



Intro to GraphRAG Workshop

AI Engineer World Fair 2025

Zach Blumenfeld | Graph Product Specialist

Agenda

1 Workshop objectives
Setting the stage

2 Groups and setup
Let's get set up

3 Module 1
Graph basics: loading,
queries, vectors

4 Module 2
Unstructured Data

5 Module 3
GraphRAG and Agents

6 Wrap up
Resources and Q&A

Workshop Rules

- Ask questions straight away, this is an interactive session
- Raise your hand if you are stuck
- Slides & notebooks will be shared
- Have fun!

Before We Start



#workshop-graphrag-intro

Connect to the notebooks:

attendee101 -> 160

<https://bit.ly/neo4j-graphrag-one>

attendee201 -> 260

<https://bit.ly/neo4j-graphrag-two>



Neo4j Browser : <https://browser.neo4j.io/preview/>

Setup Environment

```
$ git clone https://github.com/neo4j-product-examples/genai-workshop.git
```

```
$ cp ws.env genai-workshop/talent
```

Navigate to **workshop/talent** and open module. No need to pip install, everything should be good to go.

Quick Poll (by show of hands)



AI Models

Embedder

LLM

UI

Knowledge Assistant


Semantic Search, Automation,
Personalized Experiences

Knowledge Graph

Graph Analytics

Graph Enrichment

(Communities, Entity Resolution, and More)



Entity
Extraction

Graph
Retrieval

Standard ETL

Agents

Tools

Local Retriever

Graph vector search

Drift Retriever

Dynamic depth graph traversal

Global Retriever

Comprehensive communities

...

Auxiliary Search

specific cases
(RFPs, HR benefits)

Data Sources

Unstructured

Documents...

Structured

Tables...


Other External Services

E.g. Cognitive
Search

Skills & Employee Graph

Talent search ,
skill alignment & analysis,
staffing & team formation





Module 1

Graph basics: Queries, Algorithms & Vectors

Module 1

Graph Basics

- **Creating a Graph from Structured Data**
- **Basic Cypher Queries and Pattern Matching**
- **Graph Algorithms**
- **Text Embeddings for Semantic Analysis**
- **Vector Search**


Knowledge Graph = design patterns to organize & access interrelated data

Property Graph Data Model

Nodes represent entities in the graph


Relationships represent associations or interactions between nodes

Properties represent attributes of nodes or relationships



010...010
011..0011
..n
10011..00
1....010...
01001..0
011..n
10011..00
1....
010...010
011..0011
..n
10011..00
1....


Cypher: A Powerful & Expressive Query Language



RETURN p.name as person, skill

Embeddings?

- A type of data compression
- Transform messy data (text, images, audio, etc.) into a compact format for ML algorithms
- Most often numeric **vectors** (a.k.a arrays) with 100s or 1000s of elements
- Preserve information such that **“similar” items have proportionally “similar” embedding vectors**
- Similarity is measured with vector algorithms (cosine, euclidean, etc.)



Search & Vectors in Neo4j

Neo4j makes search efficient through a variety of indexes including

- **Range:** General index for predicates based on equality and range.
- **Point:** Predicts on geospatial points like distance bounding boxes, etc.
- **Text:** Predicates on strings like contains, ends with, etc.
- **FullText:** Text search based on tokenization and analyzers
- ***Vector: ANN (Approximate Nearest Neighbor) search on vectors***

How is similarity measured

What would you use for openai embedding models?

For OpenAI embedding models, **cosine similarity** is the preferred choice. Here's why:

Why Cosine Similarity?

1. OpenAI Embeddings Are Normalized

- OpenAI embedding models (like `text-embedding-ada-002`) generate **normalized vectors**, meaning each vector has a unit length (i.e., a magnitude of 1).
- In such cases, **cosine similarity and Euclidean distance are proportional**, but cosine similarity is computationally more efficient.

2. Focus on Semantic Meaning

- Since cosine similarity measures the **angle between vectors**, it captures semantic similarity between words, sentences, or documents effectively.
- Euclidean distance, on the other hand, is sensitive to magnitude differences, which aren't meaningful in normalized embeddings.

3. Industry Best Practice


- Most applications using OpenAI embeddings (e.g., search ranking, recommendation systems, and semantic clustering) rely on cosine similarity.
- OpenAI's own documentation recommends using **cosine similarity or dot product** over Euclidean distance.

