




Why Knowledge Graphs are Necessary for Enterprise GenAI

Philip Rathle, Chief Technology Officer, Neo4j



The Evolution of... Web Search



How is GenAI like Web Search?

Web Search

1. Ask question
2. Technology operates on the q
3. Returns some text
4. Text = knowledge

GenAI

1. Ask question
2. Technology operates on the q
3. Returns some text
4. Text = knowledge

Evolution of Web Search

Full text Era: 1994 - 2000

Full Text
Search

YAHOO!

Netscape

Ask Jeeves[™]
Ask.com

LYCOS[®]

AOL[®]

Google Launches World's Largest Search Engine



Google Now Enables Internet Users to Search More Than 1 Billion URLs, Providing Quick and Easy Access to 560 Million Full-Text Indexed Web Pages and 500 Million Partially Indexed URLs

MOUNTAIN VIEW, Calif. – June 26, 2000 – Google Inc., one of the fastest growing search engines on the web, today announced it has released the largest search engine on the Internet. Google's new index, comprising more than 1 billion URLs, offers users the web's most comprehensive collection of websites, which can be easily searched with Google's fast and highly relevant search technology. Available now at www.google.com, Google's portal and destination site customers can also license this new index for integration with their own websites.

“Google is based on a variety of innovative technologies, including sophisticated text matching and its advanced, patent-pending technology called PageRank™, which ensures that the most important results always come up first.”

Evolution of Web Search


PageRank Era: 2000 - 2012

Full Text
Search

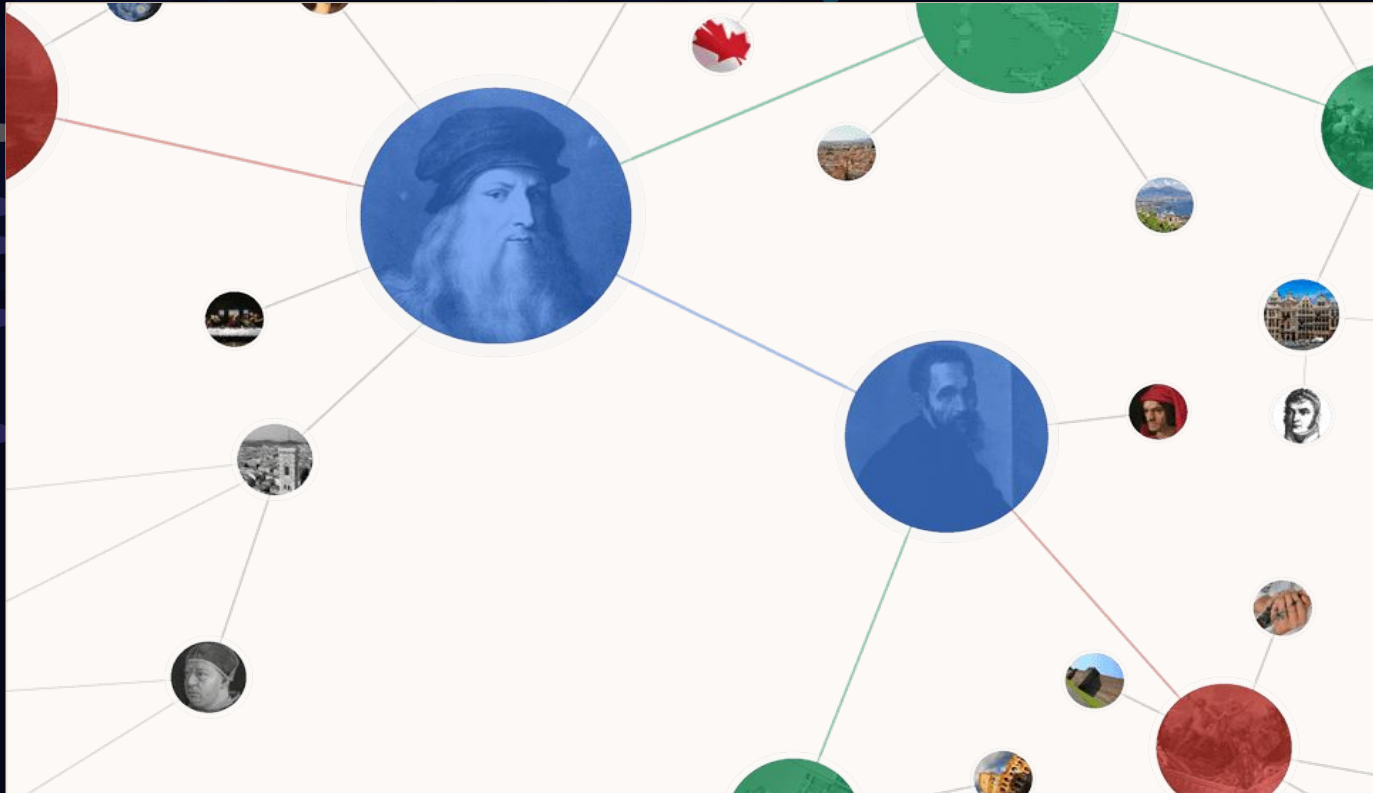


PageRank

Google



Google Knowledge Graph



SEARCH

Introducing the Knowledge Graph: things, not strings

May 16, 2012 · 4 min read

 Share



Amit Singhal
SVP, Engineering

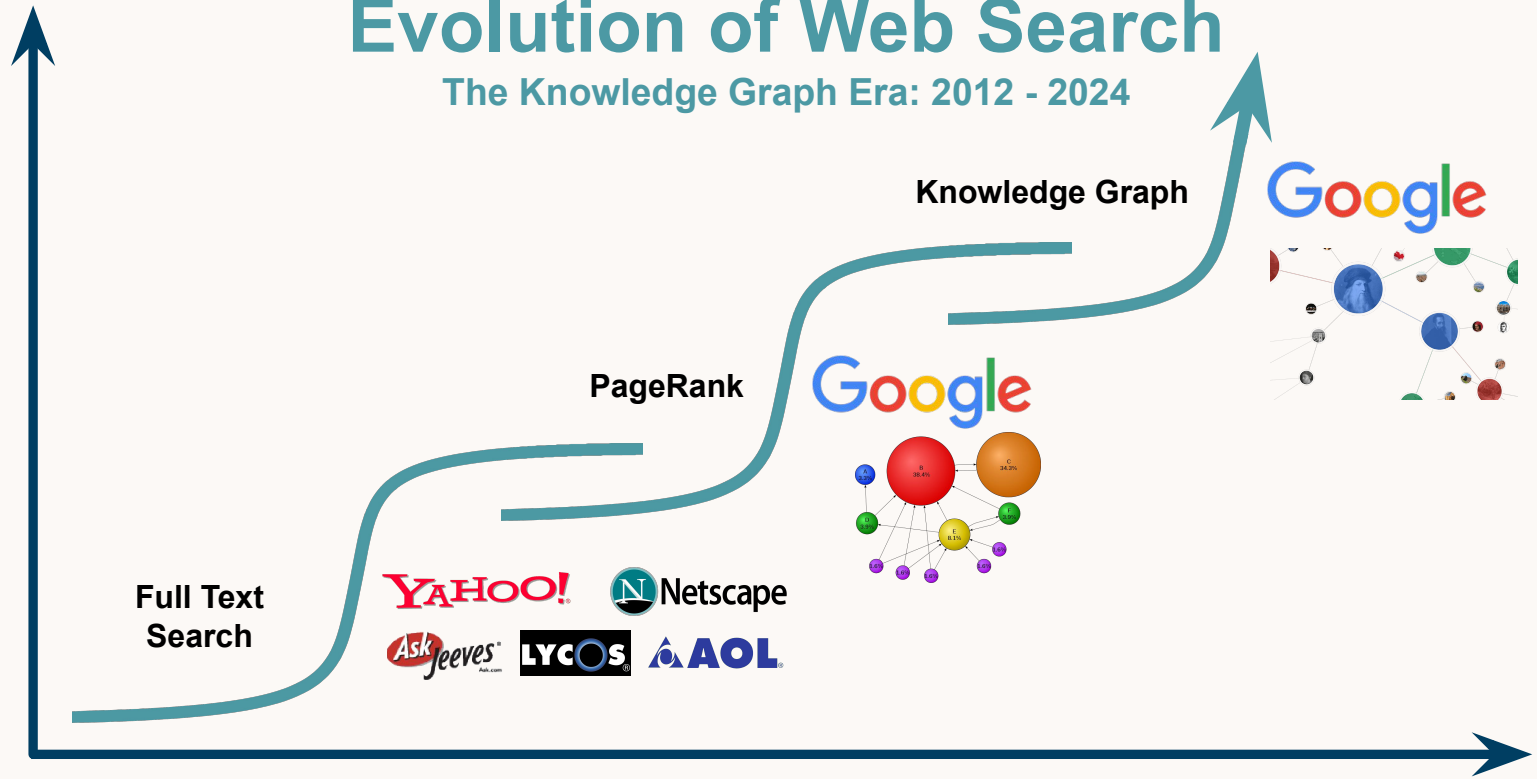
Search is a lot about discovery—the basic human need to learn and broaden your horizons. But searching still requires a lot of hard work by you, the user. So today I'm really excited to launch the Knowledge Graph, which will help you discover new information quickly and easily.

