

A Trust-Centric Neuro-Symbolic Architecture for Verifiable Large Language Model Claims

Neurosymbolic Artificial Intelligence
XX(X):3–46
©The Author(s) 2026
Reprints and permission:
sagepub.co.uk/journalsPermissions.nav
DOI: 10.1177/ToBeAssigned
www.sagepub.com/

SAGE

David Farrugia¹ and Alexiei Dingli¹

Abstract

Large Language Models (LLMs) exhibit impressive fluency yet remain prone to hallucinations, opaque reasoning, and inconsistent use of source quality, limitations that inhibit adoption in high-stakes domains. This paper formalises a trust-centric neuro-symbolic pipeline that unifies three previously separate capabilities: the translation of unconstrained LLM outputs into First-Order Logic (FOL); the verification of those formulae against provenance-rich, digitally signed Knowledge Graphs (KGs); and the generation of user-oriented explanations designed to calibrate human trust. The key architectural contribution is the integration of digital signatures and reputation scores as first-class weights in the verification decision, so that evidence is weighted by its provenance rather than treated uniformly; we also specify how these weights would enter a differentiable satisfiability objective, which the evaluated verifier does not yet implement. The architecture has been implemented as a runnable proof-of-concept covering all three stages, together with an explicit mapping from FOL formulae to knowledge-graph evidence, and evaluated on a curated pharmacological knowledge graph of 98 triples drawn from nine sources, 74 of them in signed graphs. Against KG-relative reference verdicts for 50 author-curated medical claims, the verifier reaches an accuracy of 0.72 on gold-standard FOL, compared with 0.66 for the same verifier with uniform evidence weights (a difference that is not statistically significant at this scale). All of its errors are abstentions rather than wrong verdicts. Because 49 of the 50 reference verdicts are Verified, a constant Verified answer would score 0.98, so these accuracies measure how much of the support contained in the KG the verifier recovers; assessing the rejection of false claims requires a benchmark that contains them. When claims are grounded by the LLM translator rather than supplied as gold-standard FOL, accuracy is 0.16. In both settings the verifier asserts no incorrect verdict on these claims: every error is an abstention, which identifies grounding as the main target for improvement. Verdicts are stable under moderate threshold and reputation perturbations. Evaluation of the grounding stage on the FOLIO benchmark and the medical claims shows strong Logical Equivalence ($LE = 0.853$) on FOLIO but degradation on complex relational claims (drug interactions: $LE = 0.479$). This work provides a provenance-aware, explainable verification layer and an evaluation protocol for trustworthy language technologies in safety-critical settings.

Keywords

Large language models, Neuro-symbolic AI, Knowledge graphs, Provenance, Trust calibration, First-order logic, Semantic Web

1 Introduction

Large Language Models (LLMs) have achieved strong results across natural language understanding and generation tasks (Brown et al. 2020; Vaswani et al. 2017). However, their tendency to “hallucinate” (generating factually incorrect but plausible-sounding statements) remains a major barrier to their adoption in high-stakes environments such as healthcare, finance, and legal advisory (Ji et al. 2023; Augenstein et al. 2024; Choudhury and Chaudhry 2024). In these domains, erroneous outputs can lead to real-world harm, ranging from misdiagnosis to significant financial loss (Wu et al. 2025; Moglia et al. 2024).

Current efforts to ensure LLM reliability often rely on purely statistical methods, such as Retrieval-Augmented Generation (RAG) or human-in-the-loop reviews (Patil and Gudivada 2024; Ouyang et al. 2022; Demartini et al. 2020). While these provide incremental improvements, they lack the formal rigour required for absolute verification and often treat all retrieved evidence as uniformly reliable, ignoring the provenance and cryptographic integrity of source data (Hartig 2009; Rahimzadeh Holagh and Mohebbi 2019).

1.1 Closing the Triangle: Fluency, Rigour, and Trust

We argue that the path toward trustworthy language technologies requires bridging a three-fold gap between largely independent capabilities (Marcus and Davis 2019; Besold et al. 2021). We conceptualise this as a triangle (Figure 1) where each vertex represents an essential requirement:

- **LLM Fluency:** Today’s models excel at context-rich expressiveness but lack formal guarantees of correctness or grounding (Brown et al. 2020; Ji et al. 2023).
- **Symbolic Rigour:** Neuro-symbolic and symbolic systems provide formal logic and correctness guarantees but often assume clean, curated data and struggle to scale with the noisy information of the open web (Garcez et al. 2009; Besold et al. 2021).
- **Trust and Provenance Mechanisms:** Standards in the Semantic Web and information security provide mature methods for tracking source quality and data integrity, yet these remain largely separate from AI generation loops (Artz and Gil 2007; Golbeck 2006).

¹Department of Artificial Intelligence, Faculty of ICT, University of Malta, Msida, Malta

Corresponding author:

David Farrugia, Department of Artificial Intelligence, Faculty of ICT, University of Malta, Msida MSD 2080, Malta.

Email: david.farrugia@um.edu.mt




Figure 1. The “Closing the Triangle” framework. Most existing approaches address at most two of the three vertices. Our architecture integrates all three by embedding LLM outputs within a logically grounded, provenance-aware verification loop.

Our architecture *closes this triangle* by embedding LLMs within a logically grounded, provenance-aware framework. By integrating digital signatures and reputation scores directly into the verification decision, we enable trust-weighted verification where provenance metadata determines the verdict (Artz and Gil 2007); we also specify how the same weights would enter a differentiable, LTN-style satisfiability objective (Badreddine et al. 2022), which we leave to future work.

1.2 Research Objectives

To formalise this vision, this paper addresses the following research questions through the proposed architectural blueprint:

- RQ1. **Reliable Semantic Grounding:** How can unconstrained LLM claims be reliably translated into First-Order Logic (FOL) to enable syntactically well-formed and globally consistent verification (Yang et al. 2024; Brunello et al. 2025)?
- RQ2. **Trust-Aware Verification:** To what extent can a provenance-aware verification layer over a signed knowledge graph verify such formulae correctly, and what does trust weighting add over treating all evidence uniformly (Badreddine et al. 2022; Artz and Gil 2007)?
- RQ3. **Explanations for Calibrated Trust:** Which forms of explanatory feedback most effectively foster appropriate reliance, ensuring users neither over-trust nor under-trust the system (Kim et al. 2024a; Romeo and Conti 2025)?

Together, these questions outline a roadmap for deploying language models that are fluent, validated, and trustworthy by design (Floridi and Cowsli 2019; Bellogín et al. 2024). This paper contributes a formally specified architectural design *accompanied by a runnable proof-of-concept implementation* that covers all three pipeline stages and will

be released for independent replication. It does not answer these questions in full: RQ3 in particular would require a user study, which we do not report. We provide preliminary empirical evidence from two complementary experiments (Section 5): an evaluation of the grounding stage on FOLIO and a domain-specific medical-claims dataset, and an evaluation of the verifier against KG-relative reference verdicts on a curated, digitally signed pharmacological knowledge graph, with an unweighted ablation, a cascaded run on predicted FOL and a sensitivity analysis.

1.3 Main Contributions

The main contributions of this work are:

1. A five-signal provenance-weighted trust score that integrates digital signatures, reputation, freshness, completeness, and corroboration as first-class weights in the verification decision (Definition 1), with a specification of how they would enter a differentiable, LTN-style satisfiability objective, and a signature-gated variant for regulated environments (Equation 4).
2. A conflict-aware verdict function with configurable trust thresholds that returns an explicit *Indeterminate* verdict when evidence is insufficient or contradictory, rather than forcing a binary decision (Equation 7).
3. A reusable, staged evaluation protocol with component-level and system-level checkpoints and ablation baselines, with automated scoring against curated reference data (Section 4).
4. A reference implementation of the full three-stage pipeline, to be released publicly on acceptance, including the trust weight computation, conflict-aware verdict function, and explanation trace generator, together with a curated pharmacological knowledge graph carrying digital signatures and provenance metadata for nine sources.
5. Preliminary empirical evidence from stage-level, gold-FOL and cascaded experiments: the grounding stage achieves $LE = 0.853$ on FOLIO and $LE = 0.801$ on domain-specific medical claims (with sharp degradation on drug interactions, $LE = 0.479$), and the verifier reaches accuracy 0.72 against KG-relative reference verdicts on gold FOL (0.66 with uniform weights; not significant at $n = 50$) and 0.16 on predicted FOL, with every error an abstention rather than an incorrect verdict.

Several components are adopted rather than original. The NL→FOL translator is LogicLLaMA (Yang et al. 2024); the grammar-masked decoding of Equation 2 follows Li et al. (2024) and is specified, not implemented; the satisfiability objective of Equation 5 follows Logic Tensor Networks (Badreddine et al. 2022); provenance is represented with PROV-O (Moreau et al. 2011) and signed named graphs (Carroll et al. 2005); and the three-way label set follows FEVER (Thorne et al. 2018). The original elements are the five-signal trust score and its gated variant, the conflict-aware verdict and aggregation rules, the FOL→KG mapping (Section 3.2), the evaluation protocol, and their integration into one pipeline.

Table 1 summarises the current maturity of each pipeline component to set clear expectations about what this paper delivers.

Table 1. Component maturity. Each pipeline stage is classified by its current status and the evidence provided in this paper.

Component	Status	Evidence
Stage A: $NL \rightarrow FOL$ grounding	Implemented & evaluated	LogicLLaMA (adopted) tested on FOLIO ($n=254$) and medical claims ($n=50$); LE, BLEU, parse rate, tree-edit distance; grammar mask (Eq. 2) specified, not implemented
Stage B: Trust-aware verification	Implemented & evaluated	Rule-based: five-signal trust weight, $FOL \rightarrow KG$ mapping, conflict-aware verdict rule; no learned component; KG-relative accuracy, Unweighted ablation, cascaded run, sensitivity
LTN objective (Eq. 5)	Specified only	LTN-style trust-weighted satisfiability loss; no f_θ is learned in the evaluated system
Stage C: Explanatory feedback	Implemented	Explanation trace generator emits verdict, ranked evidence, trust signals, and recommendations
Trust weight function	Implemented	Five-signal model with signature-gated variant; design properties stated; weights computed end-to-end
Evaluation protocol	Partially executed	E_1, E_2, E_3 (KG-relative), Unweighted ablation, cascaded run and sensitivity computed; E_4 , latency, conflict-robustness rate and the remaining baselines not computed

1.4 Paper Organisation

The remainder of this paper is structured as follows. Section 2 surveys the relevant literature across the Semantic Web trust stack, LLM hallucination, neuro-symbolic

AI, formal logic translation, fact verification, self-refinement methods, and socio-technical trust. Section 3 presents the three-stage architectural specification with formal definitions, mathematical formulations, and an illustrative walkthrough grounded in real outputs from the implementation. Section 4 specifies a reusable evaluation methodology with component-level and system-level checkpoints. Section 5 reports preliminary experimental results for the grounding stage (FOLIO and medical claims) and for the verifier on 50 medical claims against the signed pharmacological knowledge graph. Section 6 discusses assumptions, limitations, and adversarial considerations. Section 7 concludes the paper and outlines future research directions.

2 Background and Related Work

The development of a trust-centric neuro-symbolic architecture requires the synthesis of seven primary research strands: Semantic Web trust layers, LLM capabilities and hallucination, natural language-to-logic translation, neuro-symbolic AI frameworks, automated fact verification, self-refinement and chain-of-verification methods, and socio-technical trust calibration. This section critically reviews each strand and concludes with a gap analysis (Table 3) that motivates our contribution.

2.1 Semantic Web Stack and the Trust Layer

The Semantic Web vision aims to transform the web into a machine-processable ecosystem by annotating data with explicit semantics using standards such as the Resource Description Framework (RDF) and the Web Ontology Language (OWL) (Shadbolt et al. 2006; Koivunen and Miller 2001; Horrocks 2005). At the apex of the Semantic Web “layer cake” sits the Trust layer (Artz and Gil 2007). While the lower layers facilitate data interchange and reasoning, the trust layer is designed to verify the authenticity and reliability of both statements and sources (Artz and Gil 2007; Golbeck 2006).

Mature mechanisms exist for provenance tracking (e.g., the W3C PROV Ontology, PROV-O), cryptographic signatures, and reputation scoring (Moreau et al. 2011; Carroll et al. 2005; Golbeck 2006). Named Graphs (Carroll et al. 2005) allow individual RDF statements to carry metadata such as authorship, timestamp, and signature status, providing the foundation for per-triple trust annotation. Reputation models, such as the probabilistic trust propagation framework of Golbeck (2006), compute trust transitively across social networks, while ontology-based approaches encode trust requirements directly into OWL axioms (Huang and Fox 2006).

Despite these advances, wide-scale adoption of Semantic Web trust mechanisms remains limited due to challenges in standardisation, computational cost, and integration with real-time unstructured data streams (Harth et al. 2011; Hitzler 2021). Recent surveys on the convergence of knowledge graphs and AI note that trust and provenance remain “the missing piece” in most knowledge-augmented systems (Scherp et al. 2024; Pan et al. 2024). Our architecture addresses this gap by making these Semantic Web trust signals first-class parameters of the verification decision.

2.2 LLM Capabilities and the Challenge of Hallucination

Transformer-based Large Language Models, from BERT (Devlin et al. 2019) and GPT (Brown et al. 2020) to their successors (OpenAI et al. 2024), have achieved state-of-the-art results across Natural Language Processing (NLP) tasks (Vaswani et al. 2017; Qiu et al. 2020). However, their “black-box” nature and susceptibility to hallucinations present significant risks in high-stakes domains (Ji et al. 2023; Bender et al. 2021).

Hallucinations, where a model generates contextually plausible but factually incorrect statements, are broadly categorised into *intrinsic* hallucinations (contradicting the source material) and *extrinsic* hallucinations (fabricating unsupported claims) (Ji et al. 2023). Both types arise because LLMs generate text based on learned token-level probabilities rather than explicit knowledge bases (Augenstein et al. 2024). The problem is amplified in safety-critical domains: clinical studies report that LLM-generated medical summaries contain clinically significant hallucinations at non-trivial rates (Asgari et al. 2025), while legal applications exhibit similar vulnerabilities (Choudhury and Choudhry 2024).

Existing mitigation strategies include Reinforcement Learning from Human Feedback (RLHF) (Ouyang et al. 2022), Retrieval-Augmented Generation (RAG) (Patil and Gudivada 2024), and post-hoc fact-checking pipelines (Thorne et al. 2018). While RLHF reduces overtly harmful outputs, it does not provide formal guarantees of factual correctness. RAG improves grounding by retrieving relevant documents but treats all retrieved evidence as uniformly reliable, providing no mechanism for assessing source provenance (Lin et al. 2024). Our work goes beyond these approaches by coupling retrieval with formal logic verification and provenance-weighted trust assessment.

2.3 Natural Language to Formal Logic Translation

The reliable translation of natural language into formal logic ($NL \rightarrow FOL$) is a long-standing challenge at the intersection of computational linguistics and knowledge representation (Bos 2008). Early approaches employed Combinatory Categorical Grammars (CCG) (Stanojević and Steedman 2019) and Abstract Meaning Representations (AMR) (Banarescu et al. 2013) to produce intermediate semantic structures that could be mapped to first-order formulae. While effective for well-formed sentences, these systems struggle with the informal, noisy outputs typical of LLMs.