50+ Graph Algorithms in Neo4j




Pathfinding & Search

- Shortest Path
- Single-Source Shortest Path
- All Pairs Shortest Path
- A* Shortest Path
- Yen's K Shortest Path
- Minimum Weight Spanning Tree
- K-Spanning Tree (MST)
- Random Walk
- Breadth & Depth First Search




Link Prediction

- Adamic Adar
- Common Neighbors
- Preferential Attachment
- Resource Allocations
- Same Community
- Total Neighbors



Centrality / Importance

- Degree Centrality
- Closeness Centrality
- Harmonic Centrality
- Betweenness Centrality & Approx.
- PageRank
- Personalized PageRank
- ArticleRank
- Eigenvector Centrality




Similarity

- Euclidean Distance
- Cosine Similarity
- Node Similarity (Jaccard)
- Overlap Similarity
- Pearson Similarity
- Approximate KNN



Community Detection

- Triangle Count
- Local Clustering Coefficient
- Connected Components (Union Find)
- Strongly Connected Components
- Label Propagation
- Louvain Modularity
- K-1 Coloring
- Modularity Optimization



Embeddings

- Node2Vec
- Random Projections
- GraphSAGE

... Auxiliary Functions:

- Random graph generation
- Graph export
- One hot encoding
- Distributions & metrics


Today we will explore **Community Detection**

Evaluate how groups of nodes may be clustered or partitioned in a graph.

Community id properties assigned to node based on relationship structure.

Useful for:

- Segmentation
- Clustering
- Entity resolution
- Summarization (for AI)



Module 1

Go to Notebook





Module 2

Taming Unstructured Data

Module 2


Unstructured Data

- **Creating a graph from unstructured data**
- **Entity extraction using domain model + LLM**

Module 2

Go to Notebook





From Unstructured to Structured: An example...



RFP Generation GenAI App



Why a GraphRAG Matters in RFP GenAI App?

Challenges



Outcomes




Time consuming to read previous RFP across **multiple repositories**



Repetitive and **manual tasks** to synthesise the content




Non-standard structure of RFP making it difficult to do data modelling



Knowledge base to collect, store and retrieve **domain-specific** information



Drive efficient, accurate, contextual and **explainable** way to streamline RFP responses



Flexible storage that's adoptable to the varying structure of an RFP

Anatomy of an RFP Document

AWS RFP

Intro

About the Company

Content

Financial Result

Content

Objectives

Content

Proposal

Subsection 1

Subsection 1.1

Content

Subsection 2

Content

Anatomy of a Document

AWS RFP

Intro

About the Company

Content

Financial Result

Content

Objectives

Content

Proposal

Subsection 1


Subsection 1.1

Content

Subsection 2

Content

RFP Document as a Graph



Why Neo4j KG Matters in RFP GenAI App?

Challenges



Outcomes




Time consuming to read previous RFP across **multiple repositories**



Repetitive and **manual tasks** to synthesise the content




Non-standard structure of RFP making it difficult to do data modelling



Knowledge base to collect, store and retrieve **domain-specific** information



Drive efficient, accurate, contextual and **explainable** way to streamline RFP responses





Flexible storage that's adoptable to the varying structure of an RFP

Knowledge Graph as the Knowledge Base

Document in a KG



Knowledge Graph





Knowledge Graph as the Knowledge Base

Document in a KG



Knowledge Graph



Example Pattern

Name: Graph Enhanced Vector Search

Description: The user question is embedded using the same embedder used to create chunk embeddings. A vector similarity search is executed on the chunk embeddings to find k (number previously configured by developer/user) most similar chunks. A traversal of the Domain Graph starting at the found chunks is executed to retrieve more context.

Context: The biggest problem with basic GraphRAG patterns is finding all relevant context necessary to answer a question. The context can be spread across many chunks not being found by the search. Relating the real-world entities from the chunks to each other and retrieving these relationships together with a vector search provides additional context about these entities that the chunks refer to. They can also be used to relate chunks to each other through the entity network.


Required pre-processing: Use an LLM to execute entity and relationship extraction on the chunks. Import the retrieved triples into the graph.

Variations: Entity disambiguation, Question-guided/Schema-defined extraction, Entity embeddings, Ontology-driven traversal

```
MATCH (node)-[:PART_OF]->(d:Document)
MATCH (node)-[:__HAS_ENTITY__]->(e)
MATCH path=(e)((()-[rels:!HAS_ENTITY&!PART_OF]-()){0,2}(:!Chunk&!Document)
...
RETURN ...
```

AKA: Graph + Vector, Augmented Vector Search

Required graph pattern: Lexical Graph with Extracted Entities



Graph Enhanced Vector Search


Global Community Summary Retriever

AKA: Microsoft GraphRAG, Global Retriever
Context: Certain *questions that can be asked on a whole dataset* do not just relate to things present in some chunks but rather search for an overall message that is overarching in the dataset.