Take a query like [taj mahal]. For more than four decades, search has essentially been about



Evolution of Web Search

The Knowledge Graph Era: 2012 - 2024



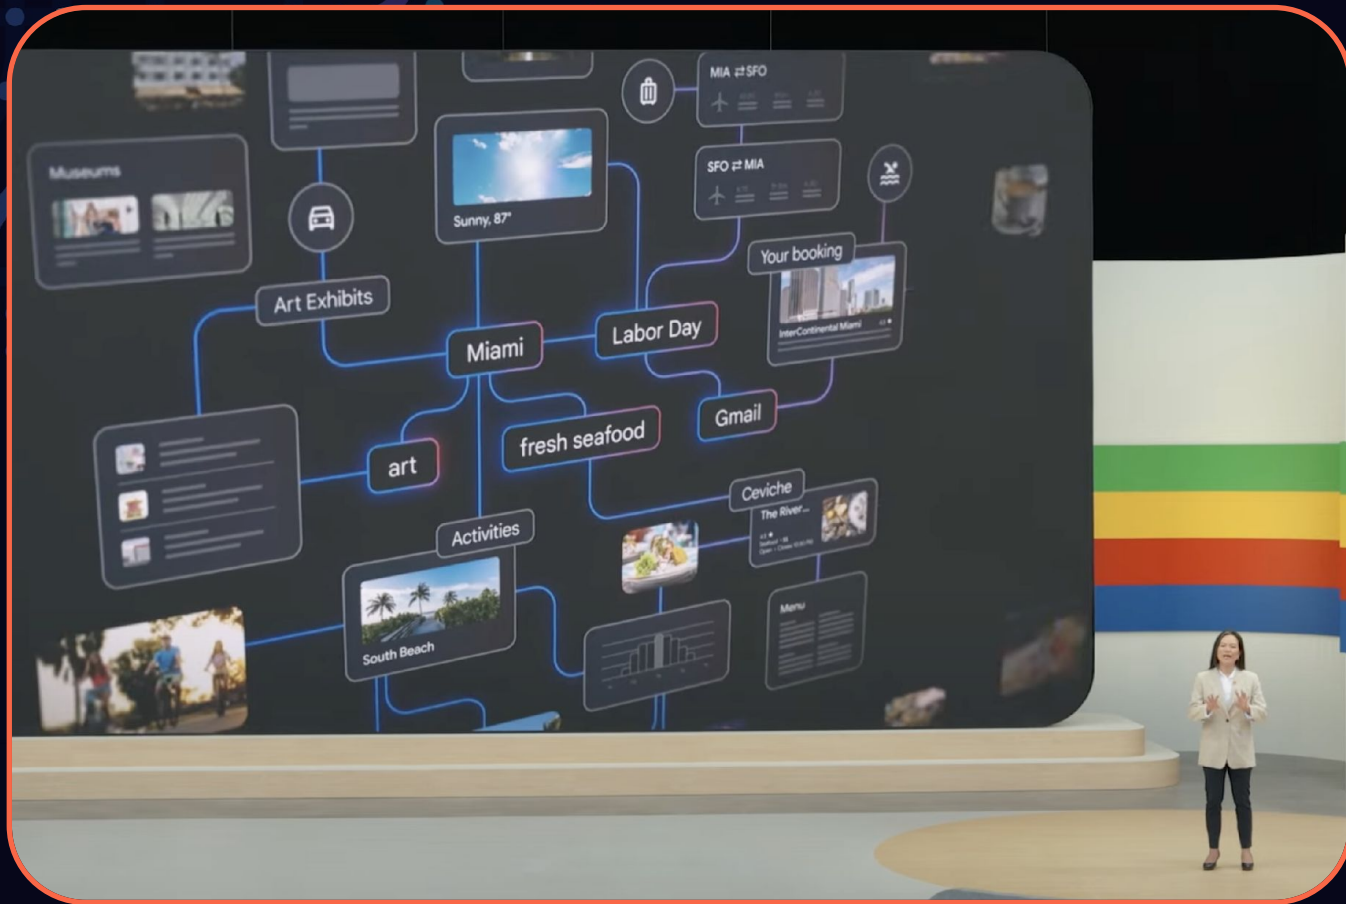
Google I/O '24 Keynote

My family and I are going to Miami for Labor Day. My son loves art and my husband really wants fresh seafood. Can you pull my flight and hotel info from Gmail and help me plan the weekend?



