# PROPAI: Propositional Reasoning Artificial Intelligence

### *A New Category of AI Architecture with Dual Representation for Truth-Critical Applications*

Francesca Lazzeri, Herbert L. Roitblat, & Chris LaCour

*Reliath AI*

## Abstract

We introduce Propositional-Reasoning Artificial Intelligence (PROPAI), a fundamentally new class of AI systems designed to reason over verifiable facts rather than predict text. While Large Language Models (LLMs) have revolutionized natural language processing through probabilistic next-token prediction, they face an inherent architectural limitation: they cannot represent truth. PROPAI addresses this limitation through a patent-pending dual-representation architecture that simultaneously represents facts as both logical units suitable for graph-based reasoning and semantic vectors ("truth profiles") enabling flexible query handling. This approach makes hallucinations impossible by design rather than merely reducing its probability. We present the technical foundations of PROPAI, including the plausibility scoring mechanism, hypothesis generation capabilities, and concrete performance characteristics. We demonstrate PROPAI's superiority over Retrieval Augmented Generation (RAG) approaches through specific examples and establish its position as a complementary technology to LLMs, optimized for accuracy and correctness in high-stakes applications where LLMs are architecturally unsuitable.

**Keywords:** Proposition-based reasoning, knowledge graphs, semantic vectors, truth representation, hallucination elimination, dual representation architecture, plausibility scoring, Large Language Models, PROPAI, AI

## 1. Introduction

The rapid adoption of LLMs has transformed how we interact with artificial intelligence. From GPT-4 to Claude, these systems demonstrate remarkable fluency in generating human-like text across diverse domains. However, their success in language generation masks a fundamental architectural limitation that becomes critical in enterprise deployments: LLMs have no representation of truth.

According to the GPT-4 Technical Report [1], "GPT-4 is a Transformer-style model pre-trained to predict the next token in a document." This architectural choice optimizes for linguistic plausibility, not factual accuracy. Individual tokens-the fundamental units of analysis in transformer-based models-have no truth value. The word "purple" can be neither true nor false. Without a representation of truth at the foundational level, these systems cannot distinguish between fluent fabrication and accurate information.

OpenAI itself acknowledged this fundamental limitation in 2024 [2], confirming that hallucinations are inevitable in transformer-based language models [3]. This admission represents a watershed moment: the leading organization in LLM development has conceded that the problem is architectural, not solvable through scaling, fine-tuning, or guardrails.

This limitation creates a critical gap in AI deployment. Enterprise surveys consistently report that the primary barrier to AI adoption is not cost or complexity. It is trustworthiness. In regulated industries like healthcare, finance, and legal services—in which accuracy is non-negotiable and errors carry severe legal and financial consequences—LLMs face an insurmountable challenge. The collection of over 900 legal cases where lawyers submitted AI-hallucinated documents [4, 5] demonstrates that current approaches to managing hallucinations are insufficient.

We introduce Proposition-Reasoning Artificial Intelligence (PROPAI) as a new category of AI systems specifically designed for truth-critical applications. Rather than competing with LLMs on language generation, PROPAI establishes a complementary paradigm: where LLMs excel at creativity and fluency, PROPAI optimizes for accuracy and correctness. This is not an incremental improvement to existing architectures but a paradigm shift-from predicting text to reasoning over facts.

## 1.1 The Hallucination Problem Is Architectural, Not Scalable

The hallucination problem in LLMs is not a bug to be patched; it is a mathematical inevitability of their architecture. Token-based language models predict the next word based on statistical patterns learned from training data. Without a representation of truth, these systems cannot distinguish between fluent fabrication and accurate information.

OpenAI's own benchmarks [2], following their attempts to address hallucinations through reasoning models, provide stark evidence. On their proprietary benchmark designed to test factual accuracy, gpt-5-thinking-mini achieved only 22% accuracy with a 52% abstention rate-meaning it refused to answer more than half the questions. The o4-mini model, which attempted to answer more questions, achieved 24% accuracy with a 75% error rate. These results, from the leading organization in LLM development using their own evaluation framework, demonstrate that hallucination is not a solvable problem within the current paradigm.