Recent work has explored using LLMs themselves as translators. Olausson et al. (2023) demonstrate that chain-of-thought prompting can guide LLMs to produce first-order logic statements with reasonable accuracy on standard benchmarks. Pan et al. (2023) propose Logic-LM, which combines LLM-generated logical forms with a symbolic solver to improve reasoning accuracy. Li et al. (2024) introduce Formal-LLM, which constrains LLM decoding through a grammar mask to ensure syntactic well-formedness of the logical output. Taking a different approach, Yang et al. (2024) present LogicLLaMA, a LoRA-finetuned LLaMA-7B model trained specifically for direct $NL \rightarrow FOL$ translation, achieving strong logical equivalence scores on the FOLIO benchmark. We adopt LogicLLaMA as the grounding model in our feasibility study (Section 5).

More recently, [Putra et al. \(2026\)](#) propose NL2Logic, which introduces an Abstract Syntax Tree (AST) as an intermediate representation between natural language and solver-ready code. Their recursive semantic parser decomposes each sentence clause-by-clause, and a deterministic AST-guided generator produces syntactically correct output. On the FOLIO, LogicNLI, and ProofWriter benchmarks, NL2Logic achieves 99% syntactic accuracy and improves semantic correctness by 30% over prior baselines, demonstrating that decoupling logical parsing from code generation substantially reduces hallucination in the translation process.

Evaluation of $NL \rightarrow FOL$ systems has been advanced by benchmarks such as FOLIO ([Han et al. 2024](#)), which provides expert-annotated natural language–FOL pairs, and by systematic evaluations of LLM-based translators across multiple logic formalisms ([Brunello et al. 2025](#); [Yang et al. 2024](#)). Despite progress, exact-match accuracy for complex multi-predicate sentences remains limited, with preliminary evaluations suggesting that current LLM-based translators struggle to achieve reliable performance on such inputs ([Brunello et al. 2025](#); [Yang et al. 2024](#)), highlighting the need for constrained decoding or structured decomposition strategies ([Putra et al. 2026](#)). Our Stage A specifies grammar-constrained decoding for this purpose; the evaluated implementation uses LogicLLaMA without it.

2.4 Neuro-Symbolic AI and Differentiable Logic

Neuro-symbolic AI aims to unite the pattern-learning abilities of neural networks with the formal expressiveness of symbolic logic ([Garcez et al. 2009](#); [Garcez and Lamb 2023](#)). A recent survey by [Cheng et al. \(2024\)](#) provides a comprehensive taxonomy of neural-symbolic methods for knowledge graph reasoning, covering KG completion, complex query answering, and logical rule learning; the survey concludes that symbolic reasoning offers interpretability and generalisability through logical rules but struggles with scalability, while neural methods scale well but fall short in modelling higher-order dependencies, motivating hybrid integration. This hybrid paradigm is increasingly positioned as a linchpin for validating LLMs, with a cross-study survey indicating that coupling LLMs with explicit symbolic or knowledge-graph verifiers typically reduces hallucination rates by approximately 45–75%, though these figures are not strictly comparable across datasets ([Pan et al. 2023](#); [Olausson et al. 2023](#); [Cao 2024](#); [Tang et al. 2024](#); [Quan et al. 2024](#); [Wadden et al. 2020](#); [Ciampaglia et al. 2015](#); [Hou et al. 2024](#); [Hanselowski et al. 2019](#); [Atanasova et al. 2020](#); [Shi and Wening 2016](#)).

Table 2 summarises representative results across these studies. Purely neural baselines yield only marginal gains, whereas knowledge-graph and hybrid neuro-symbolic approaches achieve substantially larger error reductions (median $\approx 54\%$).

Several frameworks have emerged to make logic differentiable, allowing symbolic constraints to guide neural learning via gradient-based optimisation. Logic Tensor Networks (LTNs) ([Badreddine et al. 2022](#)) compute a real-valued “satisfaction degree” for each grounded formula. DeepProbLog ([Manhaeve et al. 2021a,b](#)) embeds neural predicates within probabilistic logic programs; NeurASP ([Yang et al. 2021](#)) uses answer set programs; DeepStochLog ([Winters et al. 2022](#)) employs stochastic definite clause

Table 2. Error reduction from verification modules across representative studies. Neural-only baselines yield modest gains, whereas KG and hybrid neuro-symbolic approaches reduce error rates by approximately 45–75%. Figures are not strictly comparable across datasets and evaluation protocols.

Study	Dataset	Base LLM	Type	Base Err.%	+ Module	$\Delta\%$
Hanselowski et al. (2019)	FEVER	BiLSTM	Neural	49.1	35.8	−27.1
Atanasova et al. (2020)	FEVER	RoBERTa	Neural	31.4	30.0	−4.5
Ciampaglia et al. (2015)	Wiki Claims	—	KG	39.0	9.0	−76.9
Shi and Weninger (2016)	PolitiFact	—	KG	35.0	8.0	−77.1
Wadden et al. (2020)	SciFact	BERT	Hybrid	23.1	12.7	−45.0
Olausson et al. (2023)	ProofWriter	GPT-J	Hybrid	24.0	11.0	−54.2
Pan et al. (2023)	LogicQA	LLaMA-7B	Hybrid	19.4	9.6	−50.5
Tang et al. (2024)	FEVER	GPT-3.5	Hybrid	29.2	15.1	−48.3
Cao (2024)	FEVER	GPT-2 XL	Hybrid	26.4	11.4	−56.8
Hou et al. (2024)	WebQSP	ChatGPT	Hybrid	34.7	17.3	−50.1
Quan et al. (2024)	e-SNLI	GPT-4	Hybrid	64.0	16.0	−75.0

grammars; and Logic Neural Networks (LNNs) (Riegel et al. 2020) assign each neuron a logical connective interpretation.

The key limitation shared by all these frameworks is their treatment of input facts: every triple or ground atom is assumed to be equally reliable (Badreddine et al. 2022). In real-world knowledge graphs, however, facts vary considerably in their provenance, recency, and trustworthiness. Source credibility has been combined with soft logical reasoning before: ClaimEval (Samadi et al. 2016) jointly estimates source credibility and claim correctness with Probabilistic Soft Logic (Bach et al. 2017), expressing credibility relations as first-order rules with continuous, Łukasiewicz-style truth values. Our work differs in the provenance signals it uses and where they enter: credibility is not learned from claim agreement but derived from per-triple provenance metadata, namely digital-signature validity, source reputation, freshness, PROV-O completeness, and corroboration, and it weights the KG evidence consulted when verifying formulae translated from LLM output. We specify how such a weight would enter an LTN-style satisfiability objective (Section 3); the verifier evaluated in this paper applies it in a rule-based decision function instead.

2.5 Fact Verification and Claim Checking

Automated fact verification has received substantial attention following the release of large-scale benchmarks. The Fact Extraction and VERification (FEVER) dataset (Thorne et al. 2018) introduced a three-class task (Supported, Refuted, Not Enough Info) that closely mirrors our {Verified, Falsified, Indeterminate} verdict set. LIAR (Wang 2017) provides real-world political claims with fine-grained truthfulness labels, while SciFact (Wadden et al. 2020) targets scientific claim verification against research abstracts. More recently, MiniCheck (Tang et al. 2024) has explored lightweight LLM-based fact-checking with promising efficiency–accuracy trade-offs.

Knowledge-graph-based verification takes a complementary approach by resolving claims against structured knowledge rather than unstructured text (Ciampaglia et al.

2015; Shi and Weninger 2016). Konstantinovskiy et al. (2021) survey the pipeline components required for automated claim detection and verification, while recent work has explored combining KG-based evidence retrieval with neural reasoning (Hou et al. 2024; Cao 2024). Most recently, KG-CRAFT (Lourenço et al. 2026) constructs a knowledge graph from claims and associated reports, then formulates contrastive questions grounded in the graph structure to guide LLM-based verification. KG-CRAFT achieves state-of-the-art results on two real-world fact-checking benchmarks (LIAR-RAW and RAWFC), with F1 improvements of up to 44 percentage points over prior methods, providing strong evidence that structured knowledge substantially enhances LLM verification capabilities.

A complementary line of recent work evaluates LLM factuality at a finer granularity than binary claim labels. FActScore (Min et al. 2023) decomposes long-form generations into atomic facts and computes the percentage supported by a reliable knowledge source, revealing that even ChatGPT achieves only 58% atomic factual precision on biography generation. HaluEval (Li et al. 2023) provides a large-scale benchmark of 35,000 hallucinated samples across question answering, dialogue, and summarisation, finding that ChatGPT fabricates unverifiable information in approximately 19.5% of user queries. More recently, SAFE (Wei et al. 2024) employs LLM agents augmented with web search to evaluate individual facts in long-form responses, demonstrating that automated evaluators can match or exceed crowdsourced human annotators at a fraction of the cost. These fine-grained methods represent important advances in measuring LLM factuality; however, they share a common limitation with the claim-level systems described above: all retrieved evidence is treated as uniformly reliable, irrespective of its provenance or source credibility.

A key gap in existing fact-verification systems is their treatment of evidence quality. Most pipelines assume that if a claim is supported by retrieved evidence, it is true, regardless of whether that evidence comes from a peer-reviewed source or an anonymous web page. Our architecture fills this gap by assigning provenance-weighted trust scores to every piece of evidence before it enters the reasoning loop.

2.6 Self-Refinement and Chain-of-Verification

A recent family of methods improves LLM output quality through iterative self-correction without external supervision. Self-Consistency (Wang et al. 2023) samples multiple reasoning paths via chain-of-thought prompting and selects the most frequent answer by majority vote, exploiting the intuition that correct solutions are reachable through diverse reasoning trajectories. Self-Refine (Madaan et al. 2023) extends this idea to a generate–feedback–refine loop in which the same LLM critiques and iteratively improves its own output, achieving substantial gains across dialogue, code, and mathematical reasoning tasks. Reflexion (Shinn et al. 2023) goes further by equipping language agents with an episodic memory of verbal self-reflections, enabling learning from past failures without weight updates.

Most relevant to our evaluation philosophy, Chain-of-Verification (CoVe) (Dhuliawala et al. 2024) decomposes hallucination reduction into four stages: draft an initial response,

generate verification questions, answer them independently (preventing bias from the original output), and revise the response. CoVe demonstrates that verification questions are answered more accurately than the claims in the original generation, yielding significant hallucination reduction across list-based, span-extraction, and long-form tasks.

These methods share a common limitation: all verification is performed *within the neural model itself*. Self-Consistency relies on statistical agreement among sampled paths rather than logical proof; Self-Refine and Reflexion depend on the model’s own ability to detect and correct errors, which is bounded by the same training distribution that caused the error. CoVe’s independent answering of verification questions mitigates but does not eliminate this circularity, since the verifier and the generator share the same parametric knowledge. None of these approaches consult external structured knowledge, require formal logical consistency, or account for the provenance or reliability of evidence.

Our architecture addresses this gap by grounding verification in a symbolic knowledge base with explicit trust metadata. Rather than asking the model to check itself, we delegate verification to a formal reasoning engine operating over provenance-weighted facts. The CoVe principle does, however, inspire our evaluation protocol (Section 4): just as CoVe verifies each claim through independent sub-checks, our staged evaluation validates each pipeline component against its own gold standard before assessing system-level behaviour.

2.7 Socio-Technical Trust and Explainability

A technically correct AI solution can still falter if users do not appropriately rely on its outputs (Romeo and Conti 2025; Carnat 2024). Research on human–AI interaction reveals two complementary failure modes. “Automation bias,” the tendency to over-rely on AI decisions, can lead to uncritical acceptance of incorrect outputs (Carnat 2024; Bo et al. 2024). Conversely, “algorithm aversion,” the refusal to trust accurate AI recommendations, reduces uptake in high-value applications (Dietvorst et al. 2015). Both failure modes are exacerbated by opaque reasoning: when users cannot inspect the basis for a verdict, they lack the information needed to calibrate their trust appropriately (Lipton 2018).

Explainability research distinguishes between *global* explanations (characterising overall model behaviour) and *local* explanations (justifying individual predictions) (Doshi-Velez and Kim 2017). For fact-verification systems, local explanations are particularly relevant: users need to understand *which* evidence supported or contradicted a specific claim, and *how reliable* that evidence is. Recent surveys on explainability in LLM-based systems (Zhao et al. 2024; Lu et al. 2025) note that most explanation methods focus on attention-weight visualisation or natural language rationales, neither of which provides formal guarantees about the reasoning process.

Studies in clinical and legal domains further highlight the importance of domain-appropriate explanation design. Mirzaei et al. (2024) find that clinicians prefer structured, evidence-linked explanations over free-text summaries, while Kim et al. (2024b) report that confidence indicators significantly affect user reliance decisions. Our architecture

addresses this by producing verdict-specific explanation traces that explicitly reference the trust scores and provenance of the supporting evidence.

2.8 Summary and Gap Analysis

Table 3 synthesises the coverage of existing approaches across the capabilities required for trustworthy LLM verification. On the dimensions we selected, none of the approaches reviewed addresses all six requirements; the contribution of this work lies in integrating mechanisms that exist separately, not in a new individual mechanism. The “This work” row describes the architecture as specified; in the evaluated implementation, Stage A is not constrained to the KG vocabulary and the verifier is rule-based (Table 1). RAG-based systems provide retrieval but lack formal verification and trust weighting. Logic-LM and similar neuro-symbolic pipelines introduce formal reasoning but treat all evidence uniformly. FEVER-style fact-checkers support a three-class verdict set but operate over unstructured text without provenance metadata. Fine-grained factuality methods such as FActScore (Min et al. 2023) and SAFE (Wei et al. 2024) advance evaluation granularity to the atomic-fact level but treat all evidence sources as uniformly reliable. Self-refinement methods such as CoVe (Dhuliawala et al. 2024) improve output quality through iterative self-correction but rely entirely on the model’s own parametric knowledge, lacking formal verification, trust weighting, or provenance tracking. Our architecture brings all six capabilities together in a single, formally specified pipeline.

Table 3. Gap analysis of existing approaches. The six capability dimensions were chosen by the authors to reflect the requirements of this work, and the ratings are our reading of each approach’s published description; the table therefore frames the requirements rather than ranking systems. ✓ = fully addressed; ~ = partially addressed; – = not addressed.

Approach	NL → FOL	Formal Verif.	Trust Weights	Provenance	Explainability	Indeterminate
RAG + Citations	–	–	–	~	~	–
Logic-LM	✓	✓	–	–	~	–
LTN / DeepProbLog	~	✓	–	–	–	–
FEVER-style	–	–	–	–	~	✓
Fine-grained Factuality	–	–	–	–	~	–
KG-based Verification	~	~	–	~	–	–
Self-Refinement (CoVe)	–	–	–	–	~	–
This work	✓	✓	✓	✓	✓	✓

3 Architectural Specification

The proposed architecture is a modular framework for transforming unconstrained neural text generation into verified, provenance-weighted knowledge. It is defined by a three-stage pipeline (Figure 2) that transitions from natural language to formal logic, followed by trust-aware verification and human-centred feedback (Garcez et al. 2009; Yang et al. 2024). Table 4 summarises the notation used throughout this section.

Table 4. Summary of notation.