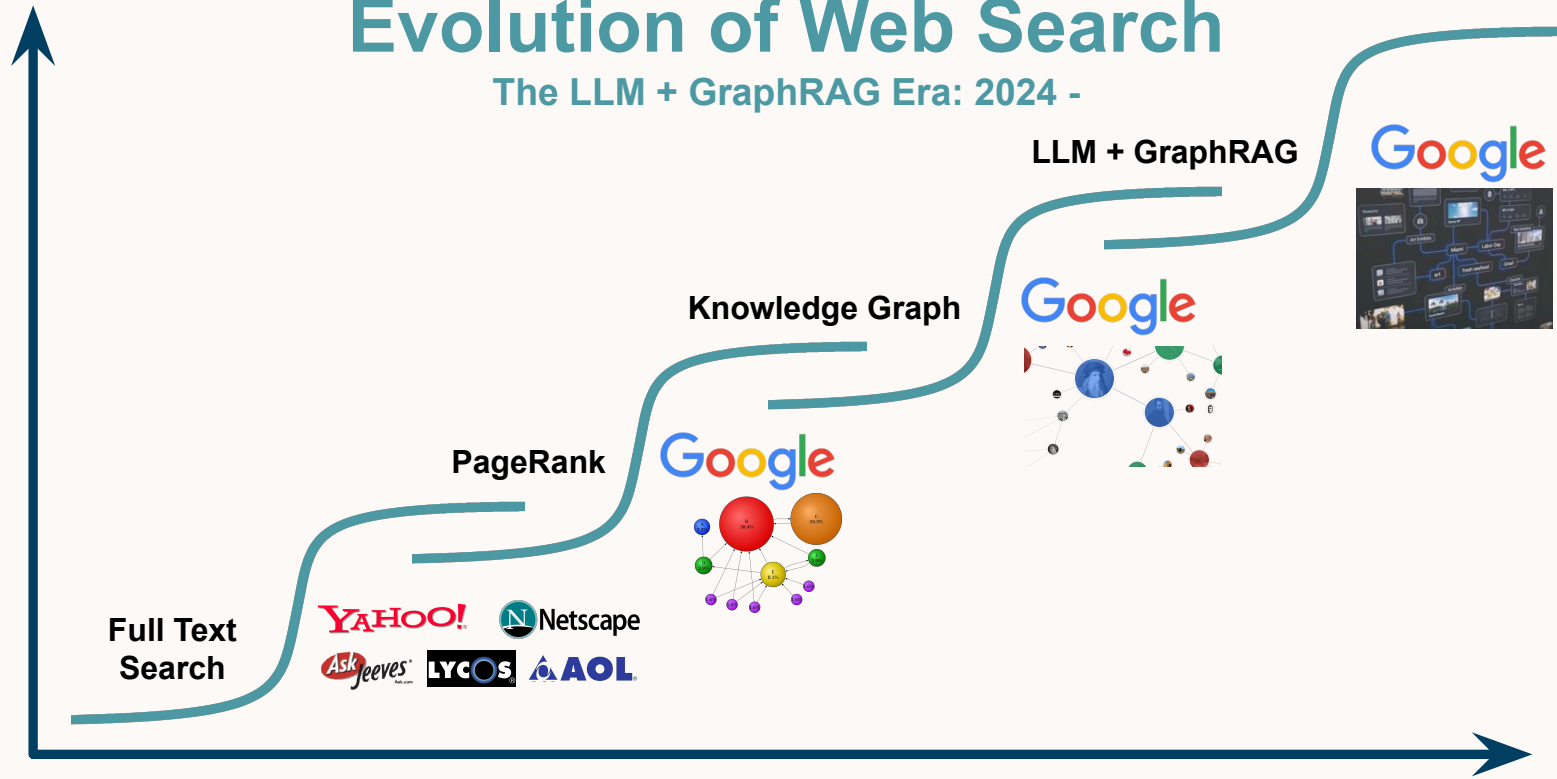
GraphRAG @Google

Google I/O '24
Keynote



Evolution of Web Search

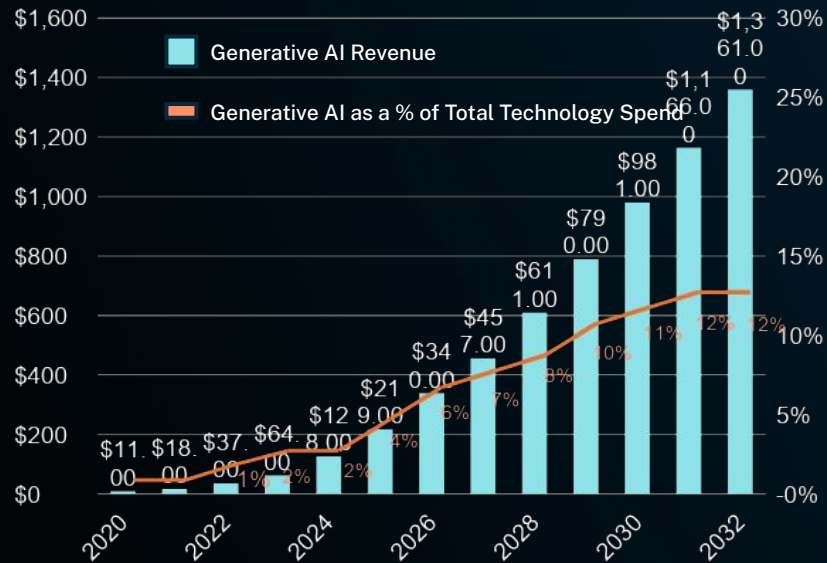
The LLM + GraphRAG Era: 2024 -





Back to GenAI...