The economics further illustrate the impossibility of scaling out of this problem. The GPT-4 technical report [1] revealed that reducing the error rate by just 1 bit required 10,000 times more computational resources than GPT-3.5. Extrapolating this exponential relationship, the next 1-bit improvement would demand approximately 100 quadrillion times more computational capacity, far exceeding even the combined $344 billion infrastructure investment the big four AI companies (Google, Microsoft, OpenAI, and Meta) plan through 2026.

Even if such computational resources were available, the fundamental problem would remain: no amount of compute can convert token probabilities into truth. As we demonstrate in Section 3, this requires a different unit of analysis entirely.

## 1.2 Why Retrieval Augmented Generation Cannot Solve the Problem

RAG [6] represents the most prominent attempt to address LLM hallucinations by constraining generation to specific source documents. However, RAG's complexity and continued high error rates demonstrate that it fundamentally cannot solve the truth representation problem.

The RAG pipeline requires multiple sequential processes, each of which must execute flawlessly:

- Document chunking: Breaking source documents into subsections of appropriate size
- Chunk embedding: Encoding each chunk into a vector representation
- Query embedding: Converting user queries into comparable vector representations
- Similarity search: Identifying chunks with vectors closest to the query vector
- Rank fusion
- Reranking
- Answer generation: Prompting an LLM to extract answers from retrieved chunks

RAG cannot produce correct output unless each of the steps are executed flawlessly.

Moreover, none of the steps in the RAG pipeline utilize any representation of truth. Chunk boundaries are arbitrary, embeddings approximate semantic similarity without truth values, and the LLM generation step reintroduces the fundamental hallucination problem. Multiple RAG extensions have been proposed, but all suffer from the same architectural limitation: they attempt to constrain a system that fundamentally cannot represent truth.

# 2. The PROPAI Architecture: Propositions as First-Class Objects

PROPAI represents a fundamental departure from probabilistic language modeling. Instead of predicting text, PROPAI systems reason over verified facts represented in a patent-pending dual form. This section establishes the core architectural principles that distinguish PROPAI from existing approaches.

## 2.1 The Nature of Truth

PROPAI does not have, nor does it require, privileged access to absolute truth. Rather, it begins with a customer-selected source of truth.

Truth has the potential to play many roles in artificial intelligence. For example, it is impossible for language models to detect that they are producing truthful statements or hallucinations if they have no way to represent the truthfulness of a statement. Instead of truth, the models are restricted to assessing the probability of a sequence of words. Uncommon facts in the training set cannot be distinguished from lies and common lies in the training set cannot be distinguished from truth, and so the models hallucinate.

## 2.2 Propositions as the Minimum Unit of Truth

Individual words, and even more so, individual tokens, do not have a truth value. The word "purple" is neither true nor false. The statement, "The sky is purple," may or may not be true.

Propositions are the shareable objects that are the bearers of truth and falsity. A proposition is the minimum unit that can have a truth value. A proposition consists of a minimum of three terms: An entity, a relation, and an entity.

(entity$_1$ → relation → entity$_2$)

There are many ways of writing these propositions, but they all express at least a relation between two entities. For example, the tuple ('Costco', 'located', 'United States') consists of two entities, "Costco" and "United States," and the relation "located."

Propositions are the first-class foundational units of PROPAI. While tokens suffice for predicting linguistic patterns, representing truth requires propositions. In less technical terms, a proposition is a "fact" if it is true, and a "factoid" if it is not or if its truth value is unknown.

Example propositions:

- From an oncology corpus:
  - ('Bufalin'→'inhibits'→'fibroblast growth factor receptor') **True**, interpreted as Bufalin inhibits fibroblast growth factor receptors."
- From an airline procedures corpus:
  - ('in-flight therapeutic oxygen'→ 'have' → 'no charge') **True**, interpreted as there is no charge for the use of in-flight therapeutic oxygen.
- From Costco SEC filings:
  - ('warehouses' → 'locate' → 'Puerto Rico') **True**, interpreted as Costco has warehouses located in Puerto Rico.
  - ('costco warehouses' → 'have' → 'bars') **False**, interpreted as Costco warehouses include bars.

## 2.3 The Dual Representation Insight

A key innovation in PROPAI is the concept of representing propositions, including facts and factoids, as a logical/semantic dual. Every proposition is represented as a logical aspect, consisting of symbolic entities and relations, and as a semantic aspect representing the semantic profile of the proposition as a vector. Both aspects of the dual can be used simultaneously and complementarily in reasoning. Two propositions can be symbolically similar because they are connected in a graph and semantically similar because they are represented by similar semantic profiles.

