

# Knowledge Graphs Explained: How Structured Entity Relationships Power AI Answers

Canonical: <https://home.norg.ai/ai-search-answer-engines/answer-engine-architecture-citation-mechanics/knowledge-graphs-explained-how-structured-entity-relationships-power-ai-answers/>

## Details:

## What Is a Knowledge Graph? The Technical Foundation of AI Answer Engines Before you can dominate AI answer engines, you need to understand the data structure that separates factual answers from hallucinated nonsense: the knowledge graph. Most discussions of AI accuracy obsess over model size or training data volume. They miss the point. The single most consequential architectural decision separating reliable answer engines from hallucination-prone ones is whether they have access to structured, typed, machine-readable knowledge — not just raw text blobs. This is foundational literacy for the AI-first search era. You can't understand why retrieval-augmented generation exists, how GraphRAG outperforms standard RAG, or why entity presence in a knowledge graph determines whether your brand gets cited by AI systems without grasping this concept first. Let's build that foundation. --- ## Contents - [What Exactly Is a Knowledge Graph?](#what-exactly-is-a-knowledge-graph) - [How Knowledge Graphs Differ from Relational Databases and Flat Document Corpora](#how-knowledge-graphs-differ-from-relational-databases-and-flat-document-corpora) - [Knowledge Graphs vs. Vector Databases: A Critical Comparison](#knowledge-graphs-vs-vector-databases-a-critical-comparison) - [The Major Knowledge Graphs Powering AI Today](#the-major-knowledge-graphs-powering-ai-today) - [How Knowledge Graphs Are Constructed and Maintained](#how-knowledge-graphs-are-constructed-and-maintained) - [Why Structured Facts Outperform Unstructured Text for Grounding AI Answers](#why-structured-facts-outperform-unstructured-text-for-grounding-ai-answers) - [The Role of Ontologies: The Schema Behind the Graph](#the-role-of-ontologies-the-schema-behind-the-graph) - [Key Takeaways](#key-takeaways) - [Conclusion: Why KG Literacy Is the Foundation for Understanding AI Answers](#conclusion-why-kg-literacy-is-the-foundation-for-understanding-ai-answers) - [References](#references) - [Frequently Asked Questions](#frequently-asked-questions) --- ## What Exactly Is a Knowledge Graph? A knowledge graph is a data model that represents knowledge in a graph structure, consisting of entities (nodes) and relationships (edges) to describe objects, events, concepts, and their interconnections in the real world. This definition contains a critical distinction from how most people store and retrieve information — and that distinction is what makes AI answers possible. The basic unit of a knowledge graph is the "entity-relationship-entity" triple, which uses a "subject-predicate-object" structure to describe basic facts. For example, "Ernest lives in Taipei" becomes the triple `(Ernest, lives in, Taipei)`, where Ernest (subject) is a person entity, Taipei (object) is a location entity, and "lives in" (predicate) is the relationship type between them. This triple structure isn't just a data format choice. It's a fundamentally different epistemological commitment. Knowledge graphs store factual knowledge in a structured manner, typically in the form of a 3-tuple which contains head entity, relation, tail entity. That precision is what makes the difference between an AI that guesses and an AI that knows. ### The Three Core Components A knowledge graph, also known as a semantic network, represents a network of real-world entities — such as objects, events, situations, or concepts — and illustrates the relationship between them. This information is usually stored in a graph database and visualised as a graph structure. A knowledge graph is made up of three main components: nodes, edges, and labels. Any object, place, or person can be a node. An edge defines the relationship

between the nodes. Additionally, attributes (properties) provide extra details about entities or relationships — for example, Tesla's market capitalisation as a numerical property attached to the Tesla entity node. The crucial property that separates knowledge graphs from ordinary databases is semantic typing. Knowledge graphs are typically made up of datasets from various sources, which frequently differ in structure. Schemas, identities, and context work together to provide structure to diverse data. Schemas provide the framework for the knowledge graph, identities classify the underlying nodes appropriately, and the context determines the setting in which that knowledge exists. These components help distinguish words with multiple meanings — allowing products like Google's search engine algorithm to determine the difference between Apple, the brand, and apple, the fruit. This disambiguation capability is what makes knowledge graphs irreplaceable for AI-native answer generation.