Symbol	Description
c	Natural language claim generated by the LLM
ψ_θ	Neural encoder (LLM) parameterised by θ
Σ	Target domain schema and ontologies
$\mathcal{C}$	Set of syntax and semantic constraints
T	Translation function $NL \rightarrow FOL$
Φ	Set of first-order logic formulae
G	Knowledge Graph
t	RDF triple $\langle s, p, o \rangle \in G$ (subject, predicate, object)
w_t	Provenance-weighted trust score for triple t
α_i	Weight coefficient for trust signal i
f_θ	LTN neural predicate grounding (specified; not implemented)
$\mathcal{L}_{\text{sat}}$	Trust-weighted satisfiability loss (specified; not implemented)
V	Verdict $\in \{\text{Verified, Falsified, Indeterminate}\}$
τ	Minimum trust threshold for verdict assertion
δ	Minimum trust margin for conflict resolution
W^+, W^-	Aggregate trust for and against a formula
$\mathcal{E}$	Explanation trace
conf	Verdict confidence $W^+ / (W^+ + W^-)$




Figure 2. Overview of the three-stage trust-centric neuro-symbolic pipeline. Stage A translates natural language claims into first-order logic under schema constraints. Stage B verifies the resulting formulae against a provenance-weighted Knowledge Graph. Stage C generates a verdict with an evidence-linked explanation trace.

3.1 Stage A: Semantic Grounding ($NL \rightarrow \Phi$)

The grounding service is the formal interface between Large Language Model (LLM) outputs and the symbolic reasoning engine. The objective is to map a set of declarative claims c , generated by a neural encoder ψ_θ , into a core first-order fragment.

Let Σ represent the target domain schema and ontologies. The translation function T maps the claim to a set of logical formulae Φ :

$$\Phi = T(c, \Sigma, \mathcal{C}) \quad (1)$$

where $\mathcal{C}$ denotes a set of syntax and semantic constraints to be enforced at decoding time (Li et al. 2024). As specified, these constraints would keep the resulting formulae syntactically well-formed and within the Knowledge Base (KB) vocabulary, and ambiguity would be managed by generating confidence-scored alternative parses. The evaluated implementation does neither: it produces a single unconstrained LogicLLaMA translation per claim (Section 5.1).

Constrained Decoding To prevent the generation of ill-formed or out-of-vocabulary predicates, the architecture specifies a grammar-masked decoding strategy adopted from Li et al. (2024). The evaluated implementation does not apply this mask: Stage A uses LogicLLaMA (Yang et al. 2024) with unconstrained generation. At each decoding step k , the token distribution is restricted by a mask M_k derived from the target grammar:

$$P(y_k \mid y_{<k}, c) = \frac{\exp(z_k/\tau_d) \odot M_k}{\sum_j \exp(z_j/\tau_d) \odot M_{k,j}} \quad (2)$$

where z_k is the logit vector, τ_d is the decoding temperature, $\odot$ denotes element-wise multiplication, and $M_k \in \{0, 1\}^{|V|}$ zeroes out tokens that would violate the syntactic grammar of the target logic at position k (Li et al. 2024).

Well-Formedness Check After translation, Stage A checks only that each formula parses under the target grammar. It does *not* test the formulae against the knowledge base: a correctly parsed claim that the KB contradicts (e.g. $P(a)$ when the KB asserts $\neg P(a)$) is not a grounding error but precisely the negative evidence Stage B is designed to weigh, and it may yield a Falsified verdict there. Treating KB contradiction as a Stage A failure would conflate an incorrect parse with a false claim. Formulae that fail to parse are returned as Indeterminate.

3.2 Stage B: Trust-Aware Neuro-Symbolic Verification

The core technical contribution of this architecture is the integration of Semantic Web trust signals as first-class parameters of the verification decision, and, in the specified learned variant, of the reasoning objective (Badreddine et al. 2022; Artz and Gil 2007). Unlike traditional theorem provers that treat all input facts as equally reliable, our verifier weighs evidence based on cryptographic and reputational metadata (Carroll et al. 2005; Golbeck 2006). In the verifier evaluated in Section 5, evidence is located by exact lookup through the declared mapping of Section 3.2, not by similarity retrieval. The

reference implementation also contains a TF-IDF retrieval path, which produced the walkthrough of Section 3.5. Learned KG embeddings (e.g., TransE-style translational embeddings (Bordes et al. 2013)) are a possible basis for softer matching that we do not evaluate.

Definition 1. Provenance-Weighted Trust Score. *For a given triple $t = \langle s, p, o \rangle$ in Knowledge Graph G (following the RDF convention of subject, predicate, and object), we assign a trust weight $w_t \in [0, 1]$. This weight is computed as a convex combination of five provenance signals:*

$$w_t = \sum_{i=1}^5 \alpha_i \cdot s_i(t), \quad \text{subject to} \quad \sum_{i=1}^5 \alpha_i = 1, \quad \alpha_i \geq 0 \quad (3)$$

where the five signals are:

1. $s_1(t) = \text{sig}(t) \in \{0, 1\}$: digital signature validity,
2. $s_2(t) = \text{rep}(t) \in [0, 1]$: source reputation score,
3. $s_3(t) = \text{fresh}(t) \in [0, 1]$: temporal freshness (exponential decay from publication date),
4. $s_4(t) = \text{comp}(t) \in [0, 1]$: provenance completeness (fraction of PROV-O fields populated),
5. $s_5(t) = \text{corr}(t) \in [0, 1]$: corroboration degree (proportion of independent sources confirming the triple). In the specification, two sources are deemed independent when they do not share a common upstream provenance chain in the PROV-O dependency graph (Moreau et al. 2011); this prevents circular corroboration in which derived sources artificially inflate the score. The evaluated implementation uses a simpler proxy: it treats distinct source institutions as independent and does not trace provenance chains.

The coefficients α_i are domain-configurable, allowing administrators to emphasise, for example, cryptographic validity in regulatory contexts or corroboration in open-data settings (Artz and Gil 2007; Carroll et al. 2005; Moreau et al. 2011).

Signature-Gated Variant In domains where cryptographic authentication is a hard prerequisite, the linear combination above may be insufficiently restrictive: an unsigned triple ($\text{sig} = 0$) with strong reputation and freshness could still receive a non-trivial trust weight. To address this, we define a *signature-gated* variant:

$$w_t^{\text{gated}} = \text{sig}(t) \cdot \sum_{i=2}^5 \frac{\alpha_i}{1 - \alpha_1} \cdot s_i(t) \quad (4)$$

Under this formulation, any triple with an invalid or absent digital signature receives $w_t^{\text{gated}} = 0$ regardless of its other provenance signals. The choice between the linear model (Equation 3) and the gated model (Equation 4) is a deployment decision: the linear model is appropriate for open-data settings where many legitimate sources

lack cryptographic infrastructure, while the gated model is preferred in regulated environments where signature validity is mandatory.

Figure 3 illustrates the trust weight computation and its integration into the verification objective.




Figure 3. Trust weight computation. Five provenance signals are combined into a scalar trust weight w_t , which modulates the neural predicate satisfaction degree $f_\theta(t)$ in the specified LTN objective (Equation 5). In the verifier evaluated here, no f_θ is learned and w_t enters the verdict rule directly.

Trust-Weighted Satisfiability Objective In a differentiable logic framework such as Logic Tensor Networks (LTNs) (Badreddine et al. 2022; Carraro et al. 2024), each grounded formula ϕ_j receives a real-valued satisfaction degree $f_\theta(\phi_j) \in [0, 1]$. We extend the standard LTN satisfiability loss by incorporating trust weights. For a knowledge base $\mathcal{K}$ containing grounded formulae derived from triples in G , the trust-weighted loss is:

$$\mathcal{L}_{\text{sat}}(\theta) = 1 - \text{SatAgg}_{\phi_j \in \mathcal{K}} [w_{t_j} \cdot f_\theta(\phi_j)] \quad (5)$$

where SatAgg is a fuzzy aggregation operator (e.g., the product t -norm or the Łukasiewicz t -norm), f_θ is the LTN grounding that maps a formula to a truth degree through learnable neural predicates (Badreddine et al. 2022), and w_{t_j} is the trust weight of the triple from which ϕ_j was derived. Were θ trained under this loss, gradient updates would prioritise knowledge authenticated by high-trust sources.

Status of Equation 5. Equation 5 specifies the *target* learned objective of the architecture; it is *not* part of the verifier evaluated in this paper. The evaluated verifier contains no learned predicate f_θ and no gradient training: it applies the trust weight w_t directly to KG evidence located by the mapping of Section 3.2, and decides the verdict with the max aggregation and thresholds of Equations 6–7. The trust weight therefore enters the *decision rule* rather than a training loss. Implementing and evaluating Equation 5 is left to future work (Section 7).

From FOL Formulae to KG Evidence Stage B must decide which triples bear on a formula and on which side. We make this step explicit as a mapping from a parsed

formula to a set of *evidence obligations*, each paired with supporting and contradicting triples (S^+ , S^-). The mapping is a deliberately partial, KG-relative interpretation of FOL: it is not full FOL entailment or OWL reasoning, and structures it cannot map safely are returned as Indeterminate with a stated reason.

- **Relation lexicon.** A fixed, declared lexicon maps each FOL predicate to KG predicates, an argument order, and a polarity: e.g. `Treats` maps to `treats` (positive) and `does_not_treat` (negative). Constants must match KG entities exactly, up to case and underscores; similarity alone never counts as evidence.
- **Atoms.** $P(a, b)$ yields S^+ = positive-polarity triples $\langle a, p, b \rangle$ and S^- = negative-polarity ones; unary class atoms $C(a)$ are checked as `is_a` triples. A negated atom $\neg P(a, b)$ swaps S^+ and S^- , so a KG triple asserting the negation supports it.
- **Restrictors.** In $\forall x (Patient(x) \wedge C(x) \rightarrow R(x, d))$, type restrictors such as $Patient(x)$ carry no evidential obligation of their own and are not checked against the KG.
- **Conditional universals (cohort policy).** A universal whose remaining condition C names a patient group and whose consequent relates that group to a drug d is checked against a KG triple stating the corresponding policy about the group, e.g. $\langle d, \text{contraindicated_in}, C \rangle$. This is a KG-relative *policy* reading: a single group-level assertion supports the claim, and it is not a proof about every patient. Universals over a drug class are checked against triples whose subject names the whole class.
- **Existentials** require a concrete witness with both class membership and the relation in the KG.
- **Abstention.** Disjunctions, nested function terms, unbound variables, predicates absent from the lexicon, and unparseable formulae yield Indeterminate with the reason recorded.

Open world. OWL and RDF adopt the open-world assumption, whereas classical FOL provers typically assume a closed world (Baader et al. 2007). The mapping operationalises the open-world reading throughout: the absence of a triple is never evidence against a claim, so missing evidence yields Indeterminate and only an explicit negative-polarity triple can falsify. A consequence is that Stage A output is KG-checkable only when its predicates and constants fall within the lexicon and the KG vocabulary; claims outside that vocabulary are Indeterminate by construction.

Worked example. The claim “Patients with chronic kidney disease should avoid NSAIDs” is grounded as $\forall x (Patient(x) \wedge ChronicKidneyDisease(x) \rightarrow ShouldAvoid(x, NSAIDs))$. $Patient(x)$ is a restrictor; $ChronicKidneyDisease$ is the cohort condition; the lexicon maps $ShouldAvoid$ to the KG predicates `contraindicated_in` and `contraindicated_for`. The obligation is therefore met by $\langle NSAIDs, \text{contraindicated_in}, ChronicKidneyDisease \rangle$, which the KG holds in the signed EMA graph, and would be contradicted by a triple such as $\langle NSAIDs, \text{safe_for}, ChronicKidneyDisease \rangle$. The trust weights of the triples on each side then enter the verdict rule below.

Conflict Detection and Resolution When the knowledge graph contains competing evidence for both ϕ and $\neg\phi$, a naive verifier might oscillate or default to whichever formula was encountered first. We formalise a conflict-aware resolution policy. Let S^+ and S^- denote the sets of triples supporting ϕ and $\neg\phi$ respectively. We compute aggregate trust for each side:

$$W^+ = \max_{t \in S^+} w_t, \quad W^- = \max_{t \in S^-} w_t \quad (6)$$

We adopt max rather than an average or sum aggregation for two reasons. First, in provenance-rich knowledge graphs, a single authoritative source (e.g., a digitally signed regulatory guideline) should be sufficient to establish a fact, regardless of how many lower-quality sources agree. Second, max is robust to knowledge-base padding attacks in which an adversary injects many low-trust triples to inflate an aggregate score; only the single highest-trust triple counts. Alternative aggregations such as weighted mean or noisy-OR (Artz and Gil 2007) are viable when evidence volume should influence the verdict; the choice is a configurable policy parameter, and empirical comparison is planned as part of future evaluation.

The verdict for a single formula is then determined by:

$$V(\phi) = \begin{cases} \text{Verified} & \text{if } W^+ > \tau \text{ and } W^+ - W^- > \delta \\ \text{Falsified} & \text{if } W^- > \tau \text{ and } W^- - W^+ > \delta \\ \text{Indeterminate} & \text{otherwise} \end{cases} \quad (7)$$

where $\tau \in (0, 1]$ is the minimum trust threshold and $\delta \in [0, 1]$ is the minimum trust margin required to resolve conflicts. Under this rule, a formula whose strongest evidence falls below τ , or whose two sides differ by at most δ , receives an **Indeterminate** verdict with a trace of the evidence on each side (Thorne et al. 2018). The rule governs only the evidence that the mapping has located; whether that evidence is sufficient in a domain sense depends on the coverage of the KG and the lexicon.

Claim-Level Verdict Aggregation When the grounding stage produces multiple formulae $\Phi = \{\phi_1, \dots, \phi_m\}$ for a single claim, the per-formula verdicts must be aggregated into a single claim-level verdict. We adopt a conservative conjunctive policy:

$$V(c) = \begin{cases} \text{Verified} & \text{if } V(\phi_j) = \text{Verified } \forall j \\ \text{Falsified} & \text{if } \exists j : V(\phi_j) = \text{Falsified and } \nexists k : V(\phi_k) = \text{Indeterminate} \\ \text{Indeterminate} & \text{otherwise} \end{cases} \quad (8)$$

Under this policy, a claim is Verified only if *every* constituent formula is verified. The policy deliberately departs from classical logic in one case. Classically, one false conjunct suffices to make a conjunction false even if another conjunct is unknown; here, if one constituent is Falsified and another is Indeterminate, the claim is Indeterminate. The rationale is epistemic rather than truth-functional. A verdict reports what the KG *establishes* about the claim as the user stated it, and in our setting the constituents of Φ

are produced by an imperfect Stage A translation, so an Indeterminate constituent may reflect a grounding or coverage gap rather than a genuine unknown. Declining to label the whole claim false keeps a Falsified verdict reserved for cases where every constituent has been resolved against the KG, which we judge the safer error in high-stakes use: a claim is not reported as refuted on the strength of a partially resolved decomposition. The trade-off is lower Falsified recall on compound claims. A classical policy, under which any Falsified conjunct falsifies the claim, is a one-line configuration change; we do not evaluate it here.

Design Properties The verdict and aggregation functions are designed to have the following properties. Each follows directly from the definitions above; we state them as design properties rather than as results.

- P1 Conservative default.** If the knowledge graph contains no relevant evidence for a formula ($S^+ = S^- = \emptyset$), then $W^+ = W^- = 0 \leq \tau$ and $V(\phi) = \text{Indeterminate}$. In particular, absence of evidence never yields Falsified.
- P2 Signature gate.** Under the gated variant (Equation 4), a triple with $\text{sig}(t) = 0$ receives $w_t^{\text{gated}} = 0$, and so cannot by itself satisfy $W > \tau$ for any $\tau > 0$.
- P3 Conservative aggregation.** For a claim decomposed into $\Phi = \{\phi_1, \dots, \phi_m\}$: (i) $V(c) = \text{Verified}$ only if every $V(\phi_j) = \text{Verified}$; (ii) if some $V(\phi_k) = \text{Indeterminate}$, then $V(c) \neq \text{Falsified}$ (Equation 8).