This dual representation bridges a historical tension in AI between symbolic systems (which excel at logical reasoning but struggle with ambiguity) and neural systems (which handle ambiguity well but lack explicit reasoning capabilities). Symbolic systems are too rigid and neural network systems are too flexible, but in combination they enable additional capabilities.

### 2.3.1 Logical Representation

The logical aspect of propositions can be connected into graphs, which allows basic reasoning by following links through the graph. Any graph analytic techniques that can be applied to knowledge graphs can be applied to PROPAI graphs, including multi-hop reasoning chains where intermediate facts serve as bridges between query concepts and answer concepts.

For example, in a medical oncology dataset, PROPAI constructs reasoning chains through graph traversal:

Query: Find path from topoisomerase 2 to t cells

Reasoning chain: topoisomerase 2 → target → treatment → t cells

Query: Find path from EGFR signaling to non-small cell lung cancer

Reasoning chain: EGFR signaling → endocytosis → cancer → non-small cell lung cancer

These reasoning chains represent hypotheses that can be evaluated against source evidence. The graph structure makes these reasoning paths explicit and auditable, unlike the black-box inference in transformer-based models.

### 2.3.2 Semantic Representation: Truth Profiles

The semantic vector aspects of propositions represent the meaning of the fact in a continuous space, enabling semantic similarity calculations that complement the discrete logical representation. These relations can be cast as knowledge maps, where similar propositions are depicted as close together in a semantic space.

## 2.4 Single-Pass Training and Continuous Learning

The PROPAI training process differs fundamentally from LLM training in both approach and efficiency, enabling capabilities that are architecturally impossible in transformer-based models. PROPAI models train with one-pass through the data without iteration. As a result, the training requires only a tiny fraction of the energy and time needed to train an LLM. Training time is measured in seconds, and the model can be updated online. Each fact can be modeled independently.

PROPAI's online learning capability is essential for enterprise applications where knowledge bases must remain current with regulatory changes, new research, or updated policies.

**Source Attribution and Auditability**

Every fact maintains a traceable link to its source document(s). Source tracking enables two critical capabilities: (1) auditability (users can verify that answers derive from approved sources), and (2) targeted update (when a source document changes, only the facts derived from that document need updating). The source tracking extends beyond simple citation; it maintains the provenance chain from natural language source through fact extraction to the dual representation.

## 2.5 Plausibility

PROPAI computes the plausibility of propositions. Plausibility is defined as consistency of a proposition with other facts in the knowledge base. Propositions may be true and have low plausibility. Their truth derives from their validated source. But plausibility is predictive of the truth of propositions derived from other sources.

Plausibility ranges from -1.0 to 1.0, but in practice it is rarely much less than 0.0. High scores (i.e., near 1.0) indicate high levels of consistency, lower scores indicate less or no consistency. Low plausibility, thus, underestimates the probability that a proposition is true, but overall, there is good calibration between predicted plausibility and observed probability of verifiable truth (Brier score of 0.027).

The capacity to estimate the plausibility of unknown propositions could also be used to allow the model to self-curate its own training data. Propositions could be extracted from sources of unknown validity. Those propositions that have high plausibility could be further evaluated for inclusion as fact in the model.

## 2.6 Hypotheses: PROPAI as a Generative Model

Beyond reasoning over extracted facts, PROPAI can generate propositions with plausibility estimates, even as their truth value is unknown. In certain settings, these generated plausible propositions can be called hypotheses. They are predicted to potentially be true, but it would take investigation of one or more authoritative sources to establish their status. These hypotheses, in other words, could be used to drive machine curiosity. This capability enables exploratory reasoning and can guide further fact extraction from additional sources.

**Concrete Examples from Medical Oncology**

From a PubMed oncology dataset, PROPAI generated the following hypotheses with plausibility scores:

(Tryptophan levels → can be rescued → asparagine): 0.426 plausibility

(Bioactive molecules → are linked to → gut-brain axis): 0.27 plausibility

For question answering, when asked "What treats CRC?" (colorectal cancer), PROPAI generated hypotheses:

(t cell infiltration → treat → crc): 0.165 plausibility

(nlrp3 → treat → crc): 0.165 plausibility

(tank-binding kinase 1 inhibitor → treat → crc): 0.606 plausibility

These plausibility scores can guide researchers toward promising therapeutic hypotheses while explicitly flagging that these are unverified factoids requiring validation through additional sources or experimental evidence.