Generative AI races toward \$1.3 trillion in revenue by 2032



Bloomberg Intelligence

However, **71%** of organizations are stuck piloting GenAI projects

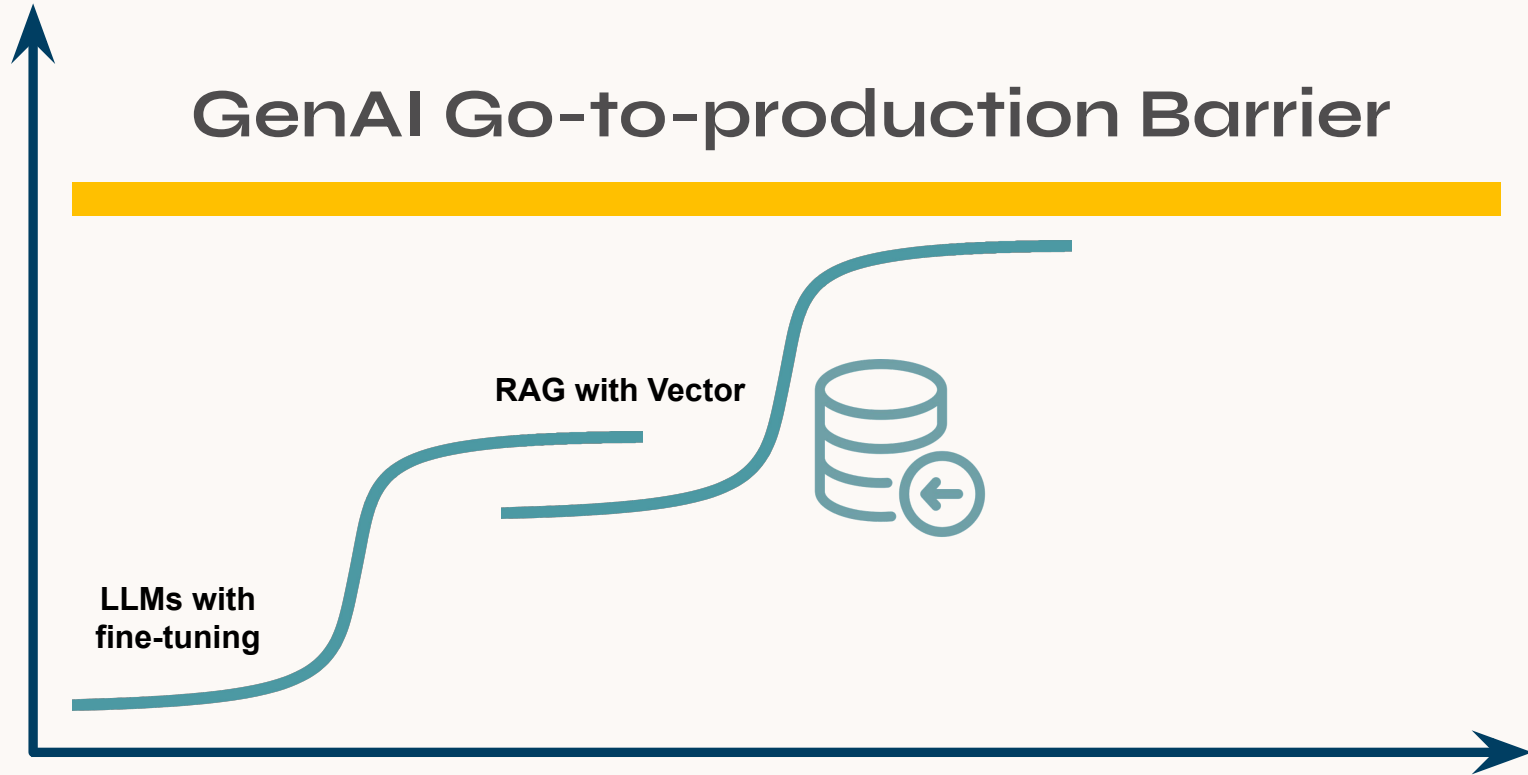
2024 IBM CEO Survey

GenAI Go-to-production Barrier



LLMs

GenAI Go-to-production Barrier



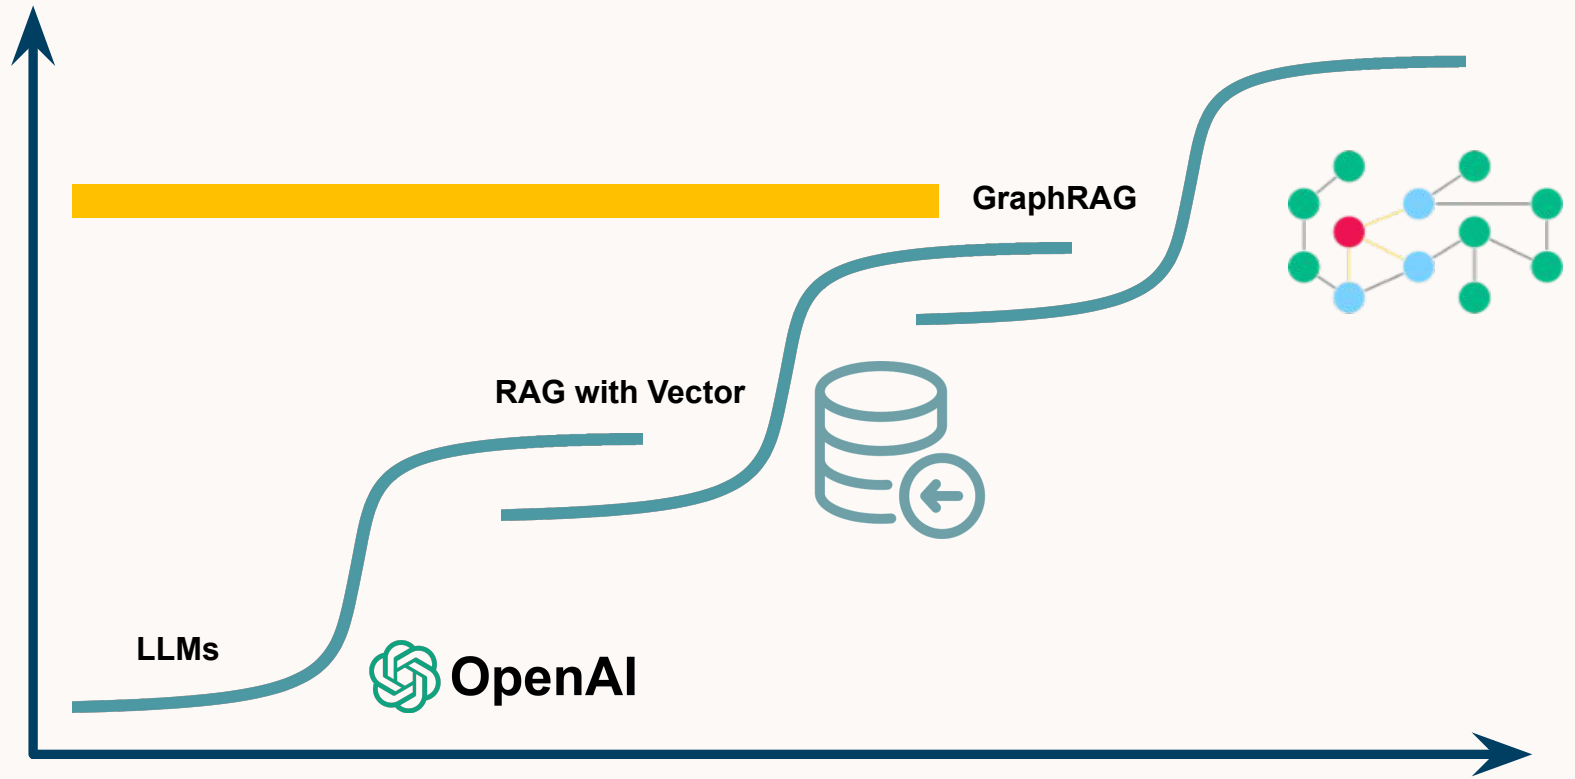
LLMs with
fine-tuning

RAG with Vector



Enter GraphRAG

*GraphRAG is **RAG** where the **Retrieval path** includes a **Knowledge Graph**.*



