors (exclusivity vs. complementarity).
4. Iterate the update; monitor $\mathcal{C}^{(t)}$; stop at τ or deadline.
5. If hybrid, call the synthesis operator with (α, I, e) as structured context.

References

1. Busemeyer, J.R.; Bruza, P.D. *Quantum Models of Cognition and Decision*; Cambridge University Press: Cambridge, UK, 2012.
2. Ghisellini, R.; Pareschi, R.; Pedroni, M.; Raggi, G.B. From Extraction to Synthesis: Entangled Heuristics for Agent-Augmented Strategic Reasoning. *arXiv* **2025**, arXiv:2507.13768. [CrossRef]
3. Peirce, C.S. Deduction, Induction, and Hypothesis. *Pop. Sci. Mon.* **1878**, *13*, 470–482.
4. Magnani, L. Cognitive Systems Monographs. In *Abductive Cognition—The Epistemological and Eco-Cognitive Dimensions of Hypothetical Reasoning*; Springer: Berlin/Heidelberg, Germany, 2009; Volume 3. [CrossRef]
5. Reid, D.A.; Knipping, C. Types of Reasoning. In *Proof in Mathematics Education: Research, Learning and Teaching*; Brill Sense: Leiden, The Netherlands, 2019; pp. 63–81.
6. Busemeyer, J.R.; Pothos, E.M.; Franco, R.; Trueblood, J.S. A quantum theoretical explanation for probability judgment errors. *Psychol. Rev.* **2011**, *118*, 193–218. [CrossRef] [PubMed]
7. Ghisellini, R.; Pareschi, R.; Pedroni, M.; Raggi, G.B. Recommending Actionable Strategies: A Semantic Approach to Integrating Analytical Frameworks with Decision Heuristics. *Information* **2025**, *16*, 192. [CrossRef]
8. Reimers, N.; Gurevych, I. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, 3–7 November 2019; Inui, K., Jiang, J., Ng, V., Wan, X., Eds.; Association for Computational Linguistics: Stroudsburg, PA, USA, 2019; pp. 3980–3990. [CrossRef]
9. Borghoff, U.M.; Bottoni, P.; Pareschi, R. Human-artificial interaction in the age of agentic AI: A system-theoretical approach. *Front. Hum. Dyn.* **2025**, *7*, 1579166. [CrossRef]
10. Pareschi, R. Beyond Human and Machine: An Architecture and Methodology Guideline for Centaurian Design. *Sci* **2024**, *6*, 71. [CrossRef]
11. Pareschi, R. Abductive reasoning with the GPT-4 language model: Case studies from criminal investigation, medical practice, scientific research. *Sist. Intell.* **2023**, *35*, 435–444. [CrossRef]
12. Sanders, R.H. Modified Newtonian Dynamics as an Alternative to Dark Matter. *Annu. Rev. Astron. Astrophys.* **2002**, *40*, 263–317. [CrossRef]
13. Milgrom, M. MOND vs. Dark Matter in Light of Historical Parallels. *arXiv* **2019**, arXiv:1910.04368. [CrossRef]
14. Korenaga, J. Initiation and Evolution of Plate Tectonics on Earth: Theories and Observations. *Annu. Rev. Earth Planet. Sci.* **2013**, *41*, 117–151. [CrossRef]
15. Mayer, M.C.; Pirri, F. First order abduction via tableau and sequent calculi. *Log. J. IGPL* **1993**, *1*, 99–117. [CrossRef]
16. Piazza, M.; Pulcini, G.; Sabatini, A. Abduction as Deductive Saturation: A Proof-Theoretic Inquiry. *J. Philos. Log.* **2023**, *52*, 1575–1602. [CrossRef]
17. Kakas, A.C.; Kowalski, R.A.; Toni, F. Abductive Logic Programming. *J. Log. Comput.* **1992**, *2*, 719–770. [CrossRef]

18. Denecker, M.; Kakas, A.C. Abduction in Logic Programming. In *Computational Logic: Logic Programming and Beyond*; Essays in Honour of Robert A. Kowalski, Part I; Kakas, A.C., Sadri, F., Eds.; Lecture Notes in Computer Science; Springer: Berlin/Heidelberg, Germany, 2002; Volume 2407, pp. 402–436. [CrossRef]
19. Ishihata, M.; Sato, T. Bayesian Inference for Statistical Abduction Using Markov Chain Monte Carlo. In Proceedings of the Asian Conference on Machine Learning, Taoyuan, Taiwain, 14–15 November 2011; pp. 81–96.
20. Raghavan, S.V. Bayesian Abductive Logic Programs: A Probabilistic Logic for Abductive Reasoning. In Proceedings of the IJCAI Workshop on Plan, Activity, and Intent Recognition (PAIR-09), Pasadena, CA, USA, 11–13 July 2011.
21. Reggia, J.A.; Nau, D.S.; Wang, P.Y. Diagnostic expert systems based on a set covering model. *Int. J. Man-Mach. Stud.* **1983**, *19*, 437–460. [CrossRef]
22. Reggia, J.A.; Nau, D.S.; Wang, P.Y. A formal model of diagnostic inference. I. Problem formulation and decomposition. *Inf. Sci.* **1985**, *37*, 227–256. [CrossRef]
23. Reggia, J.A.; Nau, D.S.; Wang, P.Y.; Peng, Y. A formal model of diagnostic inference, II. Algorithmic solution and application. *Inf. Sci.* **1985**, *37*, 257–285. [CrossRef]
24. Zadeh, L.A. Fuzzy logic and approximate reasoning. *Synthese* **1975**, *30*, 407–428. [CrossRef]
25. Klir, G.J.; Yuan, B. *Fuzzy Sets and Fuzzy Logic—Theory and Applications*; Prentice Hall: Wilmington, DE, USA, 1995.
26. Dung, P.M. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming, and n-person games. *Artif. Intell.* **1995**, *77*, 321–357. [CrossRef]
27. Garcez, A.D.; Lamb, L.C.; Gabbay, D.M. *Neural-Symbolic Cognitive Reasoning*; Springer: Berlin/Heidelberg, Germany, 2009.
28. Lamb, L.C.; d'Avila Garcez, A.S.; Gori, M.; Prates, M.O.R.; Avelar, P.H.C.; Vardi, M.Y. Graph Neural Networks Meet Neural-Symbolic Computing: A Survey and Perspective. In Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020, Online, 7–15 January 2021; Bessiere, C., Ed.; ijcai.org: San Francisco, CA, USA, 2020; pp. 4877–4884. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.