# 3. Technical Implementation Details

## 3.1 System Architecture

The PROPAI system architecture consists of four primary components operating in a deterministic pipeline:

**Fact Extraction Pipeline**

An LLM or other processing method translates natural language from reliable sources into structured facts. The extraction process identifies entity-relation-entity triples from source documents. Currently, PROPAI leverages LLMs for this extraction step with prose documents, acknowledging that extraction quality affects downstream accuracy. Structured data requires the use of more direct methods. Future research directions include developing extraction methods that reduce dependency on probabilistic language models.

Factoids extracted from vetted, reliable sources are classified as facts. Factoids generated through reasoning chains or hypothesis generation are marked as unverified unless they can be traced to source documents.

**Dual Representation Encoding**

Extracted facts are simultaneously encoded in both semantic and logical aspects as described earlier. Unlike traditional knowledge graphs, PROPAI logical encodings do not depend on a preexisting ontology. The source data prescribe the entities and relations that make up these graphs.

Similarly, PROPAI representations are generated in one pass without iteration and preserve much of the semantic information in the proposition. PROPAI embeddings are designed to preserve the truth and semantic information in a proposition, whereas standard "graph embedding techniques have shown remarkable capacity of converting high-dimensional sparse graphs into low-dimensional, dense and continuous vector spaces …, where graph structure properties are maximally preserved" [7]. Both types of embedding are different from the token embeddings used in LLMs.

**Storage Layer**

Fact representations are stored with auditable source attribution. For most enterprise applications, the entire model fits in memory, though it can also be stored in a database. The storage requirements grow sub-linearly with data volume-the number of parameters increases more slowly than the volume of text processed.

For larger datasets requiring distributed storage, the architecture is distributable across graph and vector databases. Truth profiles can be combined at inference time when reasoning requires facts stored across multiple partitions.

**Reasoning Engine**

The type of reasoning engaged by the reasoning engine depends on the specifics of the use case. Because the logical fact representations can be organized into a graph, the entire repertoire of graph processing is available to the model. The availability of the vector, semantic aspect of the propositional representation means that processes that depend on continuous representations are also available.

Propositions can be clustered following both graph and continuous space representations and both can be combined to further enhance available capabilities. Topological data analysis techniques can also be employed, for example, to find the holes in the knowledge graph and thereby suggest topics for further investigation.

The reasoning engine synthesizes results from both search methods, using logical precision where required and semantic flexibility where needed.

Unlike with LLMs, the reasoning engine is deterministic. Given the same input and fact base, PROPAI produces identical output. There is no temperature parameter, no probabilistic sampling, no randomness in reasoning. This determinism is essential for regulated applications requiring auditability and reproducibility.

## 3.2 Question-Answering as an Example Reasoning Process

One application of PROPAI is as a RAG [7] replacement. The RAG market is currently estimated to be around $2 billion with a cumulative annual growth rate around 40%.

RAG is a combination of selected source documents, search, and other processes in an attempt to improve the focus and accuracy of the question answering process. Documents in a selected corpus are chunked into sub-documents. Each chunk is converted to a vector and stored in a vector database. When a question is presented, it is converted into a comparable vector, the chunks are then ranked for their similarity to the question vector and the top few chunks are then presented to an LLM as the source of information to answer the question.

In order for such a model to be successful, the chunks have to be chosen correctly, the vector has to capture the key information in each chunk, the question vector has to match the chunk(s) with the answer in them, and the LLM has to correctly extract the answer from the selected chunks. There is a lot that can go wrong, and the processes built around this process are only adding to the complication.

PROPAI greatly simplifies this process, makes it deterministic, and eliminates the possibility of a hallucination. The fact propositions are extracted from each of the documents in the selected corpus. A PROPAI model is constructed from these facts. When a question is presented, the question is converted to a list of factoids, which may have missing information. These factoids are then used to select the potentially relevant fact proposition from the model. The question and these selected facts are then used as a prompt to an LLM, which generates the answer.