3.3 Stage C: Explanatory Feedback and Trust Calibration

The final stage closes the neuro-symbolic loop by translating formal proof traces into user-oriented explanations (Doshi-Velez and Kim 2017; Kim et al. 2024a). Given a verdict V for claim c , the system constructs an explanation trace $\mathcal{E}$ comprising:

$$\mathcal{E}(c) = \left\langle V, \{(t_i, w_{t_i}, \text{src}_i)\}_{i=1}^n, \text{conf} \right\rangle \quad (9)$$

where each tuple $(t_i, w_{t_i}, \text{src}_i)$ identifies a supporting or contradicting triple, its trust weight, and its provenance source, and $\text{conf} \in [0, 1]$ is the overall confidence defined as:

$$\text{conf}(\phi) = \frac{W^+}{W^+ + W^-} \quad (10)$$

This formulation maps directly to the verdict semantics: when supporting trust dominates ($W^+ \gg W^-$), confidence approaches 1; when contradicting trust dominates, confidence approaches 0; and when both sides are balanced ($W^+ \approx W^-$), confidence approaches 0.5, reflecting genuine uncertainty. Note that conf captures the *trust balance* rather than absolute trust magnitude. It is not a probability that the verdict is correct: its calibration is what E_4 would measure (Section 4), and it is not evaluated in this paper.

A central design principle is the explicit handling of ignorance and conflict. If the verifier detects competing evidence for both ϕ and $\neg\phi$ with similar trust weights (i.e., $|W^+ - W^-| \leq \delta$), the system returns an **Indeterminate** verdict along with a trace of the contradicting sources. This transparency is intended to foster “Trust Calibration,”

ensuring the user’s reliance on the system aligns with the actual quality of the underlying evidence (Romeo and Conti 2025; Carnat 2024).

3.4 End-to-End Pipeline Algorithm

Algorithm 1 formalises the complete decision procedure from claim input to verdict output.

Algorithm 1 Trust-Centric Verification Pipeline

Require: Claim c , Schema Σ , Constraints $\mathcal{C}$, Knowledge Graph G , Thresholds τ, δ

Ensure: Verdict V , Explanation $\mathcal{E}$

— Stage A: Semantic Grounding —

- 1: $\Phi \leftarrow T(c, \Sigma, \mathcal{C})$ ▷ constrained decoding (Eq. 2)
- 2: **if** some $\phi \in \Phi$ fails to parse **then**
- 3: **return** $\langle \text{Indeterminate}, \text{“Grounding failure”} \rangle$
- 4: **end if**

— Stage B: Trust-Aware Verification —

- 5: **for** each evidence obligation ϕ_j produced by the FOL→KG mapping of Φ **do**
- 6: $(S^+, S^-) \leftarrow$ lexicon lookup for ϕ_j ▷ Section 3.2; negation swaps sides
- 7: **if** ϕ_j cannot be mapped **then**
- 8: $V(\phi_j) \leftarrow \text{Indeterminate}$; **continue** ▷ reason recorded
- 9: **end if**
- 10: **for** each triple $t \in S^+ \cup S^-$ **do**
- 11: $w_t \leftarrow \sum_{i=1}^5 \alpha_i \cdot s_i(t)$ ▷ trust weight (Eq. 3)
- 12: **end for**
- 13: $W^+ \leftarrow \max_{t \in S^+} w_t$; $W^- \leftarrow \max_{t \in S^-} w_t$
- 14: $V(\phi_j) \leftarrow$ apply verdict rule (Eq. 7)
- 15: **end for**
- 16: $V \leftarrow \text{aggregate}(\{V(\phi_j)\})$ ▷ conjunctive: Indeterminate dominates

— Stage C: Explanatory Feedback —

- 17: $\mathcal{E} \leftarrow \langle V, \{(t_i, w_{t_i}, \text{src}_i)\}, \text{conf} \rangle$ ▷ Eq. 9
 - 18: **return** $\langle V, \mathcal{E} \rangle$
-

3.5 Illustrative Walkthrough

To ground the abstract specification in concrete behaviour, we trace two medical-domain claims through the implemented pipeline. The reported trust weights and verdicts are *actual outputs* of the reference implementation’s TF-IDF retrieval path, run on the curated pharmacological knowledge graph (98 triples from 9 sources, 74 of them in signed graphs) with default uniform weighting $\alpha_i = 0.2$, $\tau = 0.5$, and $\delta = 0.3$. The verifier evaluated in Section 5 replaces retrieval with the exact mapping of Section 3.2,

which accepts fewer triples as evidence. For both scenarios it returns the same verdict, W^+ , W^- and confidence.

Scenario 1: Verified Claim (Aspirin Classification)

Input. “Aspirin is a nonsteroidal anti-inflammatory drug.”

Stage A. We use the gold-standard formula $\phi_1 = \text{NSAID}(\text{Aspirin})$, as in Section 5.2; LogicLLaMA instead produces $\forall x (\text{Aspirin}(x) \rightarrow \text{NonsteroidalAntiInflammatoryDrug}(x))$, which the mapping cannot check (Section 5.3). The formula parses under the target grammar.

Stage B. The TF-IDF retriever surfaces three candidate triples: a digitally signed WHO classification, a digitally signed DrugBank annotation, and an unsigned Wikidata entry. Their trust signal breakdowns and aggregate weights, as computed by the running implementation, are shown in Table 5. The DrugBank triple states that aspirin is an antiplatelet agent, not an NSAID, so the mapping does not accept it as evidence. It illustrates why similarity-based retrieval alone is insufficient.

Table 5. Trust signal values and aggregate weights for evidence retrieved during verification of the Aspirin classification claim. Values are actual outputs from the implemented pipeline with default $\alpha_i = 0.2$.

Source	sig	rep	fresh	comp	corr	w_t
WHO	1	0.95	0.75	1.00	0.12	0.765
DrugBank	1	0.85	0.71	1.00	0.00	0.711
Wikidata	0	0.55	0.74	0.60	0.12	0.402

Computing $W^+ = \max\{0.765, 0.711, 0.402\} = 0.765$ and $W^- = 0$ (no contradicting evidence). Since $W^+ > \tau = 0.5$ and $W^+ - W^- = 0.765 > \delta = 0.3$, the verdict is **Verified**. Confidence is $W^+ / (W^+ + W^-) = 1.00$.

Stage C. The system emits an explanation trace listing the three supporting sources, their trust weights, and the underlying signal breakdown, accompanied by the reasoning “Supporting trust (0.765) exceeds the threshold with sufficient margin over contradicting evidence (0.000).” Note that the highest-trust WHO source dominates, but the architecture still surfaces lower-quality evidence with their relative weights, enabling the user to inspect the full evidential basis.

Scenario 2: Falsified Claim (NSAID Safety in CKD)

Input. “NSAIDs are safe for patients with chronic kidney disease.”

Stage A. We use the formula $\phi_2 = \text{SafeFor}(\text{NSAIDs}, \text{ChronicKidneyDisease})$, written by the authors for this demonstration; the claim is not one of the 50 evaluation claims. The formula parses under the target grammar.

Stage B. The mapping finds one supporting triple from an unsigned health blog and one contradicting triple from the European Medicines Agency (EMA), shown in Table 6.

Computing $W^+ = 0.244$ and $W^- = 0.739$. Since $W^- > \tau$ and $W^- - W^+ = 0.495 > \delta$, the verdict is **Falsified** with confidence $W^+ / (W^+ + W^-) = 0.248$.

Table 6. Trust signal values for evidence retrieved during verification of the NSAID safety claim. The EMA contradicting source dominates the unsigned blog through both reputation and provenance completeness, yielding a Falsified verdict.

Role	Source	sig	rep	fresh	comp	corr	w_t
Supporting	HealthBlog_A	0	0.20	0.82	0.20	0.00	0.244
Contradicting	EMA	1	0.92	0.78	1.00	0.00	0.739

Stage C. The explanation lists the contradicting EMA source as the dominant evidence and the unsigned blog as low-trust support, accompanied by the recommendation: “This claim appears to be factually incorrect based on available evidence.”

Linear vs Signature-Gated Modes We re-ran the same NSAID/CKD claim under the signature-gated variant (Equation 4). The unsigned HealthBlog support drops to $w_t^{\text{gated}} = 0$, eliminating the supporting evidence entirely. The verdict remains Falsified and the confidence $W^+/(W^+ + W^-)$ drops to 0.000, because no signed evidence supports the claim ($W^+ = 0$) while the signed EMA triple still contradicts it ($W^- = 0.674$). This demonstrates design property P2 in practice: in regulated mode, only cryptographically authenticated sources can influence a verdict.

Analytical Baseline Comparison To make the architectural advantages concrete, Table 7 contrasts how three approaches would handle the same Aspirin claim using the same evidence (WHO signed source vs unsigned Wikidata mirror).

Table 7. Analytical comparison of the Aspirin classification claim across three verification approaches using the same underlying evidence.

Property	Neural-only	Standard RAG	This architecture
Verdict	“Likely true”	Not provided	Verified
Formal grounding	None	None	FOL formula ϕ_1
Evidence sources	None cited	Both sources cited equally	WHO ($w_t=0.765$), Wikidata ($w_t=0.402$)
Trust differentiation	None	None	$W^+ = 0.765$
Conflict handling	Silent	User judge must	Explicit (δ threshold)
Confidence score	Token probability	None	1.00 (trust balance)
Provenance trail	None	Document links	Signed source, metadata

The comparison is analytical: we did not run either baseline (Section 4). By design, a neural-only LLM answers from its parametric knowledge and offers no evidence trail. A standard RAG pipeline that returns retrieved passages as citations, without per-source

weighting, leaves the user to judge their relative credibility. The trust-weighted verifier instead records that the WHO source is signed and highly reputed while the Wikidata mirror is unsigned, and makes this difference part of the verdict. Whether this yields better-calibrated confidence than either baseline is not tested here: our confidence score is a trust balance, and its calibration is not evaluated (Section 4).

Sensitivity to α Coefficients The verdicts above depend on the default uniform weighting $\alpha_i = 0.2$. To illustrate how the coefficients shape outcomes, consider an alternative policy that prioritises corroboration over reputation:

$$\alpha_{\text{sig}} = 0.15, \quad \alpha_{\text{rep}} = 0.10, \quad \alpha_{\text{fresh}} = 0.20, \quad \alpha_{\text{comp}} = 0.15, \quad \alpha_{\text{corr}} = 0.40$$

Under this policy, the WHO Aspirin triple (signals from Table 5) yields:

$$w'_{\text{WHO}} = 0.15(1) + 0.10(0.95) + 0.20(0.751) + 0.15(1.00) + 0.40(0.125) = 0.595$$

while the unsigned Wikidata triple yields:

$$w'_{\text{Wikidata}} = 0.15(0) + 0.10(0.55) + 0.20(0.737) + 0.15(0.60) + 0.40(0.125) = 0.342$$

The verdict remains Verified, but the absolute scores shift downward because both triples have low corroboration in our small KG. A systematic analysis over thresholds, coefficients and reputation errors on all 50 claims is reported in Section 5.4.

Extended Scenario: Genuine Evidential Conflict (Illustrative) The two preceding scenarios are real outputs from the implementation. To illustrate the architecture’s behaviour under *genuine* evidential conflict, where both sides of a dispute carry high provenance weight, we present a third scenario as an illustrative analytical example. The corresponding triples are not currently in our KG and the values below are computed analytically rather than measured.

Input Claim. An LLM generates: “Metformin can be safely co-administered with ACE inhibitors in patients with Stage 3 chronic kidney disease (CKD).”

Stage A.

$$\phi_3 = \forall x \left(\text{Patient}(x) \wedge \text{CKD3}(x) \rightarrow \text{SafeCo}(\text{Metformin}, \text{ACE.Inh}, x) \right)$$

Stage B (Illustrative). Three triples are retrieved:

- t_3 : $\langle \text{Metformin}, \text{safe_with}, \text{ACE_Inh} \rangle$ from the National Institute for Health and Care Excellence (NICE) 2024 guidelines, signed. Trust signals: sig = 1, rep = 0.92, fresh = 0.95, comp = 1.0, corr = 0.70. Result: $w_{t_3} = 0.914$.
- t_4 : $\langle \text{Metformin}, \text{contraindicated_with}, \text{ACE_Inh} \rangle$ when eGFR < 45, from a European Medicines Agency (EMA) 2024 safety bulletin, signed. Trust signals: sig = 1, rep = 0.90, fresh = 0.95, comp = 0.95, corr = 0.65. Result: $w_{t_4} = 0.890$.

- t_5 : <Metformin, risk_lactic_acidosis, CKD> from a peer-reviewed nephrology journal, signed. Trust signals: sig = 1, rep = 0.88, fresh = 0.85, comp = 0.90, corr = 0.75. Result: $w_{t_5} = 0.876$.

Triple t_3 supports ϕ_3 ; triples t_4 and t_5 support $\neg\phi_3$. Computing $W^+ = 0.914$ and $W^- = \max(0.890, 0.876) = 0.890$. Both exceed τ , but $W^+ - W^- = 0.024 < \delta$. The trust margin is insufficient to resolve the conflict, so the verdict is **Indeterminate**.

Stage C (Illustrative). The explanation surfaces the exact sources, their trust scores, and the reason the conflict could not be resolved, enabling the clinician to make an informed decision rather than receive a misleading definitive verdict. This constructed scenario illustrates the intended behaviour of the verdict rule: when reputable sources disagree by less than δ , the system declines to assert a verdict. It is not evidence that the implementation behaves this way on real conflicting data, which our evaluation does not contain.

4 Evaluation Methodology

A trust-centric verification pipeline cannot be assessed by a single aggregate metric: each stage introduces distinct failure modes that must be diagnosed independently before system-level conclusions are drawn. This section specifies a reusable, staged evaluation protocol, designed to be applicable to any pipeline conforming to the architecture in Section 3, that enables systematic comparison of present and future implementations. The protocol is inspired by the Chain-of-Verification (CoVe) principle (Dhuliawala et al. 2024), which reduces hallucination by decomposing verification into independent sub-checks. Analogously, we define two evaluation phases: first validating each component against its own gold standard, then assessing end-to-end behaviour, so that failures can be localised and attributed to specific stages rather than masked by downstream compensation.

As specified, all scoring is automated against curated reference data. Human domain experts contribute to the design of gold standards and annotation guidelines, and perform limited calibration checks on a stratified subset of outputs, but do not conduct large-scale manual scoring. This keeps the evaluation consistent, repeatable, and scalable while preventing the common criticism that outputs are evaluated solely by the system that produced them. The proof-of-concept evaluation in Section 5 applies the automated protocol with author-prepared reference data: the reference verdicts are defined relative to the knowledge graph, drafted by an LLM and adjudicated by the authors (Section 5.2). Expert curation and calibration of reference data belong to a clinical-scale evaluation, which is outside the scope of this paper.

4.1 Component-Level Checkpoints

The first evaluation phase isolates each pipeline stage. A component-level failure detected here prevents misleading system-level results downstream.