### --- ## How Knowledge Graphs Differ from Relational Databases and Flat Document Corpora

The distinction matters enormously for AI systems. A relational database stores data in rows and columns with a predefined schema. A flat document corpus — like the raw web text used to train LLMs — is an unstructured blob of language. A knowledge graph occupies a different category entirely. Compared to relational databases that store data in tables (rows and columns) with predefined schemas, knowledge graphs organise entities and their relationships in a graph structure. Traditional databases excel at storing structured data and handling basic queries, but they're limited in capturing complex associations and reasoning new knowledge from data. Knowledge graphs represent data as networks of nodes and edges, allowing flexible addition of new node types or relationship types without changing the overall structure, providing high schema extensibility. Each piece of knowledge is presented as a triple, enabling the system to perform bidirectional queries and reasoning through relationship edges. For example, knowing "sky has colour blue" allows reverse inference that "things with blue colour include sky." This bidirectional reasoning capability is what makes knowledge graphs irreplaceable for AI answer generation. An LLM querying a flat document corpus must infer relationships probabilistically from co-occurrence patterns. A knowledge graph makes those relationships explicit, typed, and traversable. No inference required. No probabilistic guessing. Just facts.

### --- ## Knowledge Graphs vs. Vector Databases: A Critical Comparison

In modern AI architectures, knowledge graphs and vector databases are often discussed as competing retrieval mechanisms. The distinction isn't about which is "better" — it's about what type of question each is designed to answer. Knowledge graphs represent entities and relationships explicitly, which supports structured facts, multi-step questions, and explainable answers about how things connect. Vector databases store embeddings, enabling fast similarity search where close matches matter more than exact ones, especially across unstructured data like text and images.

Dimension	Knowledge Graph	Vector Database
Data model	Typed nodes and edges (triples)	High-dimensional numerical vectors
Query mechanism	Graph traversal, SPARQL, Cypher	Cosine similarity, approximate nearest neighbour
Best for	Structured facts, multi-hop reasoning, entity disambiguation	Semantic similarity, unstructured text, fuzzy matching
Explainability	High — retrieved subgraphs show reasoning path	Low — similarity scores lack transparent justification
Weakness	Struggles with unstructured or free-form text	Cannot represent explicit entity relationships
Typical use case	Factual grounding, compliance, fraud detection	Document search, recommendation, RAG over text corpora

Knowledge graphs support precise, structured queries through graph languages (SPARQL, Cypher), provide explainability since retrieved subgraphs clearly show why data was selected, and are fit for domains with curated, structured knowledge such as biomedicine, compliance, and supply chain. Vector-based search struggles with ambiguous context, lacks explicit reasoning, and doesn't maintain structured knowledge over time — affecting reliability, especially in fields like healthcare, finance, and legal AI, where accuracy and transparency are critical. The practical implication for answer engines: when a user asks a factual question like "Who founded Tesla, and what other companies has that person led?", a vector database returns documents that \*mention\* the relevant terms. A knowledge graph traverses the entity graph and returns verified, typed relationships — `` , `` — with no ambiguity about what each relationship means. (See our guide on [\*GraphRAG vs. Standard RAG: When Knowledge Graphs Outperform Vector Search for Complex Questions\*](#graphrag-vs-standard-rag) for a deeper comparison of these retrieval architectures in practice.)