Based on a corpus drawn from a sample of medical abstracts on Alzheimer's and Parkinson's diseases, a PROPAI model and a RAG model were both constructed. The same questions were presented to both models, which gave different results. For example, one question was: "What does tau protein do?" Despite there being 102 facts in this corpus that mention tau protein, the RAG model failed to answer the question, claiming that the information was not available. The PROPAI model, on the other hand, provided a detailed description of the role of this protein.

The PROPAI approach eliminates the ambiguity inherent in RAG's document chunking, and other processes. Rather than hoping that relevant information falls within a single chunk boundary, PROPAI reasons directly over atomic facts that can be composed flexibly.

# 4. Performance Characteristics and Empirical Validation

## 4.1 Computational Efficiency

PROPAI's efficiency relative to LLMs lowers the energy requirements, and operational and infrastructure costs dramatically. Because it is a deterministic model not based on transformer architectures, it is not subject to the constraints and foibles of these machines. Because it is focused on specific customer data and uses an efficient architecture, it requires orders of magnitude fewer resources than LLMs do. Because it is incapable of hallucination, it does not require extensive guardrails to cover up misalignments.

**Training Time**

LLMs require iterative training over massive datasets, consuming weeks to months on thousands of GPUs. GPT-4 reportedly required approximately $100 million in computational costs over a three-month training period. In contrast, once the facts have been collected, PROPAI model training completes in a single pass, typically finishing in seconds. PROPAI does not require iterative optimization through gradient descent.

**Memory Footprint and Scalability**

Modern LLMs require trillions of parameters. GPT-4's parameter count exceeds one trillion. Storing and serving these models requires substantial computational infrastructure. The amount of storage needed approaches capacity that was available in the entire world only a few years ago.

PROPAI's dual representation grows sub-linearly with data volume. Even if applied to the full scope of information addressed by large LLMs, PROPAI would require only a microscopic fraction of this computational capacity and memory. The number of parameters for PROPAI grows sub-linearly with the volume of text, while language model parameters have been growing exponentially.

In the worst-case scenario where massive scale is required, the same software engineering techniques that support LLMs (distributed computation, model parallelism) could be applied to PROPAI. However, such scale is rarely necessary for enterprise deployments focused on specific knowledge domains.

**Inference Speed**

PROPAI inference operates through direct graph traversal and vector similarity search, both highly optimized operations. There is no autoregressive token generation bottleneck. Query response times scale with the complexity of the reasoning chain required, not with output length. For simple factoid retrieval, responses are near-instantaneous.

**Update Efficiency**

When source documents change, PROPAI's source tracking enables targeted updates. Only facts derived from modified documents need re-extraction and re-encoding. This contrasts sharply with LLMs, which require complete retraining or expensive fine-tuning procedures to incorporate new information. The high cost of frequent LLM updates makes them unsuitable for rapidly evolving knowledge domains.

## 4.2 Accuracy, Auditability, and Reliability Guarantees

The core architectural difference between PROPAI and LLMs manifests most clearly in accuracy guarantees, auditability, and reliability:

**Elimination of Hallucination**

PROPAI is architecturally incapable of hallucinating. It can only retrieve and reason over facts that were explicitly extracted from source documents. When information is not in the fact base, the system explicitly identifies this absence rather than fabricating plausible-sounding content. This is not achieved through guardrails or confidence thresholds, it is a fundamental property of the architecture.

The elimination of guardrail requirements stems from restricting training data to reliable, enterprise-focused sources. PROPAI models cannot be prompted to provide harmful information if they do not contain that information to begin with. This architectural safety contrasts with LLM guardrails which attempt to block access to information that the model has already learned.

**Complete Source Traceability**

Every fact maintains provenance to its source document with complete traceability:

- Source document identification
- Extraction timestamp
- Original natural language text from which the fact was extracted

Answers include attribution, enabling verification and creating complete audit trails. This auditability is essential for regulatory compliance in industries like healthcare (HIPAA), finance (Basel III, Dodd-Frank), and legal services.

**Deterministic Reasoning**

Given the same query and fact base, PROPAI produces identical results with no stochastic variation. This reproducibility is essential for:

- Regulatory compliance audits
- Quality assurance testing
- Legal discovery processes
- Scientific reproducibility

LLMs, by contrast, introduce stochastic sampling at inference time, making exact reproducibility impossible even with fixed random seeds.

**Explicit Uncertainty Acknowledgment**

When reasoning chains produce hypotheses whose truth value is unknown, PROPAI explicitly flags this uncertainty through plausibility scores. Hypotheses are never presented as verified facts. This contrasts sharply with LLMs, which present hallucinations with the same linguistic confidence as accurate information.

These characteristics make PROPAI suitable for applications where LLMs face insurmountable trust barriers: regulatory compliance, medical decision support, financial risk assessment, legal research, and any domain where being wrong has consequences.

# 5. Comparative Analysis: PROPAI vs. LLMs vs. RAG

This section provides a comprehensive comparison across multiple dimensions to clarify where each approach excels and where fundamental limitations exist.

A summary of this analysis is shown in Table 1.

*Table 1. Summary comparison of PROPAI, LLMs alone, and RAG*

| Dimension | PROPAI | LLMs | RAG |
|---|---|---|---|
| Unit of Analysis | Facts (verifiable) | Tokens (no truth value) | Document chunks + LLM |
| Hallucination | Architecturally impossible | Inevitable | Reduced but persists |
| Training Time | Minutes to hours (single pass) | Weeks to months | Hours (chunking + indexing) |
| Continuous Learning | Yes (online) | No (batch only) | Yes (re-index required) |
| Determinism | Fully deterministic | Stochastic | Stochastic (LLM) |
| Source Attribution | Complete (fact-level) | None | Partial (chunk-level) |
| Reasoning Chains | Explicit (auditable) | Implicit (black box) | Implicit (LLM-dependent) |
| Scaling Pattern | Sub-linear | Exponential | Linear (chunks) |

# 6. Applications in Regulated Industries

PROPAI's architecture makes it particularly suited for enterprise applications where accuracy, auditability, and trustworthiness are paramount. PROPAI is not an application, rather it is a foundation model on which applications are built.

Current proof-of-value pilots demonstrate practical validation in multiple sectors:

**Financial Services**

Financial institutions face stringent regulatory requirements (Basel III, Dodd-Frank, MiFID II) that mandate explainable, auditable decision-making. PROPAI applications include:

- Regulatory compliance verification with complete audit trails
- Credit risk assessment with explainable reasoning chains
- Broker-dealer supervision and trade surveillance
- Claims processing with source attribution
- ESG risk detection and monitoring

**Healthcare and Life Sciences**

Healthcare applications demand zero-hallucination guarantees due to patient safety implications and HIPAA compliance requirements. Current pilots include pharmaceutical and medical research applications:

- Clinical decision support with HIPAA-compliant reasoning
- Drug discovery and repurposing through knowledge graph reasoning
- Medical literature synthesis from PubMed with source traceability
- Adverse event detection and pharmacovigilance
- Treatment efficacy prediction with plausibility scoring

**Legal Services**

The legal domain represents arguably PROPAI's clearest use case: precision and source attribution are professional requirements, not nice-to-have features:

- Legal research with case law citations

- Contract analysis and clause extraction
- Due diligence with complete audit trails
- Regulatory compliance verification
- Patent analysis and prior art search

**Enterprise Knowledge Management**

Enterprise deployments leverage PROPAI's ability to maintain private, customer-specific knowledge bases with continuous updates. Current pilots span pharmaceutical, waste management, and retail sectors:

- Standard operating procedure Q and A
- AI governance
- Know Your Customer (KYC) verification
- Compliance monitoring against internal policies
- Regulatory requirement tracking
- Employee training and onboarding support

In each domain, the cost of hallucination is unacceptable. These applications require systems that are correct.

# 7. Discussion: Limitations and Future Directions

## 7.1 Current Limitations

While PROPAI eliminates hallucination and provides deterministic reasoning, it has inherent limitations that define its appropriate use cases:

- **Bounded Knowledge:** PROPAI can only reason over facts extracted from its training corpus. Unlike LLMs that can approximate answers based on broad training, PROPAI explicitly acknowledges when information is unavailable. This is a feature for enterprise deployments but limits general-purpose usage.
- **Extraction Dependency:** The quality of fact extraction from source documents directly affects system accuracy. Current implementations leverage LLMs for extraction of facts from unstructured data sources. Whereas hallucinations are theoretically possible during such extraction, they are very unlikely because of the structured output that is generated. The probability of a hallucination increases with the increasing length of the output as each generated token is added to the context for the next token to be generated. The facts being extracted are very short and approximately independent of one another, so there is practically no chance for such errors to accumulate. Extraction errors are contained, they do not compound through inference. Structured data sources require other tools to extract facts.
- **Natural Language Generation:** PROPAI excels at reasoning, but it is not designed to produce fluent sentences. LLMs can be used to translate propositions to fluent language, when required.
- **Creative Tasks:** PROPAI is not designed for open-ended creative generation, brainstorming, or stylistic writing. These remain appropriate use cases for LLMs.

## 7.2 Future Research Directions

Several promising research directions could extend PROPAI's capabilities:

- **Extraction Independence:** Developing fact extraction methods that reduce dependency on LLMs would eliminate the primary source of potential error propagation. Approaches might include specialized extraction models trained specifically for factoid identification or rule-based parsers for structured documents.
- **Temporal Reasoning:** Extending the dual representation to incorporate temporal validity would enable reasoning about time-dependent facts and historical changes.
- **Web-Scale Knowledge:** Scaling to web-scale knowledge bases while maintaining sub-linear growth characteristics would require distributed architectures and advanced indexing strategies. But the

ability to represent truth and curate its own data would make such models much more powerful than those existing today. PROPAI is not sufficient for artificial general intelligence, but its capabilities are essential to achieving it [8].

- **Hybrid Architectures:** Exploring hybrid systems that combine PROPAI reasoning with LLM generation could optimize for both correctness and fluency, using each architecture for its strengths.
- **Multimodal Extension:** Extending fact extraction to images, videos, and structured data (tables, graphs) would broaden the scope of knowledge PROPAI can incorporate.
- **Collaborative Reasoning:** Enabling multiple PROPAI systems to share and reason over distributed knowledge bases while maintaining privacy and access controls.

### 7.3 Implications for AI Development and Deployment

The introduction of PROPAI as a distinct category has several implications for the broader AI landscape:

**Tool Selection Framework:** Organizations need frameworks for choosing between LLMs and PROPAI based on use-case requirements. The key question shifts from "which AI is better?" to "which architecture is appropriate for this task?" Criteria include:

- Consequences of error (high-stakes vs. low-stakes)
- Regulatory requirements (audit trails, explainability)
- Knowledge scope (bounded domain vs. open-ended)
- Update frequency (static vs. rapidly evolving)

**Regulatory Evolution:** As regulators establish risk-based frameworks (EU AI Act, FDA guidance for medical AI), PROPAI's deterministic reasoning and auditability may become required for high-risk applications. The ability to provide complete reasoning chains and source attribution aligns naturally with regulatory transparency demands.

**Enterprise Economics:** PROPAI's efficiency characteristics create fundamentally different cost structures than LLMs. Single-pass training, compact models, and lower computational requirements enable enterprise AI deployment at scales previously cost-prohibitive. The combination of high gross margins (90%+) and low infrastructure costs creates a different economic model than LLM-based services.

**Architectural Diversity:** The success of PROPAI would validate the principle that different AI architectures are optimal for different tasks. This counters the current trend toward treating LLMs as general-purpose solutions and encourages development of specialized architectures optimized for specific capabilities.

## 8. Conclusion

We have introduced Propositional Reasoning Artificial Intelligence (PROPAI) as a new category of AI systems designed to address the fundamental limitation of LLMs: their inability to represent truth. Through a patent-pending dual-representation architecture that simultaneously represents facts as logical units for graph-based reasoning and semantic vectors for flexible query handling, PROPAI eliminates hallucination by design rather than attempting to reduce it probabilistically.

The technical contributions of this work include:

- A dual-representation architecture that combines the reasoning power of knowledge graphs with the semantic flexibility of neural networks
- A plausibility scoring mechanism that enables hypothesis generation and evaluation on a -1.0 to 1.0 scale
- Single-pass training with online learning capabilities for continuous knowledge base updates

- Complete source attribution and deterministic reasoning for regulatory compliance
- Demonstrated superiority over RAG approaches through elimination of chunking ambiguity and LLM-based answer extraction

PROPAI is not positioned as a replacement for all uses of LLMs but as a complementary technology optimized for specific use cases. This tool-choice boundary clarifies the AI landscape: use LLMs when fluency and creativity matter most; use PROPAI when correctness and trustworthiness are non-negotiable.

The architectural advantages of PROPAI—sub-linear scaling, deterministic inference, complete auditability, and architectural impossibility of hallucination—address critical barriers to enterprise AI adoption in regulated industries. Proof-of-value pilots in pharmaceutical, financial services, and enterprise knowledge management demonstrate practical validation of these capabilities.

As AI systems become increasingly embedded in high-stakes decision-making, the distinction between probabilistic language modeling and fact-based reasoning becomes not just technical but essential. The current paradigm—optimizing language models for scale while attempting to constrain hallucinations through guardrails—has reached its architectural limits. PROPAI establishes an alternative foundation where truth is not a post-hoc constraint but an architectural primitive.

The future of AI is not a single architecture dominating all applications, it is an ecosystem of specialized systems optimized for different capabilities. PROPAI demonstrates that when correctness matters more than creativity, when auditability is required rather than optional, and when hallucination is disqualifying rather than inconvenient, a different architectural approach is not just beneficial but necessary.

Intelligence, it turns out, really does start with facts.

# About the Authors

**Francesca Lazzeri**, Ph.D., is Principal Group Director of Data and Applied AI Science at Microsoft, where she leads a multidisciplinary organization of AI researchers, data applied scientists, and machine learning engineers. Prior to joining Microsoft, Francesca was a Research Fellow at Harvard University (Technology and Operations Management Unit) and taught Python for AI as an Adjunct Professor at Columbia University. She has authored several books, including: Machine Learning Governance for Managers (Springer Nature, 2023). She serves as Chief Editor of Data Science at Microsoft, a publication dedicated to practical insights for developers and data scientists.

**Herbert L. Roitblat**, Ph.D., is the Co-founder, Chief Technology Officer and Chief AI Officer at Reliath AI. He served in similar roles at previous technology companies and was Professor of Psychology, Second Language Acquisition, and Marine Biology at the University of Hawaii. He has authored several books including Algorithms are not Enough: Creating General Artificial Intelligence (MIT Press, 2020), and numerous papers on artificial intelligence. He is the inventor of Reliath AI technology.

**Chris LaCour** is the Co-founder and Chief Executive Officer of Reliath AI. He began his academic journey in Computer Science before pivoting to Sociology, where he discovered he was more energized by working with people than writing code - and uncovered a natural talent for sales. He launched his career at Ricoh Business Systems, selling document imaging solutions, before joining Allen Systems Group, where he sold mainframe and open systems software to Fortune 2000 IT executives. He later held a brief role in IT staffing and subsequently moved into eDiscovery with LawScribe and Hudson Legal. Following Hudson Legal, Chris became an entrepreneur, founding multiple companies, including one that curated exclusive, executive-level events bringing together senior business leaders. He successfully exited the business with its sale in 2023. Inspired by an article on his co-founder’s Substack, Chris co-founded Reliath AI with Herb Roitblat, introducing Propositional-Reasoning Artificial Intelligence (PROPAI) - a fundamentally new class of AI designed to reason over verifiable facts rather than simply predict text.

## References


[1] OpenAI. (2023). GPT-4 Technical Report. arXiv:2303.08774.

[2] OpenAI. (2024). Why Language Models Hallucinate. https://openai.com/index/why-language-models-hallucinate/

[3] Vaswani, A., et al. (2017). Attention Is All You Need. In Advances in Neural Information Processing Systems 30 (NIPS 2017).


For more information:


www.reliath.ai

info@reliath.ai

(949) 346-3846



[4] Weiser, B. (2023). Here's What Happens When Your Lawyer Uses ChatGPT. The New York Times.

[5] Charlotin, D. (2026). AI Hallucination Cases. https://www.damiencharlotin.com/hallucinations/ Retrieved 20-Feb-2026.

[6] Lewis, P., et al. (2020). Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. In Proceedings of NeurIPS 2020.

[7] Xu, M. (2021). Understanding graph embedding methods and their applications. SIAM Review, 63(4), 825-853.

[8] Roitblat, H. L. (2020). Algorithms Are Not Enough: Creating General Artificial Intelligence. MIT Press.