- **Knowledge Base Integrity (E_1):** This metric quantifies the “trust-readiness” of the symbolic substrate. It is measured by the percentage of triples in the

RDF/OWL Knowledge Graph that carry complete provenance annotations and valid cryptographic signatures (Moreau et al. 2011; Kale et al. 2023):

$$E_1 = \frac{|\{t \in G : \text{comp}(t) = 1.0 \wedge \text{sig}(t) = 1\}|}{|G|} \quad (11)$$

- **Grounding Accuracy (E_2):** To ensure downstream verification is meaningful, we measure the performance of the $NL \rightarrow FOL$ translator against a gold standard of expert-annotated Natural Language–FOL pairs (Yang et al. 2024; Brunello et al. 2025; Han et al. 2024):

$$E_2 = \frac{1}{N} \sum_{i=1}^N \mathbb{1}[T(c_i) \equiv \phi_i^*] \quad (12)$$

where ϕ_i^* is the gold-standard FOL formula for claim c_i . Since full logical equivalence is undecidable for general FOL, $\equiv$ is operationalised with the logical-equivalence (LE) metric of Yang et al. (2024): a truth-table comparison of the two formulae after predicates are softly aligned. LE is graded in $[0, 1]$, and E_2 reports its mean. As a complementary soft metric, we also report normalised tree-edit distance between parsed FOL syntax trees (Section 5.3), which captures partial-credit cases where the predicted formula is structurally close but not identical to the gold standard.

- **Verifier Core Performance (E_3):** We utilise a class-balanced macro-F1 score to assess the reasoning engine’s accuracy across three labels: $\{\text{Verified, Falsified, Indeterminate}\}$ (Thorne et al. 2018). The explicit inclusion of *Indeterminate* penalises the system for guessing when the Knowledge Base is incomplete or when contradictory high-trust evidence is present:

$$E_3 = \frac{1}{|L^*|} \sum_{l \in L^*} F_1(l), \quad L^* = \{l \in \{V, F, I\} : l \text{ occurs in the reference set}\} \quad (13)$$

With all three classes present, L^* has three elements and E_3 is the usual three-class macro-F1. A class with no reference examples has undefined recall, so it is excluded rather than scored as zero, and the classes used are always reported alongside E_3 .

4.2 System-Level Evaluation

Once components pass their individual checkpoints, the second phase evaluates end-to-end behaviour.

End-to-End Verdict Accuracy. The primary system-level metric is the proportion of claims for which the pipeline produces the correct final verdict (Verified, Falsified, or Indeterminate) when measured against the reference gold standard. This captures cascading errors across all three stages.

Confidence Calibration (E_4). A technically correct verdict is insufficient if the system’s expressed confidence does not match its actual accuracy. We define the Trust Calibration Score as:

$$E_4 = 1 - \frac{1}{B} \sum_{b=1}^B |\text{acc}(b) - \text{conf}(b)| \quad (14)$$

where B is the number of confidence bins, $\text{acc}(b)$ is the empirical accuracy of verdicts in bin b , and $\text{conf}(b)$ is the mean confidence of verdicts in bin b (Romeo and Conti 2025; Kim et al. 2024a). A perfect score of $E_4 = 1$ indicates that the system knows when it is right and, more importantly, when it should say “I do not know.”

Latency and Throughput. Practical deployment requires that the pipeline meets response-time constraints. We report mean and 95th-percentile latency per claim, as well as throughput (claims per second), to confirm that constrained decoding (Stage A) and trust-weighted reasoning (Stage B) remain tractable for the target domain.

Conflict Robustness. We construct an adversarial subset of claims for which the Knowledge Graph contains competing evidence (supporting triples for both ϕ and $\neg\phi$). The system is expected to return *Indeterminate* rather than an incorrect forced verdict. We report the rate at which conflicting evidence is correctly surfaced as indeterminate, measuring the system’s ability to fail safely.

Status in this paper. Section 5 reports E_1 , E_2 , E_3 and end-to-end verdict accuracy. Three protocol metrics are not computed, for the following reasons. E_4 is not reported: 35 of the 36 correct verdicts are Verified and there is a single reference Indeterminate claim, so binned calibration on these data would reflect the label distribution rather than the verifier. Latency is not reported because no timing measurement with a recorded script exists for the evaluated configuration. Conflict robustness is not reported as a rate: only four of the 50 claims meet genuinely competing evidence (a signed contraindication against an unsigned claim of safety), too few to estimate one. Their verdicts are analysed individually in Section 5.2.

4.3 Ablation Baselines

To isolate the contribution of each architectural component, we compare the full pipeline against progressively simpler variants:

- **Neural-Only:** LLM generation without symbolic verification or grounding (Brown et al. 2020).
- **Standard RAG:** Retrieval-augmented generation with citations but without formal logical verification (Patil and Gudivada 2024).
- **Unweighted Verifier:** A neuro-symbolic engine that performs logical checks but treats all facts as uniformly reliable ($w_t = 1 \forall t$), isolating the contribution of trust weighting (Badreddine et al. 2022).
- **Verdict-Only:** Full verification but no explanatory trace, isolating the effect of explanations on calibration (Doshi-Velez and Kim 2017).

- **Citations-Only:** Source citations without trust scores or conflict handling, representing a typical RAG-style explanation (Patil and Gudivada 2024).

In this paper we evaluate the Unweighted Verifier, which is the comparison that isolates the contribution of trust weighting (Section 5.2). The Neural-Only, Standard RAG, Verdict-Only and Citations-Only baselines, and system-level comparisons with methods such as CoVe (Dhuliawala et al. 2024) and KG-CRAFT (Lourenço et al. 2026), are part of the protocol but not evaluated here. The latter operate on unstructured text or on graphs constructed from claim reports, so a comparison on a 98-triple KG with gold FOL would not be like-for-like. The explanation-dependent baselines also require a user study of calibration.

5 Experimental Results

This section reports preliminary results. Section 5.1 evaluates the grounding stage in isolation on FOLIO and a domain-specific medical-claims dataset (E_2). Section 5.2 scores the verifier on the 50 medical claims against KG-relative reference verdicts, conditioned on gold-standard FOL, and compares it with the Unweighted Verifier (E_3). Section 5.3 repeats the verification with Stage A’s predicted FOL, and Section 5.4 analyses sensitivity to the trust coefficients, thresholds and reputation scores.

5.1 Grounding Stage Evaluation (FOLIO and Medical Claims)

We evaluate Stage A using LogicLLaMA (Yang et al. 2024), a Low-Rank Adaptation (LoRA) finetuned LLaMA-2-7B model trained for direct $NL \rightarrow FOL$ translation, on two datasets.

Datasets. The **FOLIO benchmark** (Han et al. 2024) provides expert-annotated natural-language–FOL pairs spanning general-domain reasoning. We use the pre-parsed FOLIO release distributed by the LogicLLaMA authors (`folio_parsed.json` in MALLS-v0), and the 254 examples are its first 254 items in file order, at which point the evaluation run was stopped; the subset was not otherwise selected. Because this release comes from the model’s own authors, overlap with LogicLLaMA’s fine-tuning data is likely and was not ruled out, so the FOLIO score may be optimistic and should be read as a reproduction check rather than a held-out estimate. To assess domain transfer, we additionally curate a **Medical Claims** dataset of 50 pharmacological assertions, derived from established pharmacological literature, with author-annotated gold-standard FOL translations. The same 50 claims are used for the verification experiments in Sections 5.2–5.4.

Metrics. Following Yang et al. (2024), we report three metrics: (i) **Logical Equivalence (LE)**, the primary E_2 metric, which compares the truth tables of the predicted and gold-standard formulae after soft alignment of their predicates, giving a score in $[0, 1]$; (ii) **BLEU** (Bilingual Evaluation Understudy) (Papineni et al. 2002), reported as a complementary surface-level measure; and (iii) **parse rate**, the percentage of predictions that yield syntactically valid FOL trees. BLEU systematically

underestimates translation quality for FOL because logically equivalent formulas may differ in variable naming and predicate ordering.

Results. Table 8 presents the headline comparison. LogicLLaMA achieves a mean LE of 0.853 on FOLIO, closely reproducing the performance reported by Yang et al. (2024). On the domain-specific medical claims, mean LE drops to 0.801, a measurable 6.1% decline. Parse rate remains 100% for both datasets, indicating that the model consistently produces syntactically valid FOL regardless of domain.

Table 8. Stage A grounding accuracy (E_2): FOLIO benchmark vs domain-specific medical claims using LogicLLaMA-7B.

Metric	FOLIO	Medical
Examples evaluated	254	50
Parse rate (%)	100.0	100.0
Mean LE (E_2)	0.853 ± 0.186	0.801 ± 0.244
Mean BLEU	0.396 ± 0.355	0.341 ± 0.326
Exact equivalence (LE=1)	114 (44.9%)	21 (42.0%)

Category-Level Analysis. The aggregate 6.1% decline masks a highly non-uniform pattern across claim categories. Table 9 reveals that the grounding model performs well on structurally simple medical claims but degrades sharply on relationally complex ones.

Table 9. Medical claims grounding accuracy by category. Categories involving multi-entity relationships (drug interactions) show the steepest performance degradation.

Category	N	Mean LE	Std
Monitoring	2	1.000	0.000
Contraindication	8	0.926	0.085
Side effect	5	0.925	0.150
Drug effect	11	0.886	0.135
Regulation	1	0.875	—
Classification	6	0.823	0.185
Treatment	11	0.688	0.254
Drug interaction	6	0.479	0.301

Claims with simple universal structure (e.g., “All beta-blockers reduce heart rate”) achieve near-perfect LE ($\mu = 0.984$), indicating that the logical pattern transfers well across domains even when the predicate vocabulary is novel. Contraindications and side effects, which map to conditional negation or conditional universal patterns, also transfer strongly (LE > 0.92). In contrast, **drug interactions** (LE = 0.479) and **treatments involving named entities** (LE = 0.688) exhibit substantial degradation. Qualitative analysis reveals two failure modes. First, the model over-generalises ground facts into universal statements: “Propranolol is a beta-blocker” is incorrectly translated as $\forall x (Propranolol(x) \rightarrow BetaBlocker(x))$. Second, multi-entity relational claims requiring nested functions or existential quantifiers are structurally malformed.

5.2 Verification Against Reference Verdicts

We evaluate Stage B on the same 50 medical claims, first conditioned on the gold-standard FOL (Stages B–C only) and then on LogicLLaMA’s own Stage A output (Section 5.3).

Knowledge Graph. The KG contains **98 triples** in **9 named graphs**, one per source: WHO, EMA, NICE, BNF, DrugBank, PubMed (reviews), Wikidata, and two deliberately low-reputation web sources (HealthBlog and HealthForum) included to test behaviour under low-quality evidence. Each named graph carries PROV-O provenance metadata (Moreau et al. 2011). The five graphs of WHO, EMA, NICE, BNF, and DrugBank (74 triples) are signed with Ed25519; the remaining four are unsigned. The signatures are *simulated*: the key pairs were generated by the authors for this study and are not issued by the named institutions, so they certify the integrity of our curated graph rather than institutional attestation. Reputation scores were set manually by the authors in the project’s source registry. Each source was first placed in a type: international authority, regulatory body, reference work, curated database, peer-reviewed review, community-edited resource, or unverified web source. It then received a score ordered by institutional authority and editorial oversight: WHO 0.95; EMA 0.92; NICE 0.91; BNF 0.88; DrugBank 0.85; PubMed reviews 0.82; Wikidata 0.55; HealthBlog 0.20; HealthForum 0.15. Only the ordering of these values was chosen deliberately; their exact values are the authors’ judgements, with no inter-rater check, and no external registry was used. The sensitivity analysis (Section 5.4) measures how much the verdicts depend on them. The other signals are computed from the graph itself. The signature signal is the named graph’s declared signed status: each of the five signed graphs carries an Ed25519 signature that verifies against its public key, but the evaluated verifier reads the declared status and does not re-verify the signature at verification time. Freshness from each named graph’s `prov:generatedAtTime`; completeness from the share of five expected PROV-O fields present; and corroboration from the share of the other sources that assert the same triple, where distinct institutions are treated as independent (a proxy for the PROV-O independence of Definition 1). Knowledge-base integrity (Equation 11) is $E_1 = 74/98 = 0.755$: the triples with complete provenance are exactly the signed ones. Because the authors curated the graph and generated its keys, E_1 here describes the construction of the graph rather than an external property.

Verifier Configuration. The verifier implements the mapping of Section 3.2 with $\alpha_i = 0.2$ uniformly, $\tau = 0.5$, $\delta = 0.3$, freshness decay $\lambda = 0.15$ and a fixed freshness reference date, so that the run is reproducible. The relation lexicon, configuration and code were frozen, and their hashes recorded, before any reference verdict existed; nothing was changed after the verdicts were written.

Reference Verdicts. Each claim received a KG-relative reference verdict: Verified if the 98-triple KG contains evidence entailing the claim under the mapping’s reading, Falsified if it contains evidence contradicting it, and Indeterminate otherwise, with the supporting triple identifiers recorded per claim. Candidate verdicts were drafted by an LLM from a written guideline and the KG catalogue, without access to system outputs,

and every claim was then reviewed and adjudicated by the first author, who changed 2 of the 50. Because the review was not independent of the draft, we report no agreement statistic. The labels are KG-relative, not clinical judgements, and will be released with the code on acceptance. The resulting distribution is 49 Verified, 1 Indeterminate and 0 Falsified. This imbalance follows from how the claim set was built, and it limits what the evaluation can show: Falsified performance cannot be measured on these claims, and Indeterminate is estimated from a single example.

Results. Table 10 reports accuracy, per-class results and E_3 . Because the Falsified class has no reference examples, E_3 (Equation 13) is averaged over the two classes that occur (Verified, Indeterminate). On gold FOL the trust-weighted verifier is correct on 36 of 50 claims (accuracy 0.72, macro-F1 0.479). All 35 of its Verified verdicts are correct; the 14 errors are all reference-Verified claims returned as Indeterminate.

These aggregate scores must be read against the label distribution. A constant answer of Verified, which consults no evidence, scores accuracy 0.98 and macro-F1 0.495 on this set (Table 10), higher than the verifier on both measures. On a reference set with 49 Verified claims, neither measure separates a verifier from a constant answer. The quantities that do are the ones a constant answer cannot provide. The first is whether asserted verdicts are supported: the verifier asserts Verified only when it has located supporting triples, and none of its 35 Verified verdicts is wrong. The second is the verifier’s behaviour on the one claim without support, which it returns as Indeterminate where the constant answer asserts it. That is a single claim, so we draw no rate from it. The present evaluation therefore measures how much KG-contained support the verifier recovers (35 of 49 claims) and the reasons it abstains. It does not measure discrimination between true and false claims.

Table 10. Verdict accuracy against the 50 KG-relative reference verdicts (49 V, 1 I, 0 F). Macro-F1 is E_3 over the classes present in the reference set (V, I). The majority-class row is a reference point that consults no evidence. “Pred.” is the predicted distribution. The last row feeds Stage A’s predicted FOL through the same verifier (Section 5.3).

Run	Acc.	Macro-F1	F1(V)	F1(I)	Pred. V/F/I
Trust-weighted, gold FOL	0.72	0.479	0.833	0.125	35/0/15
Unweighted ($w_t = 1$), gold FOL	0.66	0.448	0.790	0.105	32/0/18
Signature-gated, gold FOL	0.72	0.479	0.833	0.125	35/0/15
Majority class (all Verified)	0.98	0.495	0.990	0.000	50/0/0
Trust-weighted, Stage A FOL	0.16	0.148	0.250	0.045	7/0/43

Unweighted Verifier. The Unweighted Verifier runs the same mapping and verdict rule with $w_t = 1$ for every triple. It is correct on 33 claims against 36 for the trust-weighted verifier: the trust-weighted verifier is right and the unweighted one wrong on 4 claims, and the reverse holds for 1 claim (exact McNemar test, $p = 0.375$). The difference is therefore not statistically significant at $n = 50$. The four claims (13, 19, 25 and 26) share one mechanism: a signed EMA contraindication triple is contradicted by an unsigned, low-reputation blog or forum triple asserting that the drug is safe. With trust weights,

the EMA triple dominates ($W^+ = 0.739$ against $W^- \leq 0.244$) and the claim is Verified. With uniform weights both sides score 1, the margin is 0, and the verifier abstains. The opposite case is claim 11, supported only by an unsigned PubMed review triple with $w_t = 0.456 < \tau$: trust weighting abstains where the unweighted verifier accepts. The signature-gated variant returns the same verdict as the linear score on every claim.