### --- ## The Major Knowledge Graphs Powering AI Today

#### ### Google Knowledge Graph: The World's Largest Proprietary KG

Google

announced its Knowledge Graph in May 2012, presenting it as "an intelligent model, a graph in geek jargon, that encompasses real-world entities and their relationships to each other" — encompassing "things, not strings." Its growth since then has been extraordinary. As of May 2024, Google had more than 1.6 trillion facts about 54 billion entities in its Knowledge Graph — up from 500 billion facts on 5 billion entities in 2020. And it's still growing. The Knowledge Graph contains information that Google considers to be fact, including entities such as people, companies, films, and topics; attributes for those entities such as date of birth, location, and founders; and relationships between entities such as who worked for which company, who played in which music group, and which person is expert in which topics. Critically, the 2023–2024 "Killer Whale" updates revealed a strategic shift in how Google constructs this graph. Between May 2020 and June 2023, the number of Person entities in Google's Knowledge Vault increased steadily. In July 2023, the number of Person entities tripled in just four days. In March 2024, Google added an additional 17%. In less than four years, between May 2020 and March 2024, the number of Person entities in Google's Knowledge Vault increased over 22-fold. It appears Google is looking for person entities to which it can fully apply E-E-A-T credibility signals, aiming to understand who is creating content and whether they are trustworthy. This has direct implications for content creators and brands seeking citation visibility in Google AI Overviews (see our guide on [\*How Google AI Overviews Work\*](#how-google-ai-overviews-work)). Google's algorithms can now create entities in the Knowledge Vault without trusted sources like Wikipedia, if the information about the entity is clear, complete, and consistent across the web. However, almost one in five entities created in the Knowledge Vault is deleted within a year — which means KG presence isn't a one-time achievement but an ongoing maintenance requirement. If you're not in the graph, you don't exist to AI answer engines. Period. ### Wikidata: The Open-Access Global Knowledge Graph Wikidata is a collaboratively edited multilingual knowledge graph hosted by the Wikimedia Foundation, released as open data under the Creative Commons CC0 public domain dedication. As of early 2025, Wikidata had 1.65 billion item statements (semantic triples). As of August 2025, Wikidata has been described as the world's largest open-access knowledge graph. Statements are how any information known about an item is recorded in Wikidata. Formally, they consist of key-value pairs, which match a property (such as "author" or "publication date") with one or more entity values. For example, the informal English statement "milk is white" would be encoded by a statement pairing the property `colour (P462)` with the value `white (Q23444)` under the item `milk (Q8495)`. All knowledge in Wikidata is queryable through a SPARQL query interface ([query.wikidata.org/](https://query.wikidata.org/)), which also enables distributed queries across other Linked Data resources. Wikidata's significance for AI systems extends well beyond its role as a reference database. The Wikidata Embedding Project, made available in October 2025, provides a vector-based semantic search tool allowing plain-language queries and supports the Model Context Protocol standard that makes the data more readily available to AI systems. The project is a partnership between Wikimedia Deutschland, Jina.AI, and DataStax, an IBM subsidiary. This is the future: open, queryable, AI-native knowledge infrastructure. ### Industry-Specific Knowledge Graphs General-purpose knowledge graphs like Google's and Wikidata are complemented by domain-specific KGs that encode specialised ontologies. Wikidata itself has been assembled from repositories in the fields of genomics, proteomics, genetic variants, pathways, chemical compounds, and diseases, adhering to the FAIR principles of findability, accessibility, interoperability, and reusability. In biomedical research, graph databases store biomedical entities such as genes, drugs, and diseases and their relationships, enabling multi-hop reasoning and dynamic updates, while vector databases enable semantic similarity searches to match natural language queries with relevant graph data. This hybrid architecture — structured KG for entities, vectors for semantic lookup — is the current production standard in high-stakes domains. When accuracy matters — healthcare, finance, legal — knowledge graphs aren't optional. They're foundational. --- ## How Knowledge Graphs Are Constructed and Maintained ### Traditional Construction: The Knowledge Acquisition Bottleneck The traditional KG construction pipeline required relationship extraction — analysing syntax, semantics, and context to determine relationships between entities — followed by graph construction mapping extracted (Entity, Relation, Entity) tuples into a graph database. This approach had a critical bottleneck: it required large, accurately labelled training datasets. Creating these labelled corpora in specialised domains like healthcare or legal was prohibitively expensive. This "knowledge acquisition bottleneck" limited

knowledge graph construction to well-funded organisations with extensive data science teams. That bottleneck just got obliterated. ### LLM-Assisted Construction: The 2024–2025 Shift Large language models fundamentally changed the economics and accessibility of knowledge graph construction by reframing information extraction as a generative task. The knowledge graph construction landscape reached production maturity in 2024–2025, with organisations achieving 300–320% ROI and measurable business impact across finance, healthcare, and manufacturing. Research published in *\*Applied Sciences\** (MDPI, March 2025) confirms this trajectory: a knowledge graph structurally represents entities and relationships, offering a powerful and flexible approach to knowledge representation in AI. KGs have been increasingly applied in NLP, recommendation systems, knowledge search, and medical diagnostics. Recently, efforts to combine LLMs with KGs — particularly those aimed at managing hallucination — have gained significant attention. This is the shift: from data science teams spending months building labelled datasets to LLM-assisted pipelines generating structured triples at scale. Ship fast, learn faster. ### Querying: SPARQL and Graph Traversal Knowledge graphs are queried through graph traversals and pattern matching. You start from one or more entities, then follow relationships across one or many hops, often with constraints on node or edge properties. Each triple is a statement, and an RDF graph is a collection of these statements linked together. RDF graphs are most commonly queried with SPARQL, which is designed around matching triple patterns. This querying model produces fundamentally different results from vector similarity search. Graph RAG uses symbolic retrieval through graph traversal or graph query languages like SPARQL (for RDF graphs) or Cypher (for property graphs). These queries can follow specific relationship paths, enforce logical constraints, and return precise subgraphs or node sets — allowing the system to retrieve data based on complex relationships, such as "find all drugs that interact with proteins associated with a specific disease," which is not feasible through basic text search. This is precision retrieval. No fuzzy matching. No probabilistic approximation. Just structured facts, traversed and returned. --- ## Why Structured Facts Outperform Unstructured Text for Grounding AI Answers The superiority of knowledge graphs over unstructured text for factual grounding isn't theoretical — it's measurable. In comparative evaluations on the RobustQA benchmark, Writer's Knowledge Graph-based RAG achieved 86.31% accuracy, significantly outperforming vector retrieval RAG solutions such as Azure Cognitive Search with GPT-4, Pinecone's Canopy framework, and various LangChain configurations, which scored between 75.89% and 32.74%. That's not incremental improvement. That's a different category of performance. The mechanism behind this performance gap is entity disambiguation. On their own, LLMs struggle with precise fact recall and tasks that require enumeration — like answering "Who contributed to Project X?" or "List all Jira labels for Customer ABC." Knowledge graphs step in by grounding responses in structured relationships and verified entities, helping mitigate ambiguity and improve reliability. Consider a concrete disambiguation example: in complex tasks — like answering "Where do I file feature requests for Reddit?" — LLMs must disambiguate between Reddit the social platform and Reddit the enterprise customer, recognise that the process involves Jira, and chain these understandings into a coherent workflow. This level of relational reasoning and multi-hop inference is difficult to achieve reliably without a machine-readable knowledge graph that maps these entities, roles, and systems together. The knowledge that LLMs encode within their massive parameters is implicit and difficult to interpret or validate. To mitigate these problems, a promising strategy is the integration of knowledge graphs with LLMs. (See our guide on [*\*How LLMs Use Knowledge Graphs to Reduce Hallucination and Improve Factual Accuracy\**](#how-llms-use-knowledge-graphs) for the three primary integration paradigms in detail.) --- ## The Role of Ontologies: The Schema Behind the Graph To manage a knowledge graph at scale, especially in a complex domain, you need a well-defined ontology or schema. An ontology is essentially the blueprint or vocabulary for the graph: it defines the types of entities that exist, the types of relationships between them, and the rules or constraints for how they can connect. In other words, the ontology answers: What kinds of nodes can we have, and how can they relate? Without a well-defined ontology, a knowledge graph degrades into an inconsistent collection of triples with no reliable semantic meaning. This is why domain-specific KGs — in biomedicine, legal research, or financial compliance — invest heavily in ontology design before any data ingestion begins. The ontology is what makes the graph's facts *\*unambiguous\**, which is precisely the property that makes KG-grounded

answers more reliable than answers derived from probabilistic text retrieval. Ontologies are the schema layer that transforms data into knowledge. Without them, you have triples. With them, you have truth. --- ## Key Takeaways - Knowledge graphs are networks of typed entity-relationship triples, not databases of documents or high-dimensional vectors. Their basic unit — `` — makes every fact explicit, directional, and machine-queryable. - Google's Knowledge Graph now contains over 1.6 trillion facts on 54 billion entities (as of May 2024, per Kalicube/Search Engine Land analysis), making it the world's largest proprietary structured knowledge base and a primary factual grounding source for Google AI Overviews. - Wikidata, the largest open-access knowledge graph, reached 1.65 billion semantic triples as of early 2025, and now supports the Model Context Protocol, making its structured data directly accessible to AI systems. - Knowledge graphs outperform vector databases for multi-hop reasoning and entity disambiguation, but vector databases outperform knowledge graphs for semantic similarity over unstructured text — which is why production AI systems increasingly use hybrid architectures combining both. - LLM-assisted KG construction reached production maturity in 2024–2025, eliminating the "knowledge acquisition bottleneck" that previously restricted graph building to large, specialised organisations. --- ## Conclusion: Why KG Literacy Is the Foundation for Understanding AI Answers Knowledge graphs aren't a peripheral feature of modern AI answer engines — they are the structural backbone that separates factual grounding from probabilistic guessing. When Google AI Overviews cite a specific entity, when Perplexity disambiguates a brand name from a common noun, or when a medical AI correctly chains a drug to its contraindicated conditions, a knowledge graph is doing the work that unstructured text retrieval cannot reliably perform. Understanding the triple structure, the distinction between KGs and vector databases, and the scale and construction methods of the major graphs (Google, Wikidata, domain-specific KGs) gives you the conceptual vocabulary to understand every advanced topic in this series: why GraphRAG outperforms standard RAG for complex questions (see our guide on [\*GraphRAG vs. Standard RAG\*](#graphrag-vs-standard-rag)), how entity presence determines citation eligibility (see [\*Entity Authority and Knowledge Graph Presence\*](#entity-authority)), and why the hallucination problem is fundamentally a knowledge-grounding problem (see [\*The Hallucination Problem: Why Answer Engines Fabricate Citations\*](#hallucination-problem)). The shift from "strings to things" — Google's own framing from 2012 — wasn't a search feature launch. It was the beginning of a new information architecture that now determines which facts AI systems trust, which entities they recognise, and ultimately, which sources they cite. If you want visibility everywhere — in AI Overviews, in Perplexity, in ChatGPT search results — you need to understand this architecture. Because in the AI-first search era, the question isn't whether your content is good. The question is whether your entities exist in the graphs that power answer engines. Become the answer. Start with the graph. --- ## References - Google Support. "How Google's Knowledge Graph Works." \*Google Knowledge Panel Help\*, 2024. <https://support.google.com/knowledgepanel/answer/9787176> - Search Engine Land. "What Is the Knowledge Graph? How It Affects SEO and Visibility." \*Search Engine Land\*, November 2025. <https://searchengineland.com/guide/knowledge-graph> - Search Engine Land / Kalicube. "Unpacking Google's 2024 E-E-A-T Knowledge Graph Update." \*Search Engine Land\*, May 2024. <https://searchengineland.com/unpacking-google-2024-eeat-knowledge-graph-update-440224> - Wikipedia / Wikimedia Foundation. "Wikidata." \*Wikipedia\*, 2025. <https://en.wikipedia.org/wiki/Wikidata> - Turki, H. et al. "Wikidata as a Knowledge Graph for the Life Sciences." \*eLife\*, 2020. <https://elifesciences.org/articles/52614> (PMC: <https://pmc.ncbi.nlm.nih.gov/articles/PMC7077981/>) - Frontiers in Computer Science. "Practices, Opportunities and Challenges in the Fusion of Knowledge Graphs and Large Language Models." \*Frontiers in Computer Science\*, June 2025. <https://www.frontiersin.org/journals/computer-science/articles/10.3389/fcomp.2025.1590632/full> - MDPI Applied Sciences. "Knowledge Graph Construction: Extraction, Learning, and Evaluation." \*Applied Sciences\*, March 2025. <https://www.mdpi.com/2076-3417/15/7/3727> - IBM. "What Is a Knowledge Graph?" \*IBM Think\*, 2024. <https://www.ibm.com/think/topics/knowledge-graph> - Wikipedia. "Knowledge Graph (Google)." \*Wikipedia\*, 2025. [https://en.wikipedia.org/wiki/Knowledge\\_Graph\\_\(Google\)](https://en.wikipedia.org/wiki/Knowledge_Graph_(Google)) - Instaclus. "Graph RAG vs Vector RAG: 3 Differences, Pros and Cons, and How to Choose." \*Instaclus\*, November 2025. <https://www.instaclus.com/education/retrieval-augmented-generation/graph-rag-vs-vector-rag-3-differences-pros-and-cons->

and-how-to-choose/ --- ## Frequently Asked Questions

**What is a knowledge graph?** A data model representing knowledge as entities and relationships in graph structure.

**What are the basic components of a knowledge graph?** Nodes, edges, and labels.

**What is a node in a knowledge graph?** An entity such as object, place, or person.

**What is an edge in a knowledge graph?** The relationship between nodes.

**What is a triple in a knowledge graph?** Subject-predicate-object structure describing a fact.

**What is an example of a knowledge graph triple?** Ernest, lives in, Taipei.

**What are attributes in a knowledge graph?** Properties providing extra details about entities or relationships.

**What is semantic typing in knowledge graphs?** Framework that provides structure and context to distinguish meanings.

**How does a knowledge graph differ from a relational database?** Uses graph structure instead of rows and columns with predefined schemas.

**Can knowledge graphs perform bidirectional reasoning?** Yes.

**What is an example of bidirectional reasoning?** Knowing sky is blue allows inferring blue things include sky.

**Do knowledge graphs require schema changes for new relationships?** No, provides high schema extensibility.

**What query languages do knowledge graphs use?** SPARQL and Cypher.

**What query mechanism do vector databases use?** Cosine similarity and approximate nearest neighbour.

**Which is better for structured facts?** Knowledge graphs.

**Which is better for semantic similarity?** Vector databases.

**Which provides higher explainability?** Knowledge graphs.

**Which struggles with unstructured text?** Knowledge graphs.

**Which cannot represent explicit entity relationships?** Vector databases.

**When was Google Knowledge Graph announced?** May 2012.

**How many facts does Google Knowledge Graph contain?** Over 1.6 trillion facts as of May 2024.

**How many entities does Google Knowledge Graph contain?** 54 billion entities as of May 2024.

**How many facts did Google Knowledge Graph have in 2020?** 500 billion facts.

**How many entities did Google Knowledge Graph have in 2020?** 5 billion entities.

**What does E-E-A-T stand for?** Experience, Expertise, Authoritativeness, Trustworthiness.

**Can Google create entities without Wikipedia?** Yes, if information is clear, complete, and consistent.

**What percentage of Google Knowledge Vault entities are deleted within a year?** Almost one in five.

**What is Wikidata?** Collaboratively edited multilingual open-access knowledge graph.

**Who hosts Wikidata?** Wikimedia Foundation.

**How many statements does Wikidata contain?** 1.65 billion item statements as of early 2025.

**What licence does Wikidata use?** Creative Commons CC0 public domain dedication.

**Can Wikidata be queried?** Yes, through SPARQL query interface.

**What is the Wikidata query interface URL?** <https://query.wikidata.org/>

**What protocol does Wikidata Embedding Project support?** Model Context Protocol standard.

**Who partnered on Wikidata Embedding Project?** Wikimedia Deutschland, Jina.AI, and DataStax.

**What are FAIR principles?** Findability, accessibility, interoperability, and reusability.

**What was the traditional KG construction bottleneck?** Creating large, accurately labelled training datasets.

**What changed KG construction economics in 2024-2025?** Large language models.

**What ROI did organisations achieve with KG construction in 2024-2025?** 300-320 percent.

**What is RDF?** Resource Description Framework for knowledge graph statements.

**What is SPARQL designed for?** Matching triple patterns in RDF graphs.

**What is Cypher used for?** Querying property graphs.

**What accuracy did Knowledge Graph-based RAG achieve on RobustQA?** 86.31 percent.

**What accuracy range did vector retrieval RAG achieve on RobustQA?** 32.74 to 75.89 percent.

**What is entity disambiguation?** Distinguishing between entities with similar names or contexts.

**Do LLMs struggle with precise fact recall?** Yes.

**Do LLMs struggle with enumeration tasks?** Yes.

**Is LLM knowledge implicit?** Yes.

**Is LLM knowledge easy to validate?** No.

**What is an ontology in knowledge graphs?** Blueprint defining entity types, relationships, and connection rules.

**What happens without a well-defined ontology?** Graph degrades into inconsistent collection of triples.

**What was Google's 2012 framing for Knowledge Graph?** From strings to things.

**When did the Killer Whale update occur?** Between July 2023 and March 2024.

**How much did Person entities increase in July 2023?** Tripled in four days.

**How much did Person entities increase in March 2024?** Additional 17 percent.

**How much did Person entities increase between May 2020 and March 2024?** Over 22-fold.

**What determines AI citation eligibility?** Entity presence in knowledge graphs.

**Can knowledge graphs reduce hallucination?** Yes, by grounding responses in structured facts.

**What is GraphRAG?** RAG using graph traversal instead of vector similarity.

**What is the basic unit of a knowledge graph?** Entity-relationship-entity triple.

**Are knowledge graphs optional for high-stakes domains?** No, they are foundational.

**What**

makes knowledge graph answers more reliable than text retrieval? \*\* Explicit, typed, and traversable relationships. \*\*What is the primary weakness of vector databases? \*\* Cannot represent explicit entity relationships. \*\*What is the primary weakness of knowledge graphs? \*\* Struggles with unstructured or free-form text.

## Source Data (JSON):

```
{\n  \"_type\": \"article\", \n  \"title\": \"Knowledge Graphs Explained: How Structured Entity Relationships
```