

# GrapAL: Connecting the Dots in Scientific Literature

Christine Betts<sup>♣\*</sup>, Joanna Power<sup>♡</sup>, Waleed Ammar<sup>♡</sup>

<sup>♣</sup>Paul G. Allen School of Computer Science & Engineering,  
University of Washington, Seattle, WA, USA

<sup>♡</sup>Allen Institute for Artificial Intelligence, Seattle, WA, USA

chrstn@cs.washington