Results by Complexity and Error Analysis. Table 11 breaks down the trust-weighted gold-FOL run by the logical complexity of the reference formula. The 14 errors, all Verified $\rightarrow$ Indeterminate, fall into three groups. (i) *Lexical coverage*, 6 claims (5, 15, 34, 37, 49, 50): the KG holds the relevant triple but the lexicon does not declare the predicate or entity alignment, e.g. *LowersLDL* against *lowers LDLCholesterol*, or *Paracetamol_HighDose* against the constant *Paracetamol*. (ii) *Universals with no safe mapping*, 4 claims (8, 22, 36, 39), for example monitoring rules whose KG triple takes the drug, not the cohort, as its subject (22, 39). (iii) *Structures the mapping abstains on by design*, 3 claims: nested function terms (16, 47) and a disjunction (29). The remaining error is claim 11, which falls below τ as described above. No claim received a Falsified verdict, and no reference-Verified claim was contradicted. Every error is therefore an abstention caused by the mapping’s coverage, not an incorrect assertion. These results do not show how the verifier behaves when evidence genuinely contradicts a claim, because the reference set contains no Falsified claim.

Table 11. Trust-weighted verdicts on gold FOL by FOL complexity, with the number correct against the reference verdicts (gold FOL) and the number correct when Stage A’s predicted FOL is used instead (Section 5.3). No claim was Falsified in either run.

Complexity	<i>N</i>	<i>V</i>	<i>I</i>	Correct (gold)	Correct (Stage A)
atomic	7	6	1	6	0
simple_universal	8	5	3	5	4
conditional_universal	7	4	3	4	0
conditional_negation	5	4	1	4	0
conjunction	9	7	2	7	1
conjunction_negation	3	3	0	3	0
conjunction_universal	1	1	0	1	0
triple_conjunction	2	2	0	2	1
nested_function	1	0	1	0	0
existential	2	1	1	2	1
universal_conjunction	1	1	0	1	1
universal_disjunction	1	0	1	0	0
negation	1	1	0	1	0
multi_variable_universal	1	0	1	0	0
universal_nested	1	0	1	0	0
Total	50	35	15	36	8

5.3 Cascaded Run with Predicted FOL

To measure error propagation, we pass LogicLLaMA’s predicted FOL for the same 50 claims (the predictions scored in Section 5.1) through the identical frozen verifier and score the verdicts against the same reference labels. Predictions that the verifier’s parser cannot read are Indeterminate. The lexicon was not extended to cover predicted predicate names, since that would amount to tuning after the labels existed.

On predicted FOL, accuracy is 0.16, compared with 0.72 on gold FOL (Table 10). Thirty claims that are Verified on gold FOL become Indeterminate, two move from Indeterminate to Verified, and none becomes Falsified (exact McNemar test on correctness, 30 vs. 2 discordant claims, $p < 10^{-6}$). All 7 Verified verdicts on predicted FOL are correct, so in this run grounding errors propagate as abstention rather than as wrong verdicts. Of the 43 Indeterminate verdicts, 36 arise because a predicted universal has no safe mapping, 3 from parse failures, 2 from coverage gaps, 1 from evidence below τ and 1 from a missing existential witness. The dominant cause is the first failure mode identified in Section 5.1: in 20 predictions, all of them Indeterminate, a named drug is rendered as a unary predicate (e.g. $\forall x (Aspirin(x) \rightarrow \dots)$), so the formula no longer refers to a KG constant. Predicted predicates also fuse relation and argument, as in *TreatsDepression(fluoxetine)*, which a relation lexicon cannot align. By category, treatment and contraindication claims lose every correct verdict (9 of 11 and 6 of 8 correct on gold FOL, 0 on predicted FOL), while drug-effect claims are the most robust (8 and 6 of 11). The logical-equivalence score does not predict this gap. The mean LE on these predictions is 0.801, yet the mean normalised tree-edit distance between predicted and gold trees, computed with the Zhang–Shasha algorithm with unit costs, positional variable names and normalisation by the larger tree, is 0.581 ± 0.360 over the 47 parseable predictions, with 3 exact matches. LE rewards logical shape, whereas KG-checkability also depends on predicate and constant vocabulary.

5.4 Sensitivity to Trust Coefficients, Thresholds and Reputations

We vary one factor at a time around the frozen configuration on gold FOL and report accuracy against the reference verdicts and the number of verdicts that change. These runs are descriptive; the frozen configuration was not re-selected from them.

Thresholds. Over the grid $\tau \in \{0.3, \dots, 0.7\}$ and $\delta \in \{0.1, \dots, 0.5\}$, accuracy lies between 0.68 and 0.74 and at most 2 of the 50 verdicts change. Lowering τ to 0.4 or below admits claim 11 ($W^+ = 0.456$). Raising δ to 0.5 turns claims 13 and 26 Indeterminate, because their EMA support (0.739) no longer exceeds the blog contradiction (0.244) by the margin.

Coefficients. Removing any one signal and renormalising the other four changes at most one verdict (claim 11, when signature or corroboration is removed), and accuracy lies between 0.72 and 0.74. Over 500 draws of α from a flat Dirichlet distribution, median accuracy is 0.72 and 82.8% of draws reach at least 0.70, but 9% fall below 0.50. Those draws place on average 0.55 of the weight on corroboration, which is close to zero on this KG (mean 0.028, maximum 0.125 over the 98 triples). They push most W^+ below

τ and the verifier abstains on most claims. Accuracy is negatively correlated with α_{corr} across draws (Pearson $r = -0.72$). The coefficients matter mainly through how much weight they place on signals that the KG populates.

Reputation errors. Adding independent uniform noise of ± 0.1 or ± 0.2 to the reputation of each of the nine sources (200 draws each, clipped to $[0, 1]$) changed no verdict. Of the 36 pairwise swaps of source reputations, 34 changed nothing. Swapping EMA with either low-reputation web source turns claims 13 and 26 Indeterminate (accuracy 0.68), because their contradicting triples come from the swapped sources. On this KG, therefore, verdicts are insensitive to moderate reputation error but not to a gross inversion of the sources that decide a conflict. Most supported claims have W^+ between 0.71 and 0.77 and no contradicting evidence, far from the thresholds. This insensitivity is partly a property of this small, low-conflict KG, and we do not expect it to hold for KGs with more contested evidence.

Limitations of the Current Evaluation. The 50 claims were curated by the authors from established pharmacological literature. Because the evaluation asks what the knowledge graph supports rather than what is clinically true, the reference verdicts are KG-relative; they were LLM-drafted and author-adjudicated, and neither they nor the claims have undergone independent clinical validation. The KG (98 triples, 9 sources) is small, and the reference set contains no Falsified claim and one Indeterminate claim. The results therefore measure how often the verifier recovers support that the KG contains, and where it abstains; they do not measure its ability to reject false claims, and the trust-weighting comparison rests on five discordant claims. A larger, conflict-rich benchmark with independently annotated verdicts is required before accuracy or trust-weighting conclusions can generalise.

6 Discussion and Limitations

The proposed architecture provides a foundation for trust-aware verification, but putting it into practice requires defining the operational scope and boundary conditions. This section discusses the assumptions behind our design and the limitations of the current framework.

6.1 Bounded Domains and Fact-Centric Scope

To ensure that verification remains computationally tractable and that the Knowledge Base (KB) maintains high provenance coverage, this architecture initially targets fact-centric, declarative claims within bounded domains. We assume that:

- **Data Normalisation:** The target domain’s KB can be successfully normalised into RDF/OWL with strong provenance annotations and cryptographic signatures (Hogan et al. 2021).
- **Claim Form:** Input statements are limited to single or short multi-sentence propositions about entities and relations already present in, or mappable to, the KB.

- **Logic Fragment:** The system targets a restricted first-order or description-logic fragment (Baader et al. 2007). While this ensures decidability, it currently excludes higher-order reasoning, temporal/modal logics, and complex open-ended generation judgements (Brachman and Levesque 2004).

These assumptions are reasonable for initial deployment in regulated domains such as clinical pharmacology or financial reporting, where structured knowledge bases already exist and claims tend to be factual rather than speculative. However, extending the architecture to open-domain settings will require advances in knowledge base construction, ontology learning, and the handling of vague or subjective claims.

A related practical challenge is the mapping between FOL predicates produced by Stage A and the RDF/OWL representations stored in the Knowledge Graph. FOL and OWL rest on different semantic foundations: OWL adopts the open-world assumption, whereas classical FOL theorem provers typically operate under a closed-world assumption (Baader et al. 2007). Predicate names generated by the grounding model may not align with existing ontology terms, requiring an additional alignment or ontology-matching step. This mapping could be handled by restricting the grounding vocabulary to the target ontology’s class and property names through the grammar mask $\mathcal{C}$ of Equation 2, which we specify but have not implemented, though this constraint may limit expressiveness for claims that fall outside the ontology’s current scope. The cascaded run (Section 5.3) shows that without such a restriction most Stage A output is not KG-checkable.

6.2 Trust Weight Configuration

The trust weight function (Definition 1) introduces five domain-configurable coefficients α_i . While this flexibility is a strength, allowing deployment across diverse regulatory environments, it also introduces a dependency on careful domain-expert calibration. Poorly chosen weights could systematically over- or under-trust certain source categories; Section 5.4 shows that on our KG the main risk is placing weight on a signal the KG barely populates (corroboration). We recommend a three-step calibration protocol: (1) initialise with uniform weights ($\alpha_i = 0.2$), (2) perform sensitivity analysis on a validation set, and (3) adjust based on domain expert review. Future work should investigate data-driven approaches to learning optimal α_i from labelled trust judgements.

6.3 Adversarial Robustness and Data Integrity

Any trust-centric pipeline must be resilient against malicious attempts to mislead the reasoning loop (Goodfellow et al. 2015; Pan et al. 2018). We identify three principal attack vectors and analyse the architecture’s defences against each.

Knowledge Base Poisoning. An adversary who gains write access to the Knowledge Graph could inject fabricated triples designed to sway verification verdicts. The digital signature requirement (sig) provides the first line of defence: under the gated variant (Equation 4), unsigned or tampered triples receive zero weight. Under the linear score they are down-weighted but not excluded, and the four conflicts in our evaluation are

decided by the margin δ (Section 5.2). The corroboration signal (corr) is intended as a second barrier, penalising isolated assertions that no independent source confirms. On our small KG, however, it is close to zero for almost every triple (Section 5.4), so it provides little protection there. We have not evaluated poisoning attacks.

Prompt Injection and Jailbreak Attacks. Adversaries may craft inputs designed to force the LLM into producing unverifiable or malformed formal claims (Huang et al. 2024). In the architecture, the grammar mask $\mathcal{C}$ would restrict token generation to well-formed FOL expressions over the domain schema Σ . In the evaluated implementation, which has no mask, the defence is that formulae that fail to parse, or that use vocabulary outside the relation lexicon, are returned as Indeterminate rather than verified (Section 5.3).

Trust Signal Spoofing. A sophisticated adversary might attempt to manipulate the provenance metadata itself, for example by forging digital signatures or inflating reputation scores. The architecture assumes that the cryptographic infrastructure (public-key certificates, signed named graphs (Carroll et al. 2005)) is maintained by a trusted authority external to the pipeline. If this assumption is violated, the verdict rule offers only partial protection: a spoofed high-trust triple still needs a margin δ over any legitimate contradicting evidence (Equation 7), but where no contradicting evidence exists it can determine the verdict. The evaluated verifier also reads each graph’s declared signed status rather than re-verifying signatures at verification time (Section 5.2), so a deployment would need to add that check. The integrity of the external cryptographic infrastructure remains a fundamental dependency, and the architecture does not claim to solve key management or certificate revocation, which are deferred to established standards (Artz and Gil 2007).

6.4 Computational Considerations

The pipeline introduces computational overhead at three points, each with distinct scaling characteristics.

Stage A: Constrained Decoding. In the specified architecture, the grammar mask $\mathcal{C}$ would restrict the output vocabulary at each token generation step to tokens that maintain well-formed FOL syntax, adding an $O(|\mathcal{C}|)$ operation per decoding step. For typical domain ontologies with hundreds of predicates and constants, this overhead is modest compared to the LLM’s own forward pass. The well-formedness check is a single parse per formula.

Stage B: Evidence Retrieval and Trust-Weighted Reasoning. The evaluated verifier locates evidence by exact lookup; in our implementation each lookup scans the triples, so its cost is linear in the Knowledge Graph size $|G|$, and an index keyed on subject and predicate would make it close to constant. The TF-IDF path also scores every triple, and approximate nearest-neighbour indexing over learned embeddings would reduce similarity retrieval to roughly $O(k \log |G|)$ for larger graphs. Trust weighting and the verdict rule then operate over the k candidate triples with cost linear in k , which is tractable for bounded domains where k is small (typically $k \leq 50$).

Stage C: Explanation Generation. Explanation assembly involves only formatting the verdict, the top- k evidence triples with their trust scores, and the conflict analysis, so we expect its overhead to be small relative to Stages A and B; we have not measured it.

Empirical Latency. We do not report measured latency in this paper. With Stage A invoked, we expect the cost per claim to be dominated by the LLM forward pass.

6.5 Generalisability vs. Domain Specificity

While the core mechanisms of trust-weighted verification and the proposed evaluation protocol are domain-agnostic, the underlying Knowledge Base schema and gold-standard data for semantic grounding are naturally domain-specific. Future research should investigate the feasibility of cross-domain transfer learning for the $NL \rightarrow FOL$ translator to reduce the manual overhead of domain adaptation. The modular design of the pipeline supports this: Stage B and Stage C are schema-independent, requiring only that Stage A produces well-formed FOL relative to some schema Σ .

7 Conclusion and Future Research

The adoption of Large Language Models in safety-critical processes has outpaced the development of adequate verification and trust mechanisms. In this paper, we have formalised a trust-centric neuro-symbolic architecture that bridges the gap between neural fluency, symbolic rigour, and data provenance (Marcus and Davis 2019; Besold et al. 2021). By “closing the triangle,” this framework offers a formal basis for grounding generative AI outputs in verifiable facts while accounting for the reliability of their sources.

The primary technical contributions of this work are fivefold. First, we introduce a five-signal provenance-weighted trust score (Definition 1) that integrates digital signatures, reputation, freshness, provenance completeness, and source corroboration as first-class weights in the verification decision, and we specify how the same weights would enter a differentiable satisfiability objective (Badreddine et al. 2022; Carroll et al. 2005). Second, we formalise a conflict-aware verdict function (Equation 7) that explicitly handles evidential disagreement through configurable trust thresholds. Third, we propose a staged evaluation protocol, inspired by chain-of-verification principles (Dhuliawala et al. 2024), that verifies each component independently before assessing system-level behaviour. Fourth, we provide a reference implementation, to be released on acceptance, that realises all three pipeline stages and runs end-to-end on a curated, digitally signed pharmacological knowledge graph. Fifth, we provide preliminary empirical evidence at small scale. Stage A reaches $LE = 0.853$ on FOLIO and 0.801 on medical claims, with sharp degradation on drug interactions ($LE = 0.479$). Against KG-relative reference verdicts for 50 claims, the verifier reaches accuracy 0.72 on gold FOL. It outperforms the same verifier with uniform weights (0.66), though not significantly, and errs only by abstaining. With 49 of the 50 reference verdicts Verified, this measures the recovery of support contained in the KG, not discrimination between true and false claims. The

cascaded run shows that grounding errors are the main bottleneck: accuracy is 0.16 on predicted FOL, again through abstention rather than wrong verdicts.