Required pre-processing: In addition to extracting entities and their relationships, we need to form hierarchical communities within the Domain Graph. For every community, an LLM summarizes the entity and relationship information into Community Summaries.

```
MATCH (c:__Community__)  
WHERE c.level = $level  
RETURN c.full_content AS output
```

Required graph pattern: Lexical Graph with Extracted Entities and Community Summaries




Add Domain Knowledge

- There are references in the document that give more context
- These add more context to the chunks


See also:

- [Abroad](#) (definitions)
- [DBC healthcare product](#) (definitions)
- [Rate](#) (definitions)
- [Treatment](#) (definitions)
- [Year](#) (definitions)



```
1 MATCH (c:Chunk)-[:MENTIONS]-(d:Definition)
2 WHERE c.id in $chunk_ids
3 WITH DISTINCT d as d
4 RETURN d.definition as definition, d.description as description
```





Module 3


GraphRAG and Agents

Module 3

GraphRAG & Agents

- **Experiment with queries for an Agent**
- **Define Tooling**
- **Create an agents with the available tools**
- **Chatbot for an Agent**
- **Text2Cypher (if we got time)**

Ordinary RAG




GraphRAG: RAG with graph

It's actually quite broad...

GraphRAG is Retrieval Augmented Generation (RAG) using Knowledge Graphs

Improves GenAI by taking advantage of rich graph data structures



Retrieval Patterns



Community summaries & graph enrichments



Graph enhanced vector search



Dynamic graph query generation




+ more: graph vectors, parent-child retrievers, etc.

➤ Accurate, Contextual and Explainable


Who is the main respondent
of the AWS RFP?

GenAI App


➤ Accurate, Contextual and Explainable



Accurate, Contextual and Explainable




➤ Accurate, Contextual and Explainable




➤ Accurate, Contextual and Explainable

Similarity Search using
Neo4j Vector Index




➤ Accurate, Contextual and Explainable

Contextual Knowledge
Retrieval within
Neo4j KG



➤ Accurate, Contextual and Explainable


Knowledge Retrieval
to aid in
Explainability






Accurate, Contextual and Explainable


Fine Grained Access Control to prevent unwarranted Knowledge Retrieval




➤ Accurate, Contextual and Explainable



Accurate, Contextual and Explainable



Accurate, Contextual and Explainable



Module 3

Go to Notebook





Wrap Up!



AI Models

Embedder

LLM

UI

Knowledge Assistant
Semantic Search, Automation,
Personalized Experiences

Knowledge
Graph

Graph Analytics

Graph Enrichment
(Communities, Entity Resolution, and More)

Agents

Tools

Local Retriever
Graph vector search

Drift Retriever
Dynamic depth graph traversal

Global Retriever
Comprehensive communities

...

Auxiliary Search
specific cases
(RFPs, HR benefits)

Data Sources

Unstructured

Documents...


Structured

Tables...

Other External Services

E.g. Cognitive
Search

Entity
Extraction




Standard ETL

Graph
Retrieval

Agentic GraphRAG with a Graph DB


- **Accuracy** - Organize your data so agents can answer more questions accurately
- **Explainability**: Better provide & control reasoning for data retrieval
- **Acceleration**: Start with a simple knowledge graph to quickly improve retrieval over key use cases. Easily expand & iterate from there.

Graph Retrieval Augmented Generation



Easier Development

Feedback from an AI Engineer



 Here is the PR with the changes

I kinda replicated the same action-based cache already in place for Pinecone but thanks to the graph nature of Neo4J most of the operations yield better results:

- thanks to the [Neo4j graph data science](#) plugin we can store embeddings and calculate cosine similarity at the database level
- getting related actions is as simple as following the relationships between nodes
- **the cache can be visualised.** This is an extremely valuable debugging tool for us to understand if/when and how the cache might be broken/misbehaving (I actually already fixed a couple of bugs just thanks to this 🧐)

GraphRAG Resources



Github repository with
this workshop



Developer guides and
coded examples



Meetup Today (06/03) @5: AI Agent SF #3 -
Agent Protocols (Graphs & MCP) w/ Michael
Hunger


Next Workshop @ 1:00 Golden Gate Ballroom C

**Graph Intelligence: Enhance Reasoning and
Retrieval Using Graph Analytics w/ Alison Cossette**

Swing By Our Booth!

Booth # G9

- Neo4j MCP Servers
- Google ADK Examples
- Knowledge Graph Construction
- and more!



An aerial photograph of a pod of dolphins swimming in the ocean. The dolphins are dark grey and are scattered across the frame, mostly in the upper right quadrant. The water is a deep blue with visible ripples and small waves. In the top left corner, there is a light blue abstract shape resembling a stylized flower or a cluster of bubbles. In the bottom left corner, there is a darker blue abstract shape resembling a stylized flower or a cluster of bubbles. On the right side, there is a large, solid orange shape that curves along the edge of the frame.

Thank you!