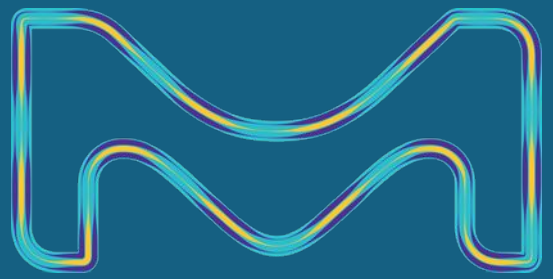
The Healthcare business of Merck KGaA, Darmstadt, Germany operates as EMD Serono in the U.S. and Canada.

From Knowledge Graph to Scientific Workflow: Operationalizing Graph Intelligence at EMD Serono

Lessons from traversing the data-to-decision links

Jesper Ryge, Aleksander (Olek) Siwek, Sergiusz (Sergio) Wesolowski

RU N&I Computational Biology team @ EMD Serono.



EMD SERONO

Agenda

01 — From Knowledge Graph to Scientific Workflow

Why Knowledge Graphs?
Why Knowledge Graphs alone are not enough?

02 — Beyond Exploration and Knowledge Retrieval

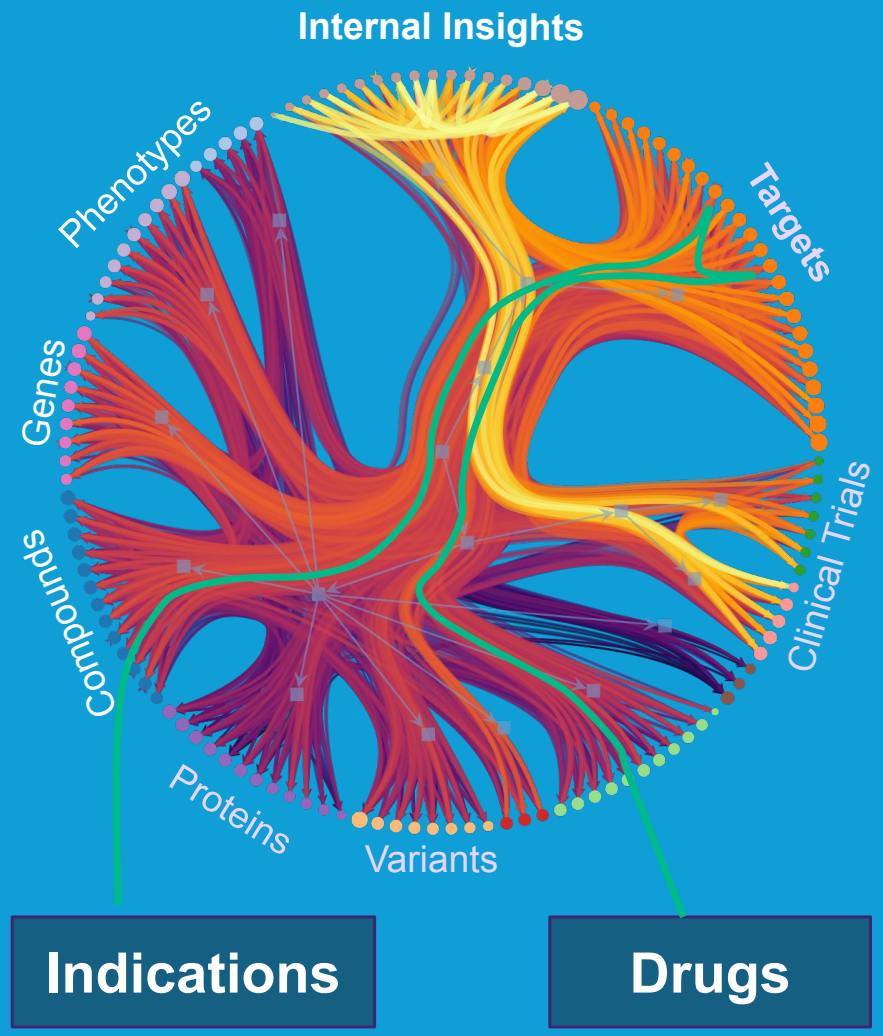
Context-specific relevance scoring
Contrastive relevance
Community detection
Link prediction

03 — Putting It to Work

From analytical outputs to a unified chat interface.

04 — Lessons Learned

What worked, what did not.



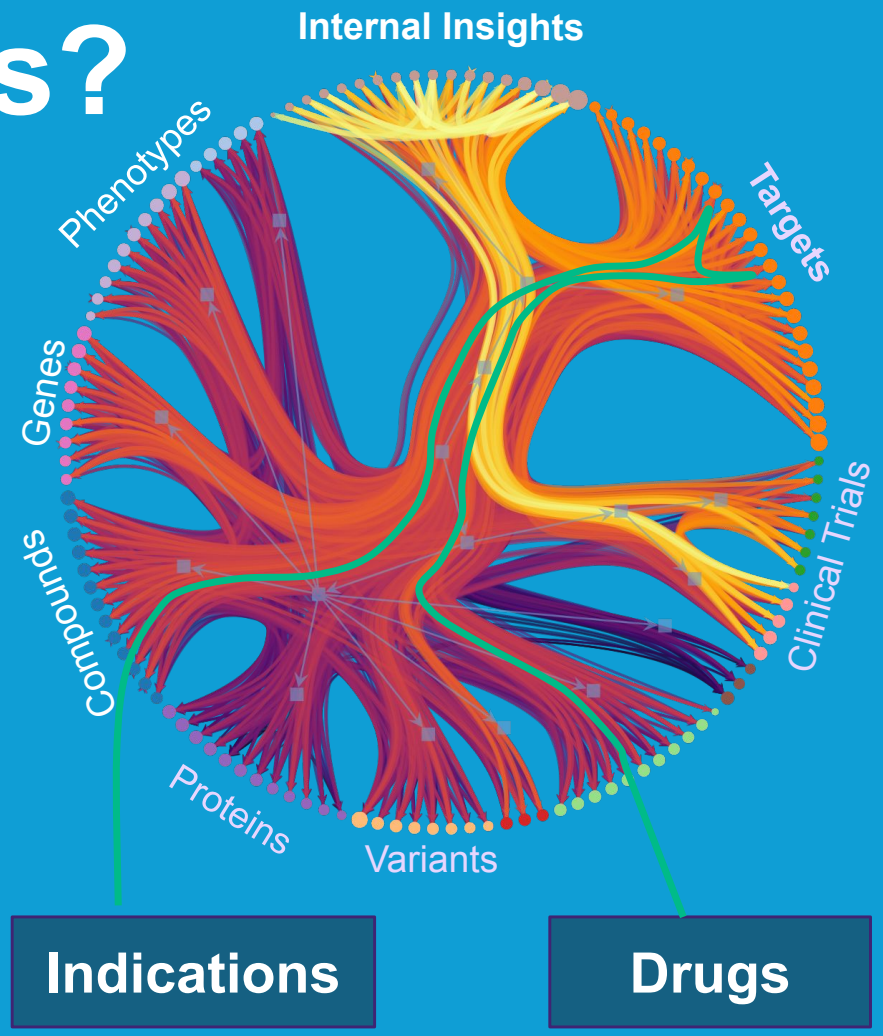
Why Knowledge Graphs?

Connect heterogeneous biomedical knowledge
Bring diseases, targets, drugs, pathways, mechanisms, multiple data sources, modalities and internal insights into one connected semantic representation.

Preserve biological context
Relationships between concepts become unified information layer forming multimodal context to a scientific question

Enable traceable analysis
Analytical results can remain connected to the underlying relationships, and supporting evidence.

Provide an AI/ML-ready representation
The same graph can support retrieval, graph analytics, machine learning and agent-based workflows.



From having knowledge graph to using It

The KG connected representation as the starting point for scientific analysis

Situation

- Drug-discovery questions span multiple modalities:
  targets, diseases, pathways, drugs, …
- Relevant evidence spans **multiple scales** (molecular, genomic, clinical)
- Knowledge graphs provide a **connected backbone for integrating and exploring this evidence in context**.
- KG semantic structure also makes the data **AI/ML-ready** — supporting both **structured retrieval** and **computational mining with graph analytics and machine learning**.

Operationalization task turns graph access into usable scientific workflows.

From having knowledge graph to using It

The KG connected representation as the starting point for scientific analysis

Situation

- Drug-discovery questions span multiple modalities:
  targets, diseases, pathways, drugs, …
- Relevant evidence spans **multiple scales** (molecular, genomic, clinical)
- Knowledge graphs provide a **connected backbone for integrating and exploring this evidence in context**.
- KG semantic structure also makes the data **AI/ML-ready** — supporting both **structured retrieval** and **computational mining with graph analytics and machine learning**.

Challenge

- A knowledge graph does not by itself tell us **what is relevant to a specific scientific question**.
- **Retrieval is useful, but not enough**
  Retrieval does not equal inference.
- LLM-driven workflows can lose **traceability and reproducibility** unless grounded in explicit graph structure, methods and evidence.
- **One graph representation does not fit every problem** — different questions may require different scopes or projections.

Operationalization task turns graph access into usable scientific workflows.

From having knowledge graph to using It

The KG connected representation as the starting point for scientific analysis

Situation

- Drug-discovery questions span multiple modalities:
  targets, diseases, pathways, drugs, …
- Relevant evidence spans **multiple scales** (molecular, genomic, clinical)
- Knowledge graphs provide a **connected backbone for integrating and exploring this evidence in context**.
- KG semantic structure also makes the data **AI/ML-ready** — supporting both **structured retrieval** and **computational mining with graph analytics and machine learning**.

Challenge

- A knowledge graph does not by itself tell us **what is relevant to a specific scientific question**.
- **Retrieval is useful, but not enough**
  Retrieval does not equal inference.
- LLM-driven workflows can lose **traceability and reproducibility** unless grounded in explicit graph structure, methods and evidence.
- **One graph representation does not fit every problem** — different questions may require different scopes or projections.

Solution

Operationalization task turns graph access into usable scientific workflows.

From having knowledge graph to using It

The KG connected representation as the starting point for scientific analysis

Situation

- Drug-discovery questions span multiple modalities:
  targets, diseases, pathways, drugs, …
- Relevant evidence spans **multiple scales** (molecular, genomic, clinical)
- Knowledge graphs provide a **connected backbone for integrating and exploring this evidence in context**.
- KG semantic structure also makes the data **AI/ML-ready** — supporting both **structured retrieval** and **computational mining with graph analytics and machine learning**.

Challenge

- A knowledge graph does not by itself tell us **what is relevant to a specific scientific question**.
- **Retrieval is useful, but not enough**
  Retrieval does not equal inference.
- LLM-driven workflows can lose **traceability and reproducibility** unless grounded in explicit graph structure, methods and evidence.
- **One graph representation does not fit every problem** — different questions may require different scopes or projections.

~~Solution~~ Lessons

Components that are necessary:

- **ML-driven decision support** — ranking, clustering and prediction on top of the graph.
- **Task-specific graph analytics** — adapt scope, projections and methods to the scientific question.
- **Chat-based interface** — make graph queries and analytical workflows accessible
- **Explainable outputs** — keep results connected to graph structure, methods and evidence.

Operationalization task turns graph access into usable scientific workflows.

Exploration is the starting point

Practical ways to interrogate knowledge graphs exist

Useful entry points exist and are readily available

- **Neo4j Browser** for direct inspection and exploration.
- **Neo4j MCP** for schema-aware access and natural-language interaction.
- These lower the barrier to asking questions of the graph.

But large graphs can be complex

- Large, complex schemas.
- Efficiency vs Readability tradeoffs.
- Even relatively simple exploration starts to require graph-specific knowledge.

Adding domain guidance

Useful: **KG-specific skills / subagents** that encode schema knowledge, example questions and reusable query patterns

a. They help translate scientific intent into more consistent graph interactions
b. Different scientific teams can have different starting points and workflows
c. Simple tasks can be offloaded to cheaper models (write cypher, execute cypher, …)

Exploration helps scientists understand the available context [retrieval, not analysis].

Beyond Exploration

KG retrieval:

What is connected to the question?
What evidence is available around this concept?

KG inference:

What knowledge is relevant for the question?
What is a different or unique factor?
What is similar, between indications, species, targets … ?
What might be missing and what might be wrong?

Establishing relevance in a knowledge graph

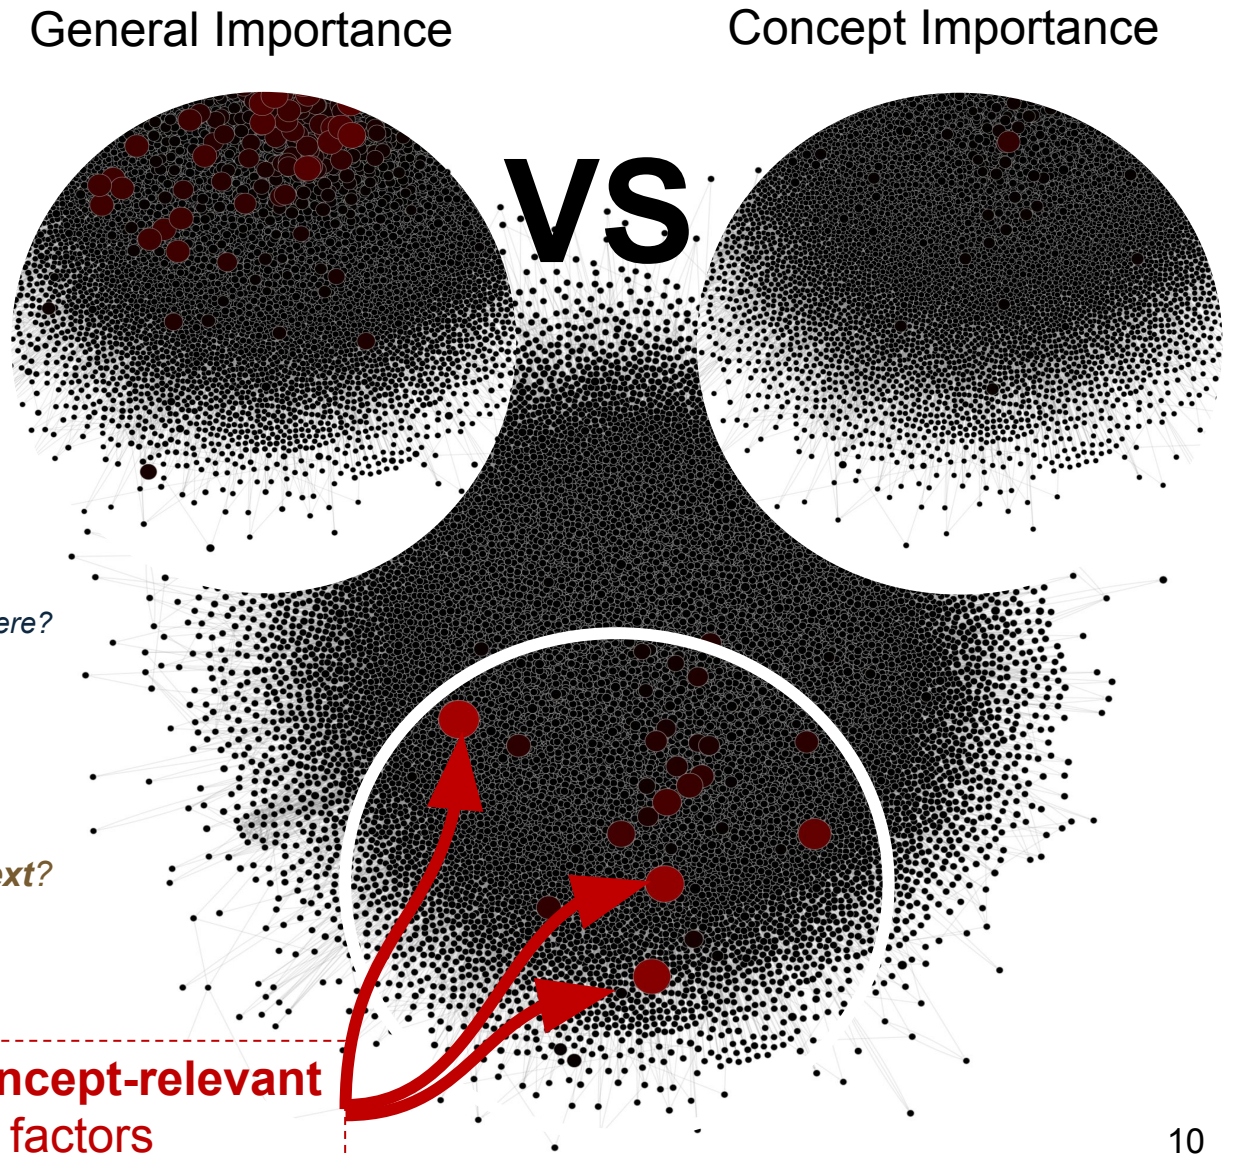
Borrowing from web search: relevance by centrality

PageRank was designed to identify important web-pages from patterns of connectivity.
The same idea can help identify **important biomedical concepts in a KG.**

- **Global relevance**
  PageRank over the full graph
  → *What is generally prominent?*
- **Local relevance**
  Personalized PageRank seeded on a disease, target, drug, etc.
  → *What becomes prominent around this context?*
- **Relative relevance**
  Local relevance / global relevance
  → *What is unusually relevant to this context, rather than simply prominent everywhere?*

Question Examples:

- *Which targets are particularly relevant to **Disease A**?*
- *How many times more relevant is pathway Z for Target X??*
- *Which concepts become important around a particular **program context**?*



Relative relevance distinguishes indication-specific signal

Contrasting relevance scores across biological contexts

What is shared and what is different?

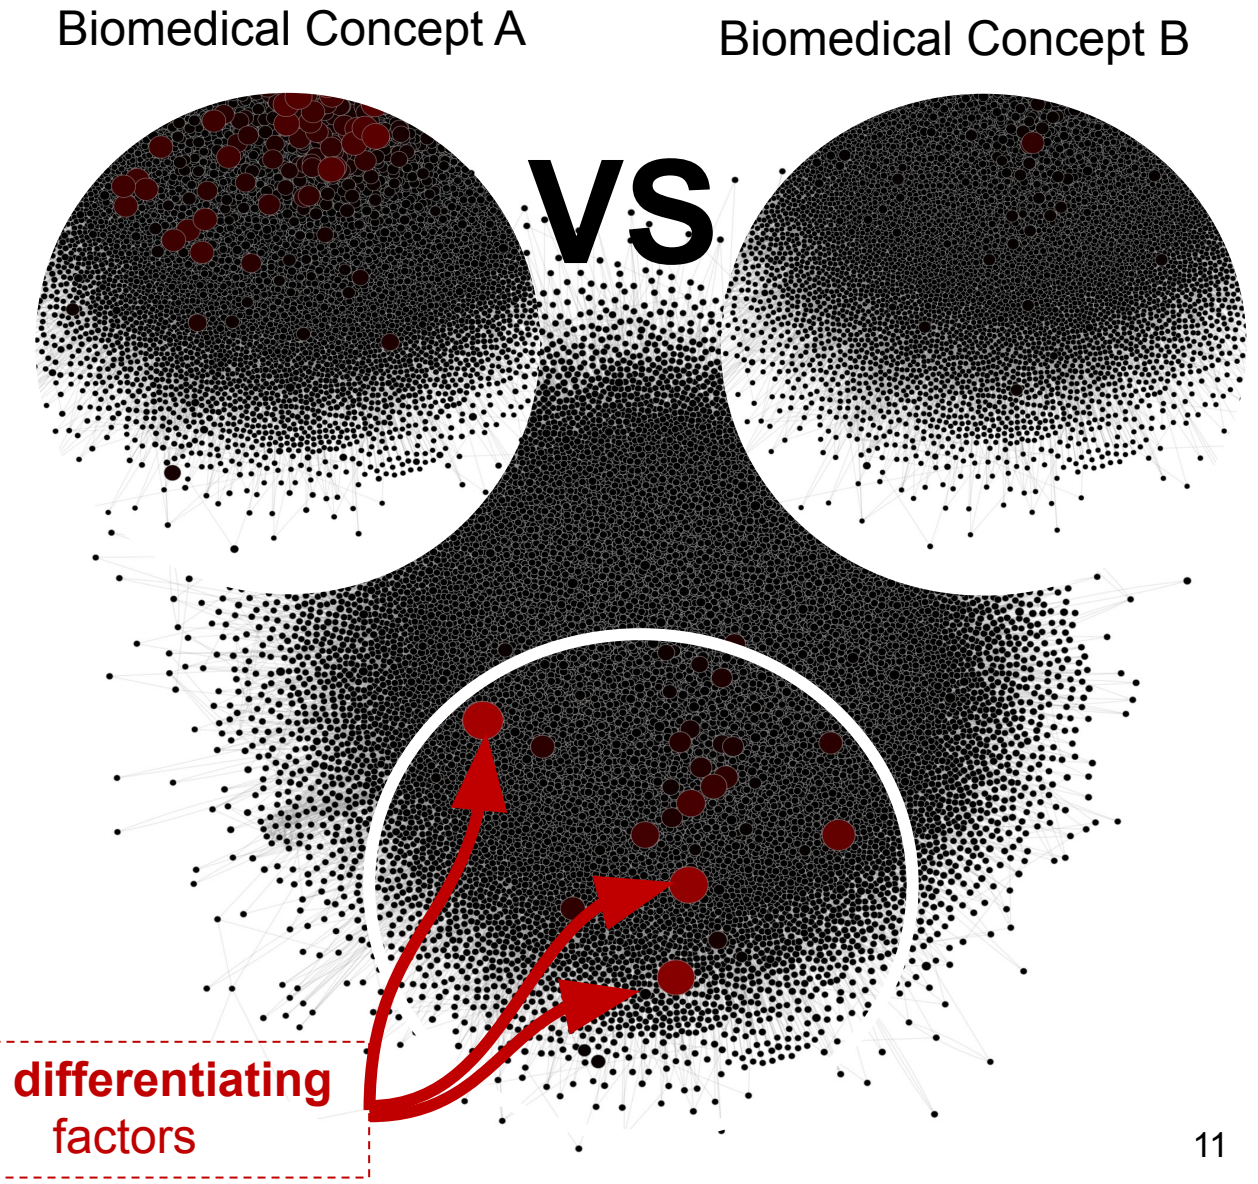
Comparing relevance profiles reveals what is shared — and what is differentiating

- Calculate context-specific relevance profiles for two diseases, targets, drugs, species or other biological contexts.
- Compare the profiles to identify **shared and differentiating concepts**.
- A simple relative comparison can turn graph relevance into an interpretable scoring tool.

Question Examples:

- *How much more relevant is **Target X in Disease A vs Disease B**?*
- *What is relevant in **oncology but not immune-mediated disease**?*
- *Which biological concepts differentiate **Drug A from Drug B**?*
- *How does relevance change across **species, or indications?***

The useful signal is often not absolute relevance, but how relevance changes between contexts.



Simple analytics, reusable workflow

Two Cypher GDS calls + a division yield context-specific relevance results

Global PageRank:

```
CALL gds.pageRank.stream($graphName, {scaler: 'L1Norm'})
YIELD nodeId, score
RETURN nodeId, score AS global_pr
```

Personalized PageRank:

```
CALL gds.pageRank.stream($graphName,
  {sourceNodes: $sourceNodes, scaler: 'L1Norm'})
YIELD nodeId, score
RETURN nodeId, score AS personalized_pr
```

Relative PageRank Ratio:

```
merged = global_df.merge(personalized_df, on="nodeId")
merged['ratio'] = merged['personalized_pr'] / (merged['global_pr'] + 1e-10)
```

Small analytical core → easy to package as a reusable skill through Neo4j MCP.

The computation is simple,
The scientific value comes from how the context and comparison are defined.

Community structure reveals higher-order biological context

Like social networks, biomedical graphs contain communities formed by patterns of connectivity

- In a social network, communities emerge because some people are more densely connected to one another than to the rest of the network.
- In a biomedical knowledge graph, diseases, genes, drugs, pathways and other concepts can form analogous structural communities.
- Community detection identifies groups of concepts with similar connectivity patterns.
- An LLM can then summarize the computed module with a biologically meaningful title and description.

Question Examples:

- Which biomedical concepts form coherent modules?
- Which diseases, targets, drugs or pathways share connection context?
- What entities could explain my MoA? What entities could share the MoA (similar connectivity pattern)?

Community Subgraphs provides biological context beyond neighborhood nodes and relationships.

From community structure to biological interpretation.

Like social networks, biomedical graphs contain communities formed by patterns of connectivity

1. Detect communities in the biomedical knowledge graph with community detection of your choice

2. LLM-Inspect community composition (names, modalities, descriptors)

3. Generate Explanation of the observed community pattern

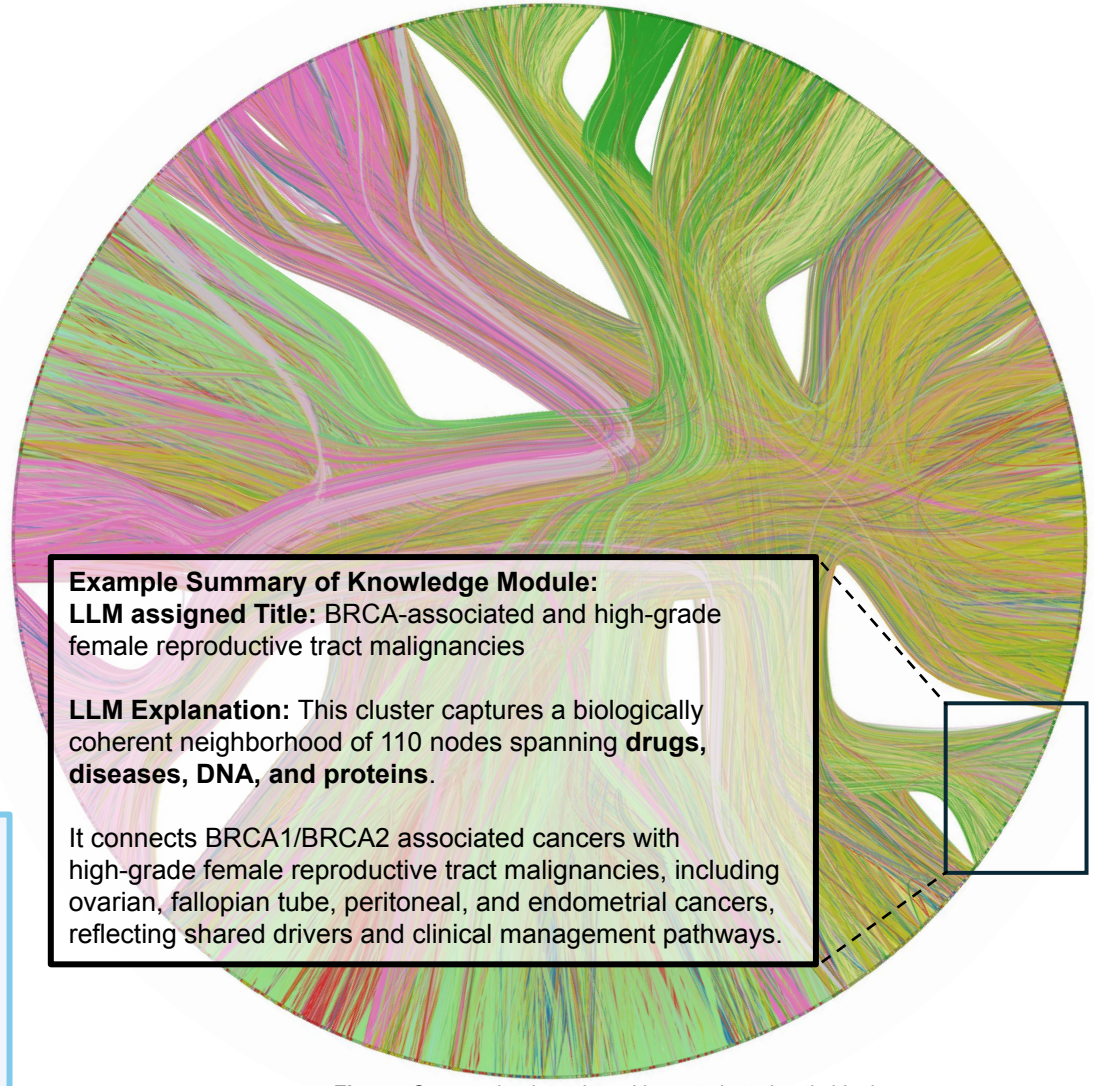


Figure:
over 170k biomedical concepts
spanning ~2m semantic connections:
grouped into >1,000 modules in 6 levels,
- Colors indicate biomedical modality
- Nodes represent a biomedical term
- Groups reflect shared connection patterns
- Edges represent curated semantic relationships

Figure: Community detection with nested stochastic block model (Peixoto, 2014) over Human subgraph

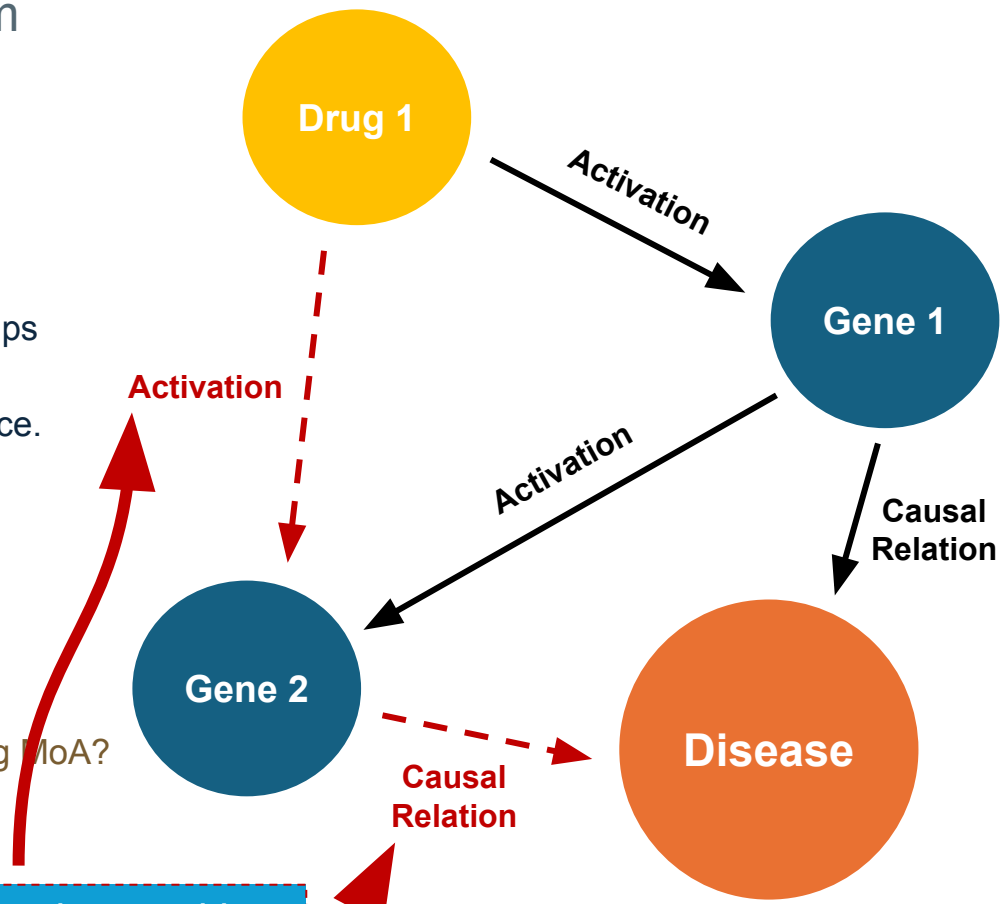
Beyond observed relationships: predicting what may be missing

Drug-discovery questions reframed as a missing-link problem

- Knowledge graphs capture known relationships, but potentially relevant relationships may be **unobserved or missing**.
- N**ovel target discovery, indication expansion, drug repurposing** can be formulated as missing link-prediction problems.
- Knowledge Graph Embedding models learn patterns from observed relationships and **score plausible missing links**.
- Predictions can then be ranked and investigated against existing graph evidence.

Questions

- Which **target–disease relationships** are plausible but currently missing?
- Which diseases are most plausible **novel indications** for drug A?
- What potential side effects could a new treatment have that spans the following MoA?



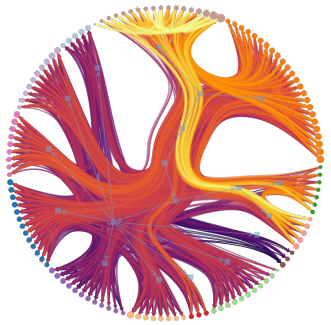
Link prediction uses the semantic structure present in the graph to prioritize plausible relationships that are not yet explicitly observed.

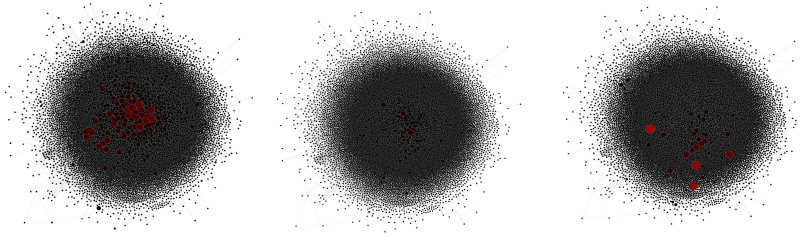
Predicted edges unobserved in the original knowledge graph, ranked by link-prediction score

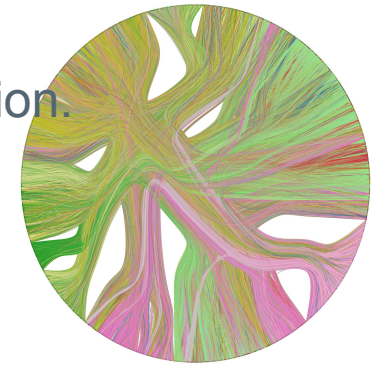
Putting it Together

Bringing the analytical layers together

Different graph methods answer different scientific questions.
No single method is the answer — the useful combination depends on the scientific question.







Claude

Modality

Concept Name

Global Relevance Score

Local Relevance Score

Relative Relevance Score

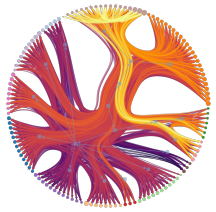
LLM Cluster Descriptors

Community Membership

node_label	node_name	global_r	IL35_narro	IL35_na	IL35_na	IL35_na	cluster_level_0_title	cluster_level_0_expl	level_0
Disease	tic disorder	0.000005	0.000302	56.233	2.57	-67	Dopaminergic CNS Disorders and Antipsychotic/In	This cluster centers on	15,6
Disease	neuroinflammatory disease	0.000001	0.000039	30.179	1.15	-81	Multiprotein Signaling Complexes: Kinase, Recept	This cluster represents	45,1
Disease	IgG2a response	0.000001	0.000039	29.689	1.54	-253	Multiprotein Signaling Complexes: Kinase, Recept	This cluster represents	45,1
Disease	immune response	0.000129	0.003585	27.812	1.92	-2	ER Stress, Mitophagy, and Metabolic Stress Signal	This cluster centers on	4,2
Disease	viral encephalitis	0.000002	0.000040	23.059	1.15	-77	Innate/Adaptive Immune Signaling: NF-κB, Inflam	This cluster represents	82,1
Disease	autoimmune inflammation	0.000002	0.000041	17.769	1.51	-244	Innate/Adaptive Immune Signaling: NF-κB, Inflam	This cluster represents	82,1
Disease	autoimmune response	0.000003	0.000042	15.183	1.15	-72	Innate/Adaptive Immune Signaling: NF-κB, Inflam	This cluster represents	82,1

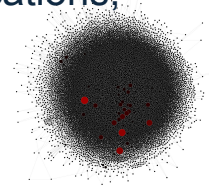
Architecture to access the graph and its analytics

Different graph methods answer different scientific questions.
No single method is the answer — the useful combination depends on the scientific question.



Talk directly to the graph via integrated Graph-schema aware Neo4j MCP connector.

Rank and Prioritize targets, indications, pathways with graph based interpretable scores.
(e.g. Relative PageRank)



Agent-Ready Graph Intelligence Modules



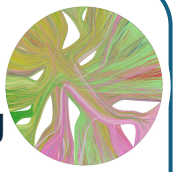
PyKEEN v1.9
Predictions for the People

PyG

Find new, leveraging KG embedding models for missing link prediction.
(e.g. ComplExN3, Rotate, GNNs)

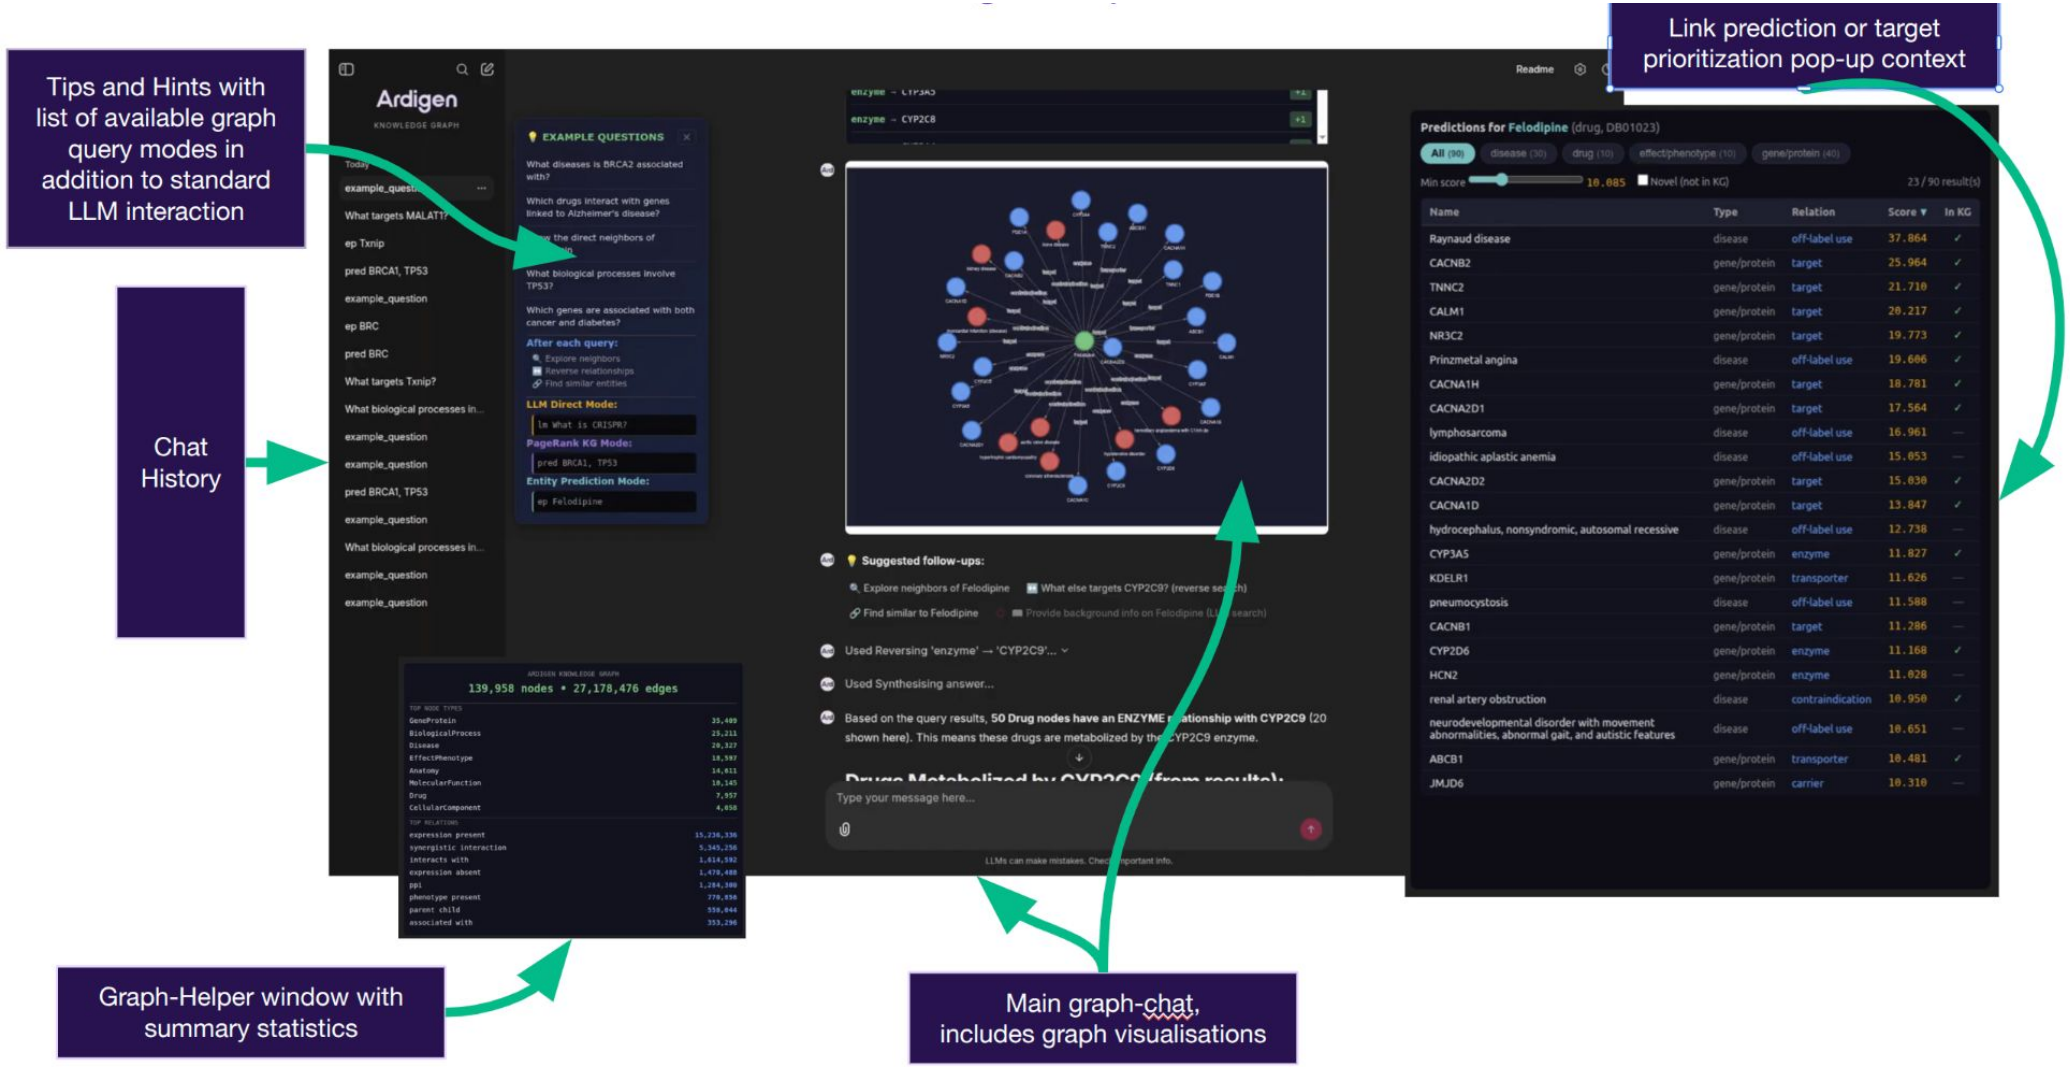
Your idea here

Provide context to findings with community-detection-based clustering module



One Interface to the Graph and Its Analytics

A common access layer that brings retrieval, analysis and exploration into one workflow.



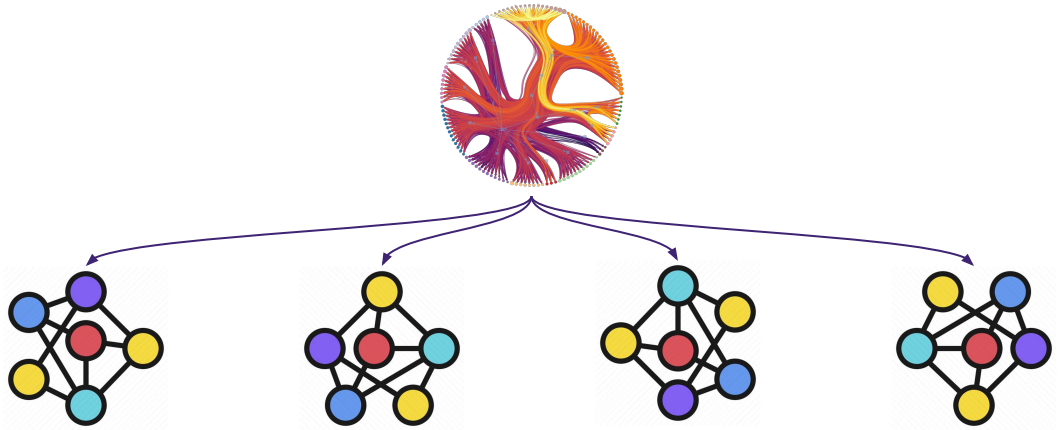
Lessons Learned

Lessons from operationalization

The scientific task shapes the graph, tools and interface

- **You may not need to build another interface.**
  Focus on well-defined tools, skills and an MCP layer; existing chat/agent interfaces can provide the user experience.

- **One central graph does not have to serve every question.**
  Use **task-specific subgraphs and graph projections** matched to the task, modality, species or analytical problem.



Purpose-scoped subgraph model **different, but overlapping** aspect of Drug Discovery
Specialized graph models, not LLMs, calculate **interpretable scores** within each **dedicated context**

Lessons from operationalization

The scientific task shapes the graph, tools and interface

- **You may not need to build another interface.**
  Focus on well-defined tools, skills and an MCP layer; existing chat/agent interfaces can provide the user experience.

- **One central graph does not have to serve every question.**
  Use **task-specific subgraphs and graph projections** matched to the task, modality, species or analytical problem.

- **Neo4j is the core backbone, not the boundary.**
  A lot of analytics can run directly on the graph, but **heavy ML workflows, specialized graph algorithms or integration with tabular data** may be better handled externally and connected back to the graph.
  Conduct analysis out of core -> integrate results into Neo4j structure.

- **BioCypher and BioLink**
  Helpful for reproducible ingestion and schema/ontology alignment

- **Graph size and schema matter operationally.**
  Very large graphs can make interactive queries and agent workflows noticeably slower; **purpose-scoped projections can improve both relevance and response time**.

Lessons from operationalization

The scientific task shapes the graph, tools and interface

- **You may not need to build another interface.**
  Focus on well-defined tools, skills and an MCP layer; existing chat/agent interfaces can provide the user experience.

- **One central graph does not have to serve every question.**
  Use **task-specific subgraphs and graph projections** matched to the task, modality, species or analytical problem.

- **Neo4j is the core backbone, not the boundary.**
  A lot of analytics can run directly on the graph, but **heavy ML workflows, specialized graph algorithms or integration with tabular data** may be better handled externally and connected back to the graph.
  Conduct analysis out of core -> integrate results into Neo4j structure.

- **BioCypher and BioLink**
  Helpful for reproducible ingestion and schema/ontology alignment

- **Graph size and schema matter operationally.**
  Very large graphs can make interactive queries and agent workflows noticeably slower; **purpose-scoped projections can improve both relevance and response time**.

What is your lesson?
What would you bring to a KG system?

The Healthcare business of Merck KGaA, Darmstadt, Germany operates as EMD Serono in the U.S. and Canada.

Thank You

Acknowledgements:

EMD Serono RUNI CB Team
Nick, Wai-Kit, Carmina, Peter, Ahmet, ...

Qiagen DKB Team
Venkat, Andreas, Ovidiu, Ivana, …

Neo4j
Alex, Stuart, Konrad, …

Ardigen
Maciej, Przemek, Ada, Marek, Dawid, ...