These evaluations are small and leave several questions open. First, the trust-weighted satisfiability objective (Equation 5) remains to be implemented and evaluated, for example in LTNtorch (Carraro et al. 2024). Second, the present reference set contains no false claims and a single Indeterminate claim, and the KG contains few genuine conflicts. Whether trust weighting improves verification in general, and how the verifier behaves when evidence contradicts a claim, therefore requires a larger benchmark with Falsified and Indeterminate claims, sources of mixed reliability, and independently annotated verdicts, together with calibration (E_4) and the remaining baselines. Third, the cascaded run shows that Stage A output is rarely KG-checkable, so grounding must be constrained to the KG vocabulary. Fourth, cross-domain generalisability remains to be tested with an independent gold standard. Additional technical directions include: improving Stage A grounding accuracy by adopting structured decomposition approaches such as the AST-guided translation of Putra et al. (2026); integrating knowledge-graph-grounded verification techniques drawing on recent advances in KG-enhanced fact-checking (Lourenço et al. 2026); and investigating data-driven approaches to learning optimal trust weight coefficients from labelled judgements. Ultimately, this research aims to transition LLMs from fluent but opaque generators into trustworthy, accountable participants in safety-critical knowledge systems (Floridi and Cowls 2019; Bellogín et al. 2024).

Declaration of conflicting interests

The Authors declare that there is no conflict of interest.

Funding

This research is conducted as part of a PhD programme at the University of Malta, Department of Artificial Intelligence, Faculty of ICT. The authors received no specific external funding for this work.

Supplemental material

The reference implementation of the trust-centric verification pipeline, including the source code for all three pipeline stages, the curated medical claims dataset with gold-standard FOL annotations and reference verdicts, the annotation guideline, the signed pharmacological knowledge graph with its public keys, the saved Stage A predictions, the scripts that produce every reported result, evaluation notebooks, and instructions for replication, will be made publicly available upon acceptance. The FOLIO benchmark used in this study is publicly available (Han et al. 2024). AI-assisted tools were used in two ways: to support $\text{\LaTeX}$ formatting and language editing of this manuscript; and to draft candidate reference verdicts for the 50 claims, which the first author then reviewed and adjudicated claim by claim (Section 5.2). The authors take full responsibility for the content of this publication.

References

- Artz D and Gil Y (2007) A survey of trust in computer science and the Semantic Web. *Journal of Web Semantics* 5(2): 58–71. DOI:10.1016/j.websem.2007.03.002.
- Asgari E, Montaña-Brown N, Dubois M, Khalil S, Balloch J, Yeung JA and Pimenta D (2025) A framework to assess clinical safety and hallucination rates of LLMs for medical text summarisation. *npj Digital Medicine* 8(1): 274. Publisher: Nature Publishing Group.
- Atanasova P, Simonsen JG, Lioma C and Augenstein I (2020) Generating Fact Checking Explanations. In: *Proceedings of the 58th Annual Meeting of the ACL*. Online: Association for Computational Linguistics, pp. 7352–7364. DOI:10.18653/v1/2020.acl-main.656.
- Augenstein I, Baldwin T, Cha M, Chakraborty T, Ciampaglia GL, Corney D, DiResta R, Ferrara E, Hale S, Halevy A, Hovy E, Ji H, Menczer F, Miguez R, Nakov P, Scheufele D, Sharma S and Zagni G (2024) Factuality challenges in the era of large language models and opportunities for fact-checking. *Nature Machine Intelligence* 6(8): 852–863. DOI: 10.1038/s42256--024--00881-z. Publisher: Nature Publishing Group.
- Baader F, Calvanese D, McGuinness DL, Nardi D and Patel-Schneider PF (eds.) (2007) *The Description Logic Handbook: Theory, Implementation and Applications*. 2 edition. Cambridge: Cambridge University Press. ISBN 978–0–521–15011–8. DOI:10.1017/CBO9780511711787.
- Bach SH, Broecheler M, Huang B and Getoor L (2017) Hinge-loss Markov random fields and probabilistic soft logic. *J. Mach. Learn. Res.* 18(1): 3846–3912.
- Badreddine S, d’Avila Garcez A, Serafini L and Spranger M (2022) Logic Tensor Networks. *Artificial Intelligence* 303: 103649. DOI:10.1016/j.artint.2021.103649.
- Banarescu L, Bonial C, Cai S, Georgescu M, Griffitt K, Hermjakob U, Knight K, Koehn P, Palmer M and Schneider N (2013) Abstract meaning representation for sembanking. In: *Proceedings of the 7th linguistic annotation workshop and interoperability with discourse*. pp. 178–186.
- Bello  n A, Grau O, Larsson S, Schimpf G, Sengupta B and Solmaz G (2024) The EU AI Act and the Wager on Trustworthy AI. *Commun. ACM* 67(12): 58–65.
- Bender EM, Gebru T, McMillan-Major A and Shmitchell S (2021) On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? In: *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, FAccT ’21. New York, NY, USA: Association for Computing Machinery. ISBN 978–1–4503–8309–7, pp. 610–623. DOI:10.1145/3442188.3445922.
- Besold TR, d’Avila Garcez A, Bader S, Bowman H, Domingos P, Hitzler P, Kuhnberger KU, Lamb LC, Lima PMV, de Penning L, Pinkas G, Poon H and Zaverucha G (2021) Neural-Symbolic Learning and Reasoning: A Survey and Interpretation. In: *Neuro-Symbolic Artificial Intelligence: The State of the Art*. IOS Press, pp. 1–51. DOI:10.3233/FAIA210348.
- Bo JY, Wan S and Anderson A (2024) To Rely or Not to Rely? Evaluating Interventions for Appropriate Reliance on Large Language Models. DOI:10.48550/arXiv.2412.15584. ArXiv:2412.15584 [cs].
- Bordes A, Usunier N, Garcia-Duran A, Weston J and Yakhnenko O (2013) Translating Embeddings for Modeling Multi-relational Data. In: *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc.

- Bos J (2008) Wide-Coverage Semantic Analysis with Boxer. In: Bos J and Delmonte R (eds.) *Semantics in Text Processing. STEP 2008 Conference Proceedings*. College Publications, pp. 277–286.
- Brachman R and Levesque H (2004) *Knowledge Representation and Reasoning*. 1st edition edition. Amsterdam Boston: Morgan Kaufmann. ISBN 978–1–55860–932–7.
- Brown T, Mann B, Ryder N, Subbiah M, Kaplan JD, Dhariwal P, Neelakantan A, Shyam P, Sastry G, Askell A, Agarwal S, Herbert-Voss A, Krueger G, Henighan T, Child R, Ramesh A, Ziegler D, Wu J, Winter C, Hesse C, Chen M, Sigler E, Litwin M, Gray S, Chess B, Clark J, Berner C, McCandlish S, Radford A, Sutskever I and Amodei D (2020) Language Models are Few-Shot Learners. In: *Advances in Neural Information Processing Systems*, volume 33. Curran Associates, Inc., pp. 1877–1901.
- Brunello A, Ferrarese R, Geatti L, Marzano E, Montanari A and Saccomanno N (2025) Evaluating LLMs Capabilities at Natural Language to Logic Translation: A Preliminary Investigation. In: *Artificial Intelligence and Formal Verification, Logic, Automata, Synthesis, CEUR WORKSHOP PROCEEDINGS*, volume 3904. CEUR-WS, p. 8.
- Cao L (2024) GraphReason: Enhancing Reasoning Capabilities of Large Language Models through A Graph-Based Verification Approach. In: *Proceedings of the 2nd Workshop on Natural Language Reasoning and Structured Explanations (@ACL 2024)*. Bangkok, Thailand: Association for Computational Linguistics, pp. 1–12.
- Carnat I (2024) Human, all too human: accounting for automation bias in generative large language models. *International Data Privacy Law* : ipae018DOI:10.1093/idpl/ipae018.
- Carraro T, Serafini L and Aioli F (2024) LTNtorch: PyTorch Implementation of Logic Tensor Networks. ArXiv:2409.16045 [cs].
- Carroll JJ, Bizer C, Hayes P and Stickler P (2005) Named graphs, provenance and trust. In: *Proceedings of the 14th international conference on World Wide Web, WWW '05*. New York, NY, USA: Association for Computing Machinery, pp. 613–622.
- Cheng K, Gao J, Jiang J and Choudhary A (2024) Neural-Symbolic Methods for Knowledge Graph Reasoning: A Survey. *ACM Transactions on Knowledge Discovery from Data* 18(9). DOI: 10.1145/3686806.
- Choudhury A and Chaudhry Z (2024) Large Language Models and User Trust: Consequence of Self-Referential Learning Loop and the Deskilling of Health Care Professionals. *Journal of Medical Internet Research* 26: e56764. DOI:10.2196/56764.
- Ciampaglia GL, Shiralkar P, Rocha LM, Bollen J, Menczer F and Flammini A (2015) Computational Fact Checking from Knowledge Networks. *PLOS ONE* 10(6): e0128193. DOI: 10.1371/journal.pone.0128193. Publisher: Public Library of Science.
- Demartini G, Mizzaro S and Spina D (2020) Human-in-the-loop Artificial Intelligence for Fighting Online Misinformation: Challenges and Opportunities. *IEEE Data Eng. Bull.* 43: 65–74.
- Devlin J, Chang MW, Lee K and Toutanova K (2019) BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In: *North American Chapter of the Association for Computational Linguistics*.
- Dhuliawala S, Komeili M, Xu J, Raileanu R, Li X, Celikyilmaz A and Weston J (2024) Chain-of-Verification Reduces Hallucination in Large Language Models. In: *Findings of the Association*

- for *Computational Linguistics: ACL 2024*. Bangkok, Thailand: Association for Computational Linguistics, pp. 3563–3578. DOI:10.18653/v1/2024.findings-acl.212.
- Dietvorst BJ, Simmons JP and Massey C (2015) Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General* 144(1): 114–126. DOI:10.1037/xge0000033.
- Doshi-Velez F and Kim B (2017) Towards A Rigorous Science of Interpretable Machine Learning. *arXiv: Machine Learning*.
- Floridi L and Cowls J (2019) A Unified Framework of Five Principles for AI in Society. *Harvard Data Science Review* 1(1). DOI:10.1162/99608f92.8cd550d1. Publisher: The MIT Press.
- Garcez Ad and Lamb LC (2023) Neurosymbolic AI: the 3rd wave. *Artificial Intelligence Review* 56(11): 12387–12406. DOI:10.1007/s10462--023--10448-w.
- Garcez ASd, Lamb LC and Gabbay DM (2009) *Neural-Symbolic Cognitive Reasoning*. Cognitive Technologies. Berlin, Heidelberg: Springer. ISBN 978–3–540–73245–7 978–3–540–73246–4.
- Golbeck J (2006) Combining Provenance with Trust in Social Networks for Semantic Web Content Filtering. In: Moreau L and Foster I (eds.) *Provenance and Annotation of Data*. Berlin, Heidelberg: Springer. ISBN 978–3–540–46303–0, pp. 101–108. DOI:10.1007/11890850.12.
- Goodfellow IJ, Shlens J and Szegedy C (2015) Explaining and Harnessing Adversarial Examples. In: Bengio Y and LeCun Y (eds.) *3rd International Conference on Learning Representations, ICLR 2015*. URL <http://arxiv.org/abs/1412.6572>.
- Han S, Schoelkopf H, Zhao Y, Qi Z, Riddell M, Zhou W, Coady J, Peng D, Qiao Y, Benson L, Sun L, Wardle-Solano A, Szabó H, Zubova E, Burtell M, Fan J, Liu Y, Wong B, Sailor M, Ni A, Nan L, Kasai J, Yu T, Zhang R, Fabbri A, Kryscinski WM, Yavuz S, Liu Y, Lin XV, Joty S, Zhou Y, Xiong C, Ying R, Cohan A and Radev D (2024) FOLIO: Natural Language Reasoning with First-Order Logic. In: *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*. Miami, Florida, USA: Association for Computational Linguistics, pp. 22017–22031. DOI:10.18653/v1/2024.emnlp-main.1229.
- Hanselowski A, Stab C, Schulz C, Li Z and Gurevych I (2019) A Richly Annotated Corpus for Different Tasks in Automated Fact-Checking. In: *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*. Hong Kong, China: Association for Computational Linguistics, pp. 493–503. DOI:10.18653/v1/K19--1046.
- Harth A, Janik M and Staab S (2011) Semantic Web Architecture. In: Domingue J, Fensel D and Hendler JA (eds.) *Handbook of Semantic Web Technologies*. Berlin, Heidelberg: Springer. ISBN 978–3–540–92913–0, pp. 43–75. DOI:10.1007/978--3--540--92913--0.2.
- Hartig O (2009) Querying Trust in RDF Data with tSPARQL. In: Aroyo L, Traverso P, Ciravegna F, Cimiano P, Heath T, Hyvönen E, Mizoguchi R, Oren E, Sabou M and Simperl E (eds.) *The Semantic Web: Research and Applications*. Berlin, Heidelberg: Springer, pp. 5–20.
- Hitzler P (2021) A review of the semantic web field. *Communications of the ACM* 64(2): 76–83. DOI:10.1145/3397512.
- Hogan A, Blomqvist E, Cochez M, D’amato C, Melo GD, Gutierrez C, Kirrane S, Gayo JEL, Navigli R, Neumaier S, Ngomo ACN, Polleres A, Rashid SM, Rula A, Schmelzeisen L, Sequeda J, Staab S and Zimmermann A (2021) Knowledge Graphs. *ACM Comput. Surv.* 54(4): 71:1–71:37. DOI:10.1145/3447772.

- Horrocks I (2005) OWL: A Description Logic Based Ontology Language. In: van Beek P (ed.) *Principles and Practice of Constraint Programming - CP 2005*. Berlin, Heidelberg: Springer, pp. 5–8.
- Hou K, Li J, Liu Y, Sun S, Zhang H and Jiang H (2024) KG-EGV: A Framework for Question Answering with Integrated Knowledge Graphs and Large Language Models. *Electronics* 13(23): 4835. DOI:10.3390/electronics13234835. Number: 23; Publisher: Multidisciplinary Digital Publishing Institute.
- Huang J and Fox MS (2006) An ontology of trust: formal semantics and transitivity. In: *Proceedings of the 8th international conference on Electronic commerce, ICEC '06*. New York, NY, USA: Association for Computing Machinery. ISBN 978–1–59593–392–8, pp. 259–270. DOI:10.1145/1151454.1151499.
- Huang Y, Sun L, Wang H, Wu S, Zhang Q, Li Y, Gao C, Huang Y, Lyu W, Zhang Y, Li X, Sun H, Liu Z, Liu Y, Wang Y, Zhang Z, Vidgen B, Kailkhura B, Xiong C, Xiao C, Li C, Xing EP, Huang F, Liu H, Ji H, Wang H, Zhang H, Yao H, Kellis M, Zitnik M, Jiang M, Bansal M, Zou J, Pei J, Liu J, Gao J, Han J, Zhao J, Tang J, Wang J, Vanschoren J, Mitchell J, Shu K, Xu K, Chang KW, He L, Huang L, Backes M, Gong NZ, Yu PS, Chen PY, Gu Q, Xu R, Ying R, Ji S, Jana S, Chen T, Liu T, Zhou T, Wang WY, Li X, Zhang X, Wang X, Xie X, Chen X, Wang X, Liu Y, Ye Y, Cao Y, Chen Y and Zhao Y (2024) Position: TrustLLM: Trustworthiness in Large Language Models. In: *Proceedings of the 41st International Conference on Machine Learning*. PMLR, pp. 20166–20270. ISSN: 2640–3498.
- Ji Z, Lee N, Frieske R, Yu T, Su D, Xu Y, Ishii E, Bang YJ, Madotto A and Fung P (2023) Survey of Hallucination in Natural Language Generation. *ACM Comput. Surv.* 55(12): 248:1–248:38. DOI:10.1145/3571730.
- Kale A, Nguyen T, Harris FC Jr, Li C, Zhang J and Ma X (2023) Provenance documentation to enable explainable and trustworthy AI: A literature review. *Data Intelligence* 5(1): 139–162.
- Kim J, Maathuis H and Sent D (2024a) Human-centered evaluation of explainable AI applications: a systematic review. *Frontiers in Artificial Intelligence* 7. DOI:10.3389/frai.2024.1456486. Publisher: Frontiers.
- Kim SSY, Liao QV, Vorvoreanu M, Ballard S and Vaughan JW (2024b) “I’m Not Sure, But...”: Examining the Impact of Large Language Models’ Uncertainty Expression on User Reliance and Trust. In: *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency, FAccT ’24*. New York, NY, USA: Association for Computing Machinery. ISBN 9798400704505, pp. 822–835. DOI:10.1145/3630106.3658941.
- Koivunen MR and Miller E (2001) The Semantic Web “layer cake” graphic. URL <https://www.w3.org/2001/12/semweb-fin/w3csw>.
- Konstantinovskiy L, Price O, Babakar M and Zubiaga A (2021) Toward Automated Factchecking: Developing an Annotation Schema and Benchmark for Consistent Automated Claim Detection. *Digital Threats* 2(2): 14:1–14:16. DOI:10.1145/3412869.
- Li J, Cheng X, Zhao X, Nie JY and Wen JR (2023) HaluEval: A Large-Scale Hallucination Evaluation Benchmark for Large Language Models. In: *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, pp. 6449–6464. DOI:10.18653/v1/2023.emnlp-main.397.

- Li Z, Hua W, Wang H, Zhu H and Zhang Y (2024) Formal-LLM: Integrating Formal Language and Natural Language for Controllable LLM-based Agents. DOI:10.48550/arXiv.2402.00798. ArXiv:2402.00798 [cs].
- Lin Z, Guan S, Zhang W, Zhang H, Li Y and Zhang H (2024) Towards trustworthy LLMs: a review on debiasing and dehallucinating in large language models. *Artificial Intelligence Review* 57(9): 243. DOI:10.1007/s10462--024--10896-y.
- Lipton ZC (2018) The Mythos of Model Interpretability: In machine learning, the concept of interpretability is both important and slippery. *Queue* 16(3): 31–57. DOI:10.1145/3236386.3241340.
- Lourenço VN, Paes A, Weyde T, Depeige A and Dubey M (2026) KG-CRAFT: Knowledge Graph-based Contrastive Reasoning with LLMs for Enhancing Automated Fact-checking. In: *Proceedings of the 19th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*. March 24–29, 2026: Association for Computational Linguistics, pp. 6419–6439. DOI:10.18653/v1/2026.eacl-long.302.
- Lu Q, Li R, Sagheb E, Wen A, Wang J, Wang L, Fan JW and Liu H (2025) Explainable Diagnosis Prediction through Neuro-Symbolic Integration. DOI:10.48550/arXiv.2410.01855. ArXiv:2410.01855 [cs].
- Madaan A, Tandon N, Gupta P, Hallinan S, Gao L, Wiegrefe S, Alon U, Dziri N, Prabhunoy S, Yang Y, Gupta S, Majumder BP, Hermann K, Welleck S, Yazdanbakhsh A and Clark P (2023) Self-Refine: Iterative Refinement with Self-Feedback. In: *Advances in Neural Information Processing Systems*, volume 36. Curran Associates, Inc., pp. 46534–46594.
- Manhaeve R, Dumančić S, Kimmig A, Demeester T and De Raedt L (2021a) Neural probabilistic logic programming in DeepProbLog. *Artificial Intelligence* 298: 103504. DOI:10.1016/j.artint.2021.103504.
- Manhaeve R, Marra G and Raedt LD (2021b) Approximate Inference for Neural Probabilistic Logic Programming. *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning* 18(1): 475–486.
- Marcus G and Davis E (2019) *Rebooting AI: Building Artificial Intelligence We Can Trust*. Vintage.
- Min S, Krishna K, Lyu X, Lewis M, Yih Wt, Koh P, Iyyer M, Zettlemoyer L and Hajishirzi H (2023) FActScore: Fine-grained Atomic Evaluation of Factual Precision in Long Form Text Generation. In: *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, pp. 12076–12100. DOI: 10.18653/v1/2023.emnlp-main.741.
- Mirzaei T, Amini L and Esmailzadeh P (2024) Clinician voices on ethics of LLM integration in healthcare: a thematic analysis of ethical concerns and implications. *BMC Medical Informatics and Decision Making* 24(1): 250. DOI:10.1186/s12911--024--02656--3.
- Moglia A, Georgiou K, Cerveri P, Mainardi L, Satava RM and Cuschieri A (2024) Large language models in healthcare: from a systematic review on medical examinations to a comparative analysis on fundamentals of robotic surgery online test. *Artificial Intelligence Review* 57(9): 231. DOI:10.1007/s10462--024--10849--5.
- Moreau L, Clifford B, Freire J, Futrelle J, Gil Y, Groth P, Kwasnikowska N, Miles S, Missier P, Myers J, Plale B, Simmhan Y, Stephan E and den Bussche JV (2011) The Open Provenance

- Model core specification (v1.1). *Future Generation Computer Systems* 27(6): 743–756. DOI: 10.1016/j.future.2010.07.005.
- Olausson T, Gu A, Lipkin B, Zhang C, Solar-Lezama A, Tenenbaum J and Levy R (2023) LINC: A Neurosymbolic Approach for Logical Reasoning by Combining Language Models with First-Order Logic Provers. In: *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*. Singapore: Association for Computational Linguistics, pp. 5153–5176. DOI:10.18653/v1/2023.emnlp-main.313.
- OpenAI, Achiam J, Adler S, Agarwal S, Ahmad L, Akkaya I, Aleman FL, Almeida D, Altenschmidt J, Altman S et al. (2024) GPT-4 Technical Report. DOI:10.48550/arXiv.2303.08774. ArXiv:2303.08774 [cs].
- Ouyang L, Wu J, Jiang X, Almeida D, Wainwright C, Mishkin P, Zhang C, Agarwal S, Slama K, Ray A, Schulman J, Hilton J, Kelton F, Miller L, Simens M, Askell A, Welinder P, Christiano PF, Leike J and Lowe R (2022) Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems* 35: 27730–27744.
- Pan L, Albalak A, Wang X and Wang W (2023) Logic-LM: Empowering Large Language Models with Symbolic Solvers for Faithful Logical Reasoning. In: *Findings of the Association for Computational Linguistics: EMNLP 2023*. Singapore: Association for Computational Linguistics, pp. 3806–3824. DOI:10.18653/v1/2023.findings-emnlp.248.
- Pan S, Hu R, Long G, Jiang J, Yao L and Zhang C (2018) Adversarially regularized graph autoencoder for graph embedding. In: *Proceedings of the 27th International Joint Conference on Artificial Intelligence, IJCAI'18*. Stockholm, Sweden: AAAI Press. ISBN 978-0-9992411-2-7, pp. 2609–2615.
- Pan S, Luo L, Wang Y, Chen C, Wang J and Wu X (2024) Unifying Large Language Models and Knowledge Graphs: A Roadmap. *IEEE Transactions on Knowledge and Data Engineering* 36(7): 3580–3599. DOI:10.1109/TKDE.2024.3352100.
- Papineni K, Roukos S, Ward T and Zhu WJ (2002) BLEU: A method for automatic evaluation of machine translation. In: *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, pp. 311–318.
- Patil R and Gudivada V (2024) A Review of Current Trends, Techniques, and Challenges in Large Language Models (LLMs). *Applied Sciences* 14(5): 2074. DOI:10.3390/app14052074.
- Putra RR, Basuki RSP, Cheng Y and Gao P (2026) NL2LOGIC: AST-Guided Translation of Natural Language into First-Order Logic with Large Language Models. In: *Findings of the Association for Computational Linguistics: EACL 2026*. Association for Computational Linguistics, pp. 6035–6051. DOI:10.18653/v1/2026.findings-eacl.317.
- Qiu X, Sun T, Xu Y, Shao Y, Dai N and Huang X (2020) Pre-trained models for natural language processing: A survey. *Science China Technological Sciences* 63(10): 1872–1897. DOI: 10.1007/s11431--020--1647--3.
- Quan X, Valentino M, Dennis LA and Freitas A (2024) Verification and Refinement of Natural Language Explanations through LLM-Symbolic Theorem Proving. In: *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*. Miami, Florida, USA: Association for Computational Linguistics, pp. 2933–2958. DOI:10.18653/v1/2024.emnlp-main.172.

- Rahimzadeh Holagh S and Mohebbi K (2019) A glimpse of Semantic Web trust. *SN Applied Sciences* 1(12): 1732. DOI:10.1007/s42452--019--1598--6.
- Riegel R, Gray A, Luus F, Khan N, Makondo N, Yunus Akhalwaya I, Qian H, Fagin R, Barahona F, Sharma U, Ikbal S, Karanam H, Neelam S, Likhyan A and Srivastava S (2020) Logical Neural Networks.
- Romeo G and Conti D (2025) Exploring automation bias in human–AI collaboration: a review and implications for explainable AI. *AI & SOCIETY*.
- Samadi M, Talukdar P, Veloso M and Blum M (2016) ClaimEval: Integrated and flexible framework for claim evaluation using credibility of sources. In: *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30. pp. 222–228. DOI:10.1609/aaai.v30i1.9996.
- Scherp A, Groener G, Škoda P, Hose K and Vidal ME (2024) Semantic Web: Past, Present, and Future. *Transactions on Graph Data and Knowledge (TGDK)* 2(1): 3:1–3:37. DOI: 10.4230/TGDK.2.1.3.
- Shadbolt N, Berners-Lee T and Hall W (2006) The Semantic Web Revisited. *IEEE Intelligent Systems* 21(3): 96–101. DOI:10.1109/MIS.2006.62.
- Shi B and Weninger T (2016) Fact Checking in Heterogeneous Information Networks. In: *Proceedings of the 25th International Conference Companion on World Wide Web, WWW '16 Companion*. Republic and Canton of Geneva, CHE: International World Wide Web Conferences Steering Committee. ISBN 978–1–4503–4144–8, pp. 101–102. DOI:10.1145/2872518.2889354.
- Shinn N, Cassano F, Gopinath A, Narasimhan K and Yao S (2023) Reflexion: Language Agents with Verbal Reinforcement Learning. In: *Advances in Neural Information Processing Systems*, volume 36. Curran Associates, Inc., pp. 8634–8652.
- Stanojević M and Steedman M (2019) CCG Parsing Algorithm with Incremental Tree Rotation. In: *Proceedings of the 2019 Conference of the NAACL: Human Language Technologies, Volume 1*. Minneapolis, Minnesota: Association for Computational Linguistics, pp. 228–239. DOI: 10.18653/v1/N19--1020.
- Tang L, Laban P and Durrett G (2024) MiniCheck: Efficient Fact-Checking of LLMs on Grounding Documents. In: *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*. Miami, Florida, USA: Association for Computational Linguistics, pp. 8818–8847. DOI:10.18653/v1/2024.emnlp-main.499.
- Thorne J, Vlachos A, Christodoulopoulos C and Mittal A (2018) FEVER: a Large-scale Dataset for Fact Extraction and VERification. In: *Proceedings of the 2018 Conference of the NAACL: Human Language Technologies, Volume 1 (Long Papers)*. New Orleans, Louisiana: Association for Computational Linguistics, pp. 809–819. DOI:10.18653/v1/N18--1074.
- Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, Kaiser L and Polosukhin I (2017) Attention is All you Need. In: *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Wadden D, Lin S, Lo K, Wang LL, van Zuylen M, Cohan A and Hajishirzi H (2020) Fact or Fiction: Verifying Scientific Claims. In: *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Online: Association for Computational Linguistics, pp. 7534–7550. DOI:10.18653/v1/2020.emnlp-main.609.

- Wang WY (2017) “Liar, Liar Pants on Fire”: A New Benchmark Dataset for Fake News Detection. In: *Proceedings of the 55th Annual Meeting of the ACL (Volume 2: Short Papers)*. Vancouver, Canada: Association for Computational Linguistics, pp. 422–426. DOI: 10.18653/v1/P17--2067.
- Wang X, Wei J, Schuurmans D, Le QV, Chi EH, Narang S, Chowdhery A and Zhou D (2023) Self-Consistency Improves Chain of Thought Reasoning in Language Models. In: *Proceedings of the Eleventh International Conference on Learning Representations*.
- Wei J, Yang C, Song X, Lu Y, Hu N, Tran D, Peng D, Liu R, Huang D, Du C and Le QV (2024) Long-form factuality in large language models. In: *Advances in Neural Information Processing Systems*, volume 37. DOI:10.52202/079017-2567.
- Winters T, Marra G, Manhaeve R and Raedt LD (2022) DeepStochLog: Neural Stochastic Logic Programming. *Proceedings of the AAAI Conference on Artificial Intelligence* 36(9): 10090–10100.
- Wu K, Wu E, Wei K, Zhang A, Casasola A, Nguyen T, Riantawan S, Shi P, Ho D and Zou J (2025) An automated framework for assessing how well LLMs cite relevant medical references. *Nature Communications* 16(1): 3615. Publisher: Nature Publishing Group.
- Yang Y, Xiong S, Payani A, Shareghi E and Fekri F (2024) Harnessing the Power of Large Language Models for Natural Language to First-Order Logic Translation. In: *Proceedings of the 62nd Annual Meeting of the ACL (Volume 1: Long Papers)*. Bangkok, Thailand: Association for Computational Linguistics, pp. 6942–6959. DOI:10.18653/v1/2024.acl-long.375.
- Yang Z, Ishay A and Lee J (2021) NeurASP: embracing neural networks into answer set programming. In: *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI’20*. ISBN 978–0–9992411–6–5, pp. 1755–1762.
- Zhao H, Chen H, Yang F, Liu N, Deng H, Cai H, Wang S, Yin D and Du M (2024) Explainability for Large Language Models: A Survey. *ACM Trans. Intell. Syst. Technol.* 15(2): 20:1–20:38. DOI:10.1145/3639372.

Author biography

David Farrugia is a data scientist and Ph.D. candidate in Artificial Intelligence at the University of Malta. He has worked across diverse industries including gaming, manufacturing, customer relationship management, affiliate marketing, and anti-fraud. His research explores the intersection of applied data science and academic inquiry, with a focus on neuro-symbolic AI, trustworthy language technologies, and provenance-aware knowledge systems.

Author biography

Alexiei Dingli is a Professor of Artificial Intelligence at the University of Malta with over two decades of experience in the field. His work has been recognised as world-class by international bodies, earning awards from the European Space Agency, the World Intellectual Property Organization, and the United Nations. He has helped numerous organisations implement AI solutions and has been a core member of the Malta.AI task force. He holds an MBA from Grenoble Business School specialising in Technology Management.