LLMs




OpenAI

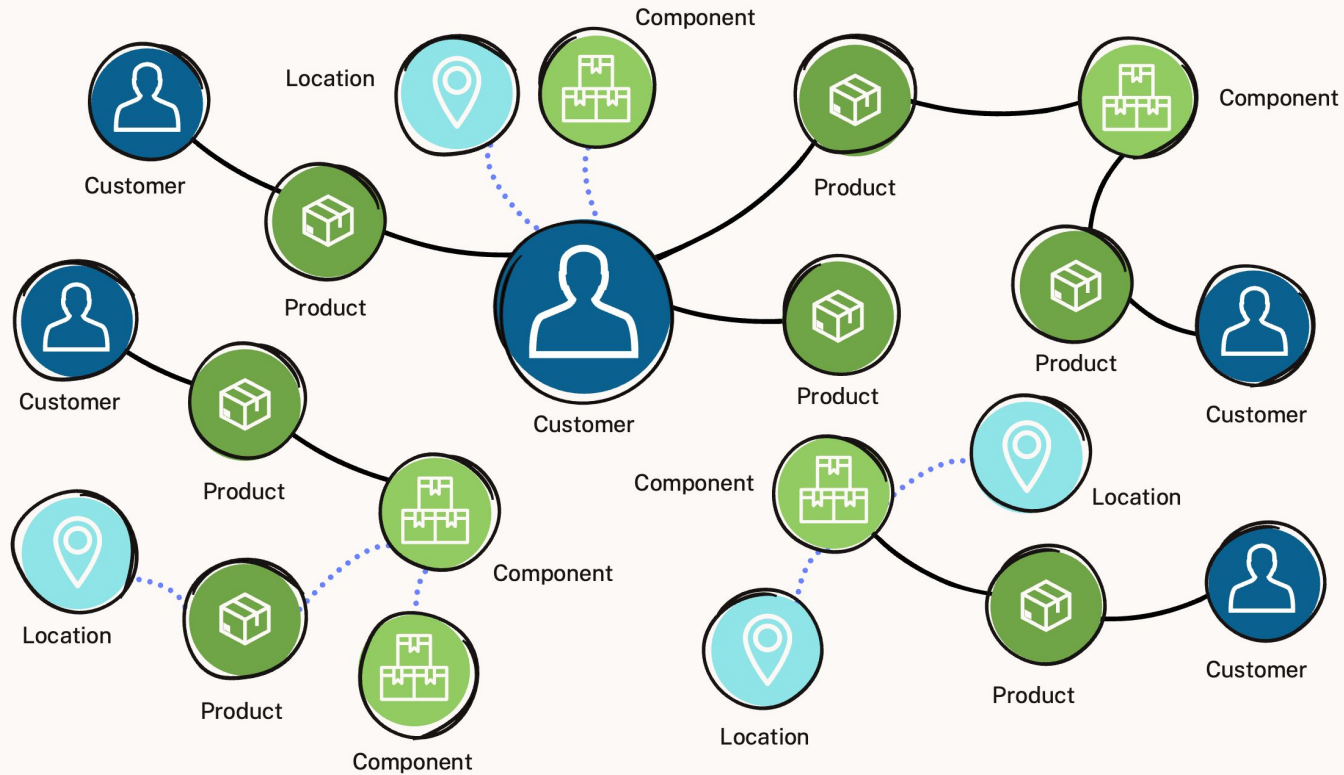
RAG with Vector



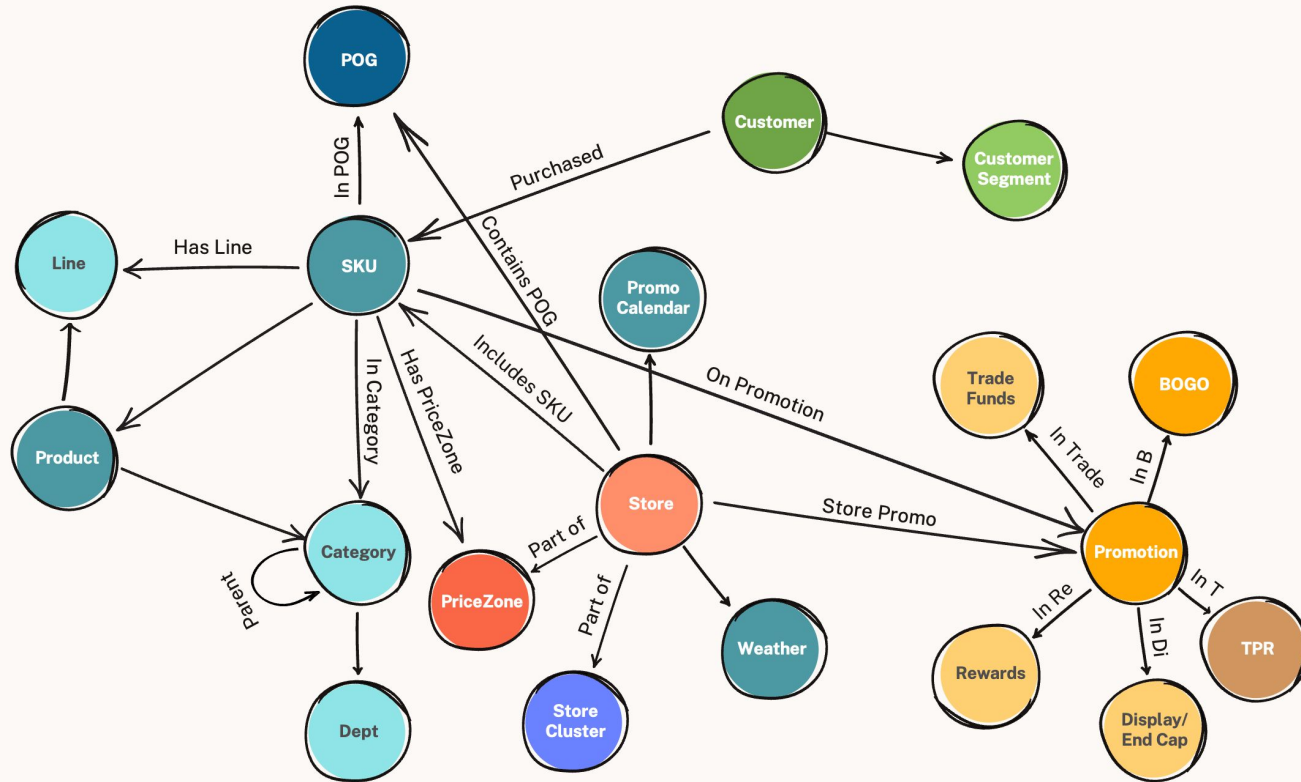
GraphRAG



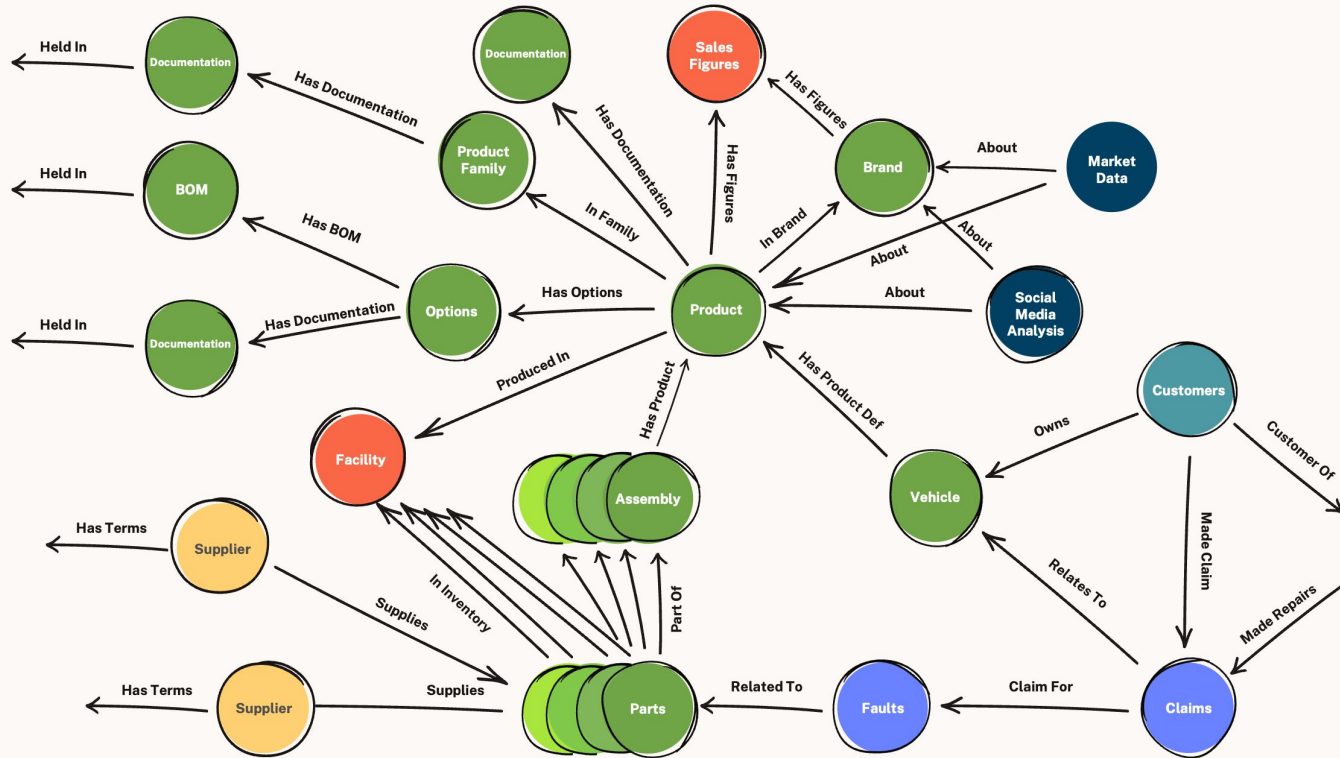
Commerce Transactions Graph



Product and Parts Graph



Digital Twin Graph





1

**Better
Answers**

2


**Easier
Development**

3

**Explainability
& Governance**

1. Better Answers

Vector-only RAG



**Better accuracy though
matching of concepts**


Which asset managers are vulnerable to lithium shortage?

Based on the context provided, here is a summary of the asset managers that may be vulnerable to lithium storage:

- The company that produces batteries and other products incorporating battery technology is vulnerable, as increased demand for batteries could place capacity constraints on their lithium supply chain.
- The company that sources steel and copper from Chile and Australia is vulnerable, as those countries are expected to experience severe weather due to climate change.
- The company that relies on suppliers for components and systems necessary for their products has some vulnerability. If there are lithium shortages it could result in supply constraints, cost pressures, or component shortages.
- The company uses lithium-ion batteries or other battery technology in their products is vulnerable to lithium shortages which could reduce battery supply, increase costs, or otherwise negatively impact their business.

1. Better Answers

GraphRAG Vector + Graph



**More actionable and
concrete answers**

Which asset managers are vulnerable to lithium shortage?

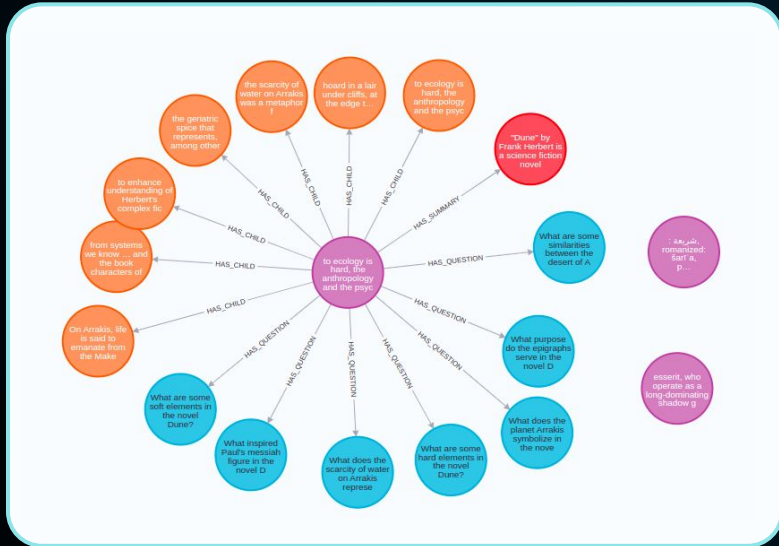
Based on the context provided, the asset managers that may be vulnerable to lithium shortage are:

- Quantamental Technologies LLC
- DEPRINCE RACE & ZOLLO INC
- ATRIA Wealth Solutions, Inc.
- CITIGROUP INC
- CHARLES SCHWAB INVESTMENT MANAGEMENT INC
- ROYAL LONDON ASSET MANAGEMENT LTD
- STIFEL FINANCIAL CORP
- TUCKER ASSET MANAGEMENT, LLC
- Patriot Financial Group Insurance Agency, LLC
- SeaCrest Wealth Management, LLC

The context indicates that Stanley Black & Decker uses lithium batteries in many of its products. With increasing demand for lithium, there could be supply constraints and increased costs for lithium. This could negatively impact the profitability of Stanley Black & Decker. The asset managers listed above have holdings in Stanley Black & Decker, and therefore may be vulnerable to the risks of lithium shortage

2. Easier Development

Transparent & Explainable

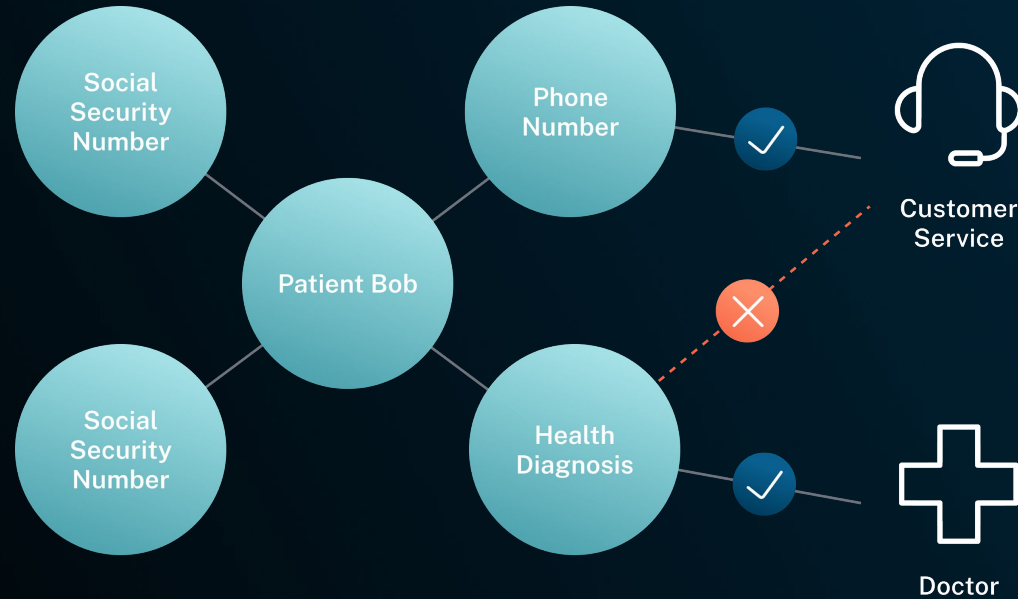


Opaque & Implicit

[3.9150659e-03	2.6659777e-03	1.0298982e-03	-2.7156321e-03
1.9977870e-03	3.1204436e-03	1.2055682e-04	1.0450699e-03
-6.4308796e-04	3.0822519e-03	2.1972554e-03	5.1480172e-05
-3.7099270e-03	3.9439583e-03	6.8276987e-04	7.7137066e-04
2.3698520e-03	-7.8547641e-04	6.0383842e-04	4.6370425e-03
-1.6786088e-03	1.7417425e-03	2.4216413e-03	3.6545738e-03
-1.9871239e-03	2.9489421e-03	-1.2810023e-03	-4.9174053e-04
-3.9743204e-03	-2.7023794e-03	-3.0541950e-04	-1.5724347e-03
-2.1029566e-03	-2.1624754e-03	2.1620055e-04	-1.4000515e-03
-4.0824865e-03	4.6588355e-04	3.5028579e-03	4.8283348e-03
-2.8737928e-03	-4.5569306e-03	-7.6568732e-04	-3.3311991e-03
3.5790715e-03	4.2424244e-03	3.3478225e-03	-7.4140396e-04
1.0030111e-03	-5.2394503e-04	5.8383477e-04	-4.8430995e-03
2.6972082e-03	-4.8002079e-03	-2.3011414e-03	8.0388715e-04
3.1952575e-05	-8.1621204e-04	-3.8127291e-03	-6.7428290e-04
-1.7713077e-03	-3.0159748e-03	1.7178850e-03	-1.9258332e-03
-2.4637436e-03	3.3779652e-03	2.7676420e-03	1.8853768e-03
-2.4718521e-03	-1.9754141e-03	2.6104036e-03	-2.1335895e-03
2.4405334e-03	-3.2013952e-04	3.9961869e-03	4.0419102e-03
2.0586823e-03	4.9897884e-03	4.5599132e-03	-1.0976522e-03
1.5563263e-03	3.9063310e-03	-2.9308300e-03	-4.8254002e-03
-8.7642738e-06	3.9748671e-03	5.2895391e-04	6.3330121e-04
-1.2614765e-03	-8.5018738e-04	3.7659388e-03	3.0237564e-03

AI Engineer "I actually already fixed a couple of bugs thanks to this!"

3. Explainability & Governance






When are Knowledge Graphs *Necessary* for GenAI?

Higher Stakes = Higher Bar



Low

Business peripheral
No single right answer
Leeway for creativity



Medium

Involves core business
Mitigations in place
(e.g. human in the loop)



High

Wrong answers result
in lost \$\$/ reputation/
fines/ etc.

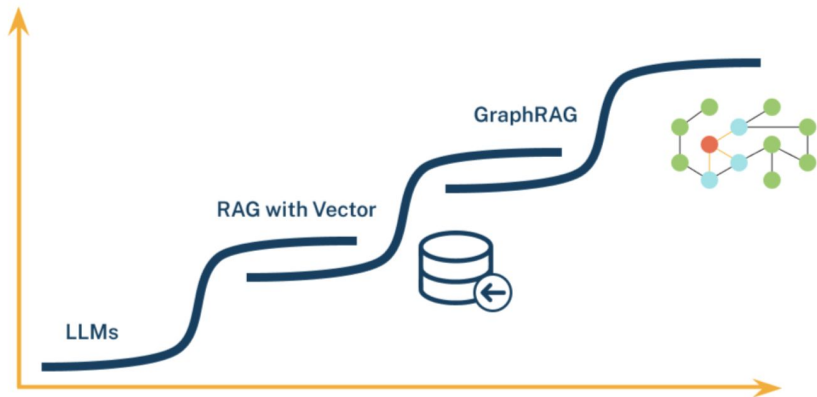
The GraphRAG Manifesto: Adding Knowledge to GenAI




Philip Rathle, Chief Technology Officer, Neo4j

Jul 11 · 22 mins read

We're Entering the "Blue Links" Era of RAG



KG



LLM

Better together

neo4j

Thank you

Join Alison Cossette tomorrow @ 2:30 PM
For Getting Started with GraphRAG





<https://www.deeplearning.ai/short-courses/knowledge-graphs-rag/>

GraphAcademy — Free, Self-Paced, Hands-on Online Training

<https://graphacademy.neo4j.com/>



Building Knowledge Graphs: A Practitioner's Guide

By Jesús Barrasa & Jim Webber


Publisher: O'Reilly

Available Formats: PDF - EN US

O'REILLY®

Building Knowledge Graphs

A Practitioner's Guide



Jesús Barrasa
& Jim Webber

Free e-book!

<https://neo4j.com/knowledge-graphs-practitioners-guide/>



What's Next?

Fine
tuning

RAG

Graph
RAG

GraphRAG Retrieval Patterns

The core GraphRAG pattern is very simple, yet powerful.

1. Do a vector search to find an initial set of nodes
2. Traverse the graph around those nodes to add context

Optional: Rank the results using the graph and pass the top-k documents to the LLM

The Benefits of GraphRAG

**1. Higher
Accuracy**

**2. Easier
Development**

**3. Explainability
& Governance**



The Benefits of GraphRAG

1. Higher Accuracy

3X

A Knowledge Graph improves the accuracy of LLM responses by 54.2%, an average of 3x




2. Easier Development

Feedback from an AI Engineer

09:05

Redacted

Ciao team  exciting

news from our caching infrastructure! I've been experimenting with `Neo4J` as a candidate for our cache storage and the results are very promising!

 [Here is the PR](#) with the changes

2.


Easier Development: Why

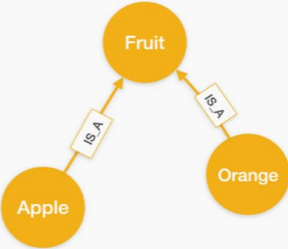
Natural language description: "Apples & oranges are both fruits"

From The Going Meta Series

DATA SEMANTICS

EXPLICIT (SYMBOLIC)

 Representation



```
graph BT; Apple((Apple)) -- IS A --> Fruit((Fruit)); Orange((Orange)) -- IS A --> Fruit
```

2.

Easier Development: Why

Natural language description: "Apples & oranges are both fruits"

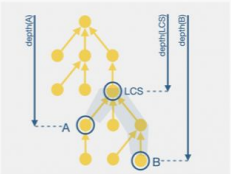
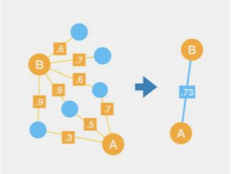
From The Going Meta Series

DATA SEMANTICS

EXPLICIT (SYMBOLIC)

Similarity calculation

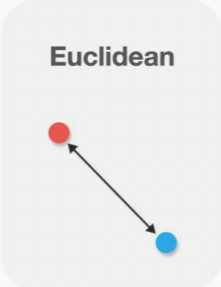
- Structural**
 - ▶ Node similarity
 - ▶ Overlap
 - ▶ Jaccard
- Taxonomy based**
 - ▶ Path
 - ▶ Leacock-Chodorow
 - ▶ Wu-Palmer



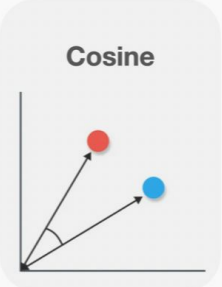
IMPLICIT (SUB-SYMBOLIC)

Similarity calculation

Euclidean



Cosine



Vector distance based

2.

Easier Development: Why

Natural language description: "Apples & oranges are both fruits"

From The Going Meta Series

DATA SEMANTICS

EXPLICIT (SYMBOLIC)

Search

Graph exploration

IMPLICIT (SUB-SYMBOLIC)

Search

Indexing:

- IVF
- HNSW

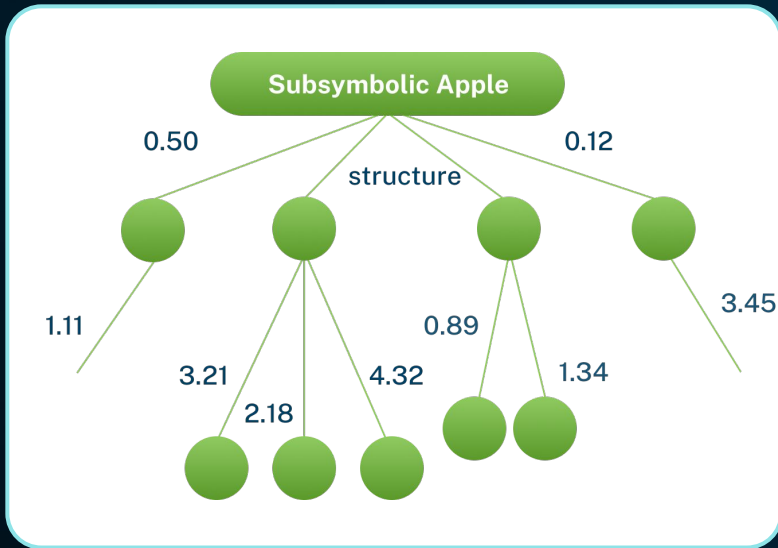
ANN (Approx. Nearest Neighbours)

2. Easier Development: Why

Human View of an Apple



LLM + Vector View of an Apple



Detour! Knowledge Graph Creation

Unstructured data

Typically PDFs or other text documents

Watch This Space!



Mix of structured & unstructured data

Structured data with long-form text

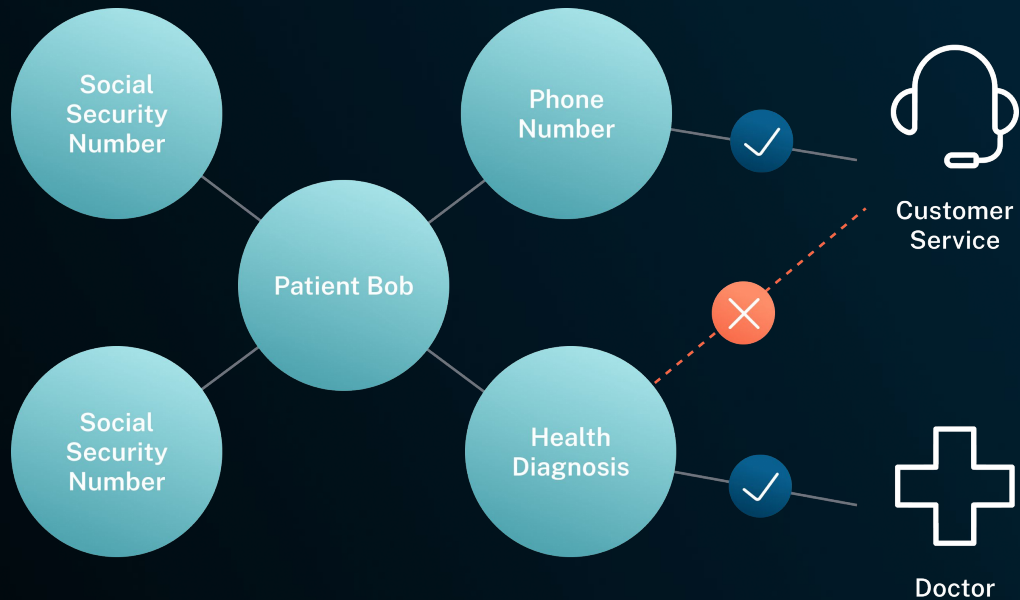


Structured data

Structured data with short text values



3. Explainability & Governance




The GraphRAG Manifesto: Adding Knowledge to GenAI



Philip Rathle, Chief Technology Officer, Neo4j

Jul 11 · 22 mins read


We're Entering the "Blue Links" Era of RAG



ChatGPT marked the beginning of a
GenAI hockey stick moment

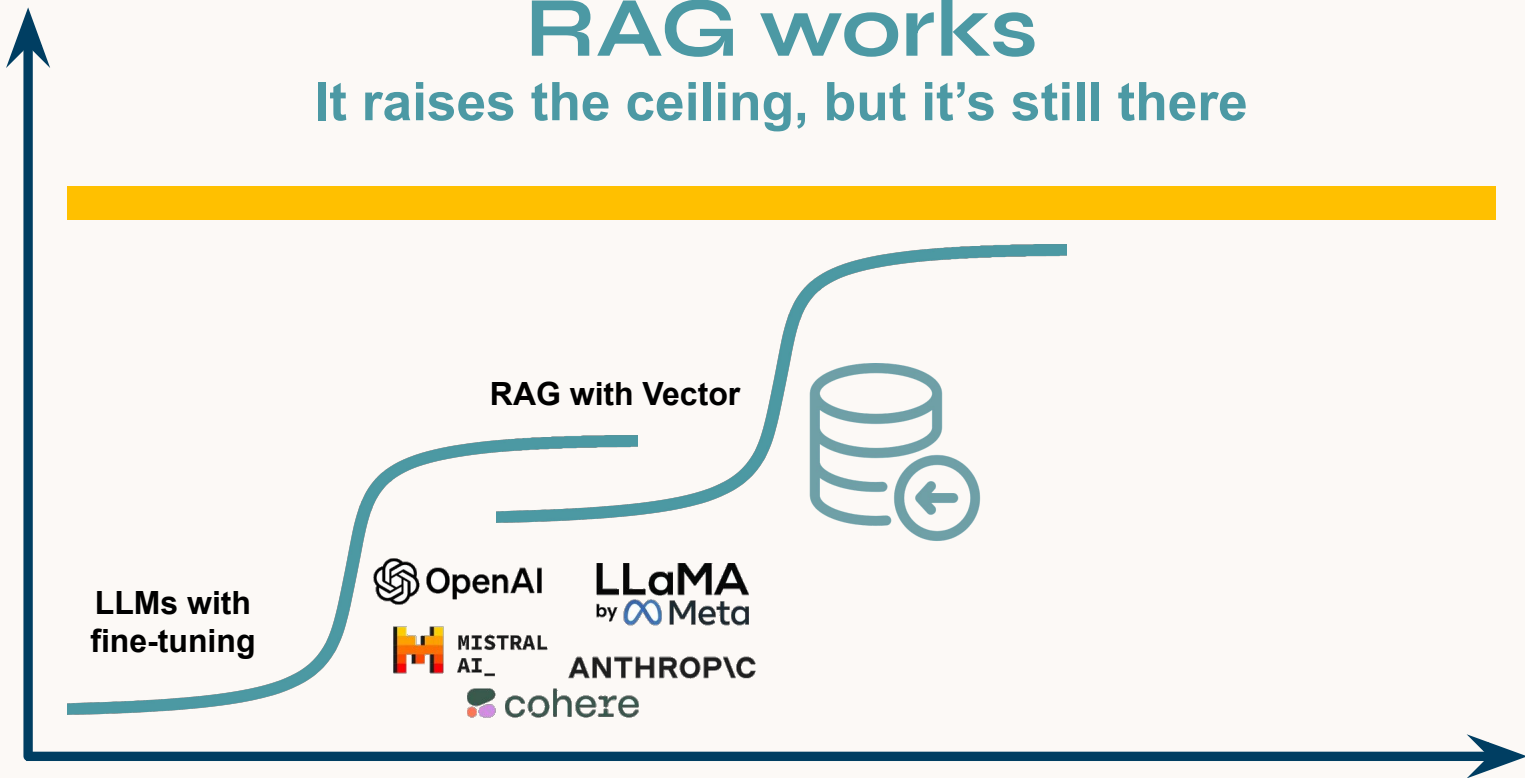


GenAI Go-to-production Barrier



RAG works

It raises the ceiling, but it's still there





Breaking through the next ceiling



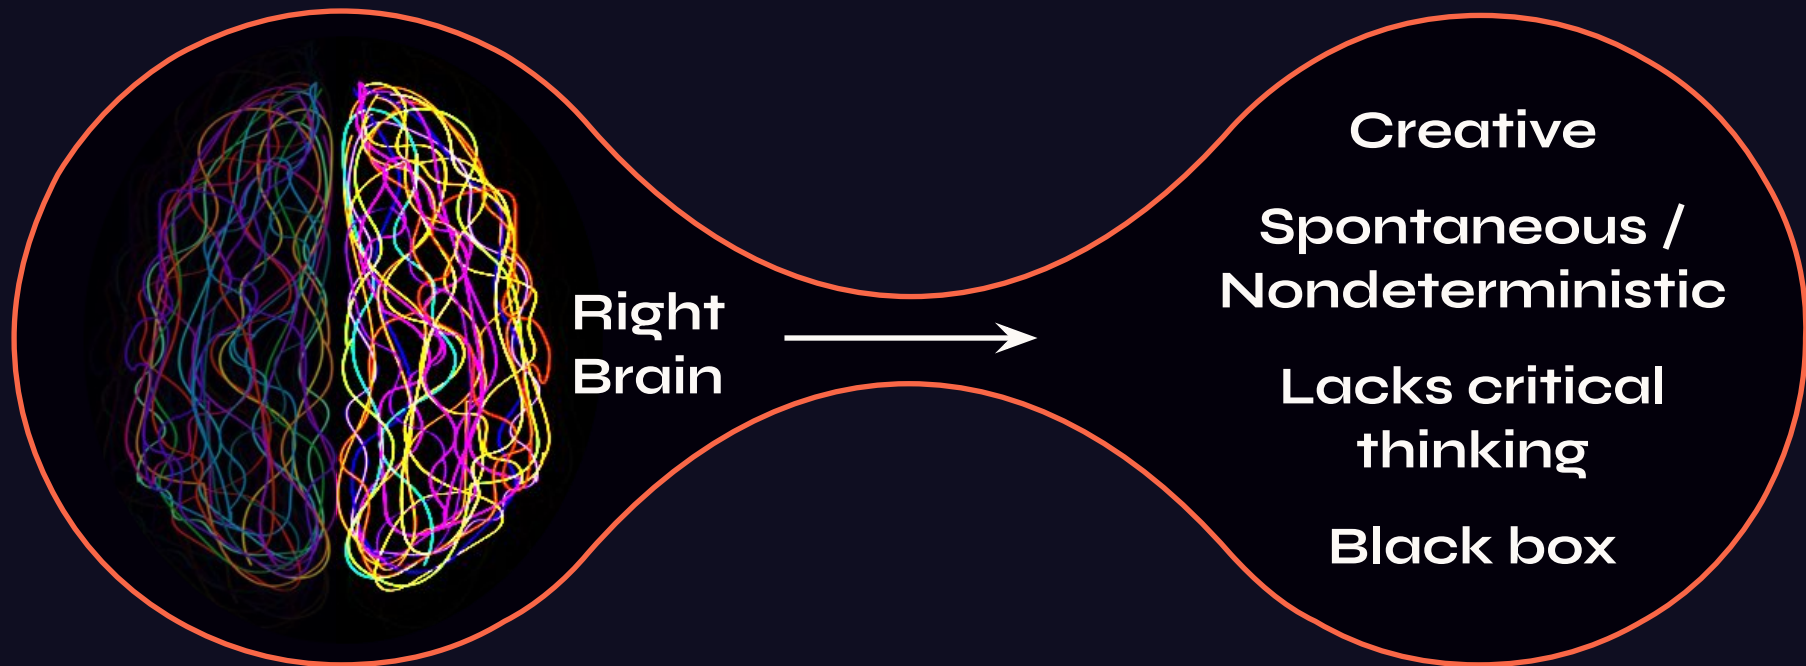
An Analogy



Why do knowledge graphs help with GenAI

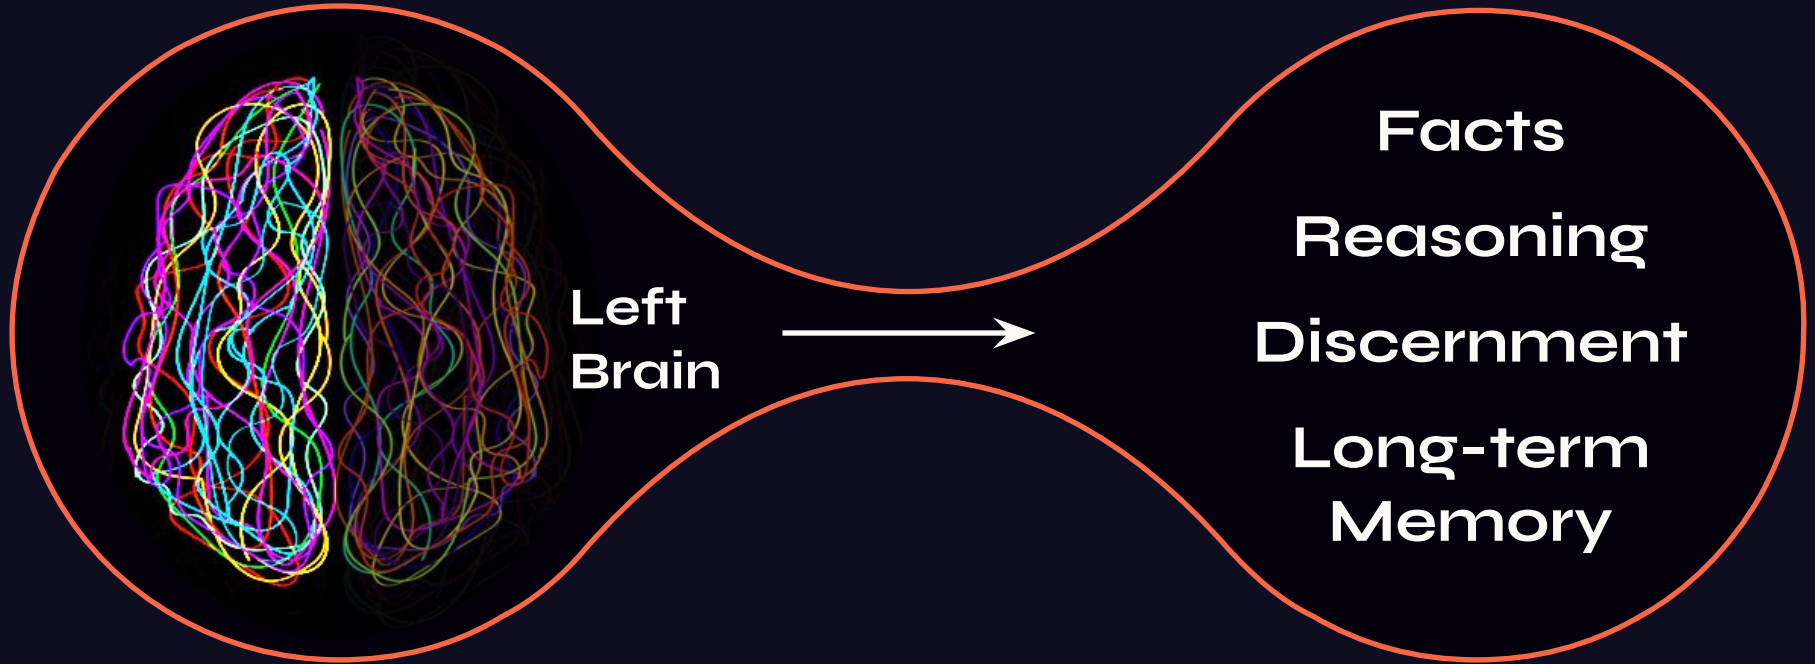
LLMs with Vector-Based RAG

a 'word model' view




Knowledge Graph

a 'world model' view



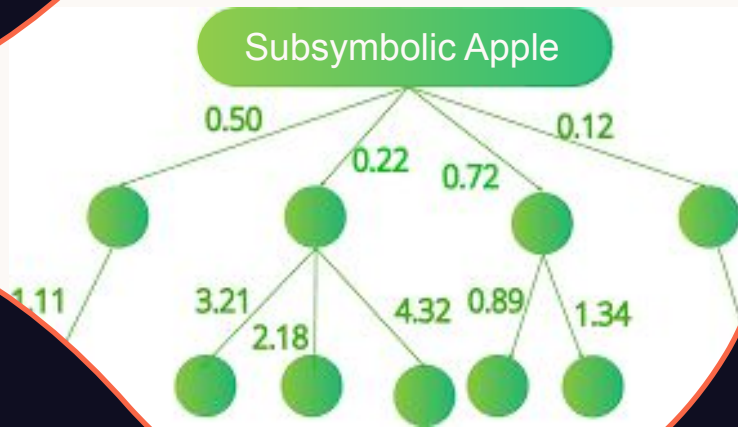
LLM + Vector View of an Apple



Human View of an Apple



LLM + Vector View of an Apple



Lack explicit domain understanding



Deal in words, not meaning



Answers emerge from language statistics



Provide nondeterministic answers




Lack explainability




ETHICAL, BIAS, ACCURACY CONCERNS

LLM + Vector View of an Apple



Knowledge Graph View of an Apple



Deal with concepts and meaning



Support human & machine reasoning



Relate concepts, reveals emergent statistics



Answers are deterministic




Support explainability & human understanding




CREATES A SHARED WORLD VIEW

Knowledge Graph View of an Apple



GraphRAG @Nvidia

Jensen Huang's Nvidia GTC Keynote



GraphRAG
@Microsoft

teamgraphrag@
microsoft.com



The image shows a screenshot of a Microsoft Research webpage. At the top, there is a navigation bar with a hamburger menu icon, a search icon, and the Microsoft logo. Below the navigation bar, the word "Research" is displayed with a dropdown arrow. The main content area features a large, vibrant, multi-colored network graph visualization with nodes and edges in shades of green, blue, purple, and orange. Below the graph, the text "Project GraphRAG" is prominently displayed in a large, bold, black font. Underneath this, the subtitle "LLM-Derived Knowledge Graphs" is written in a smaller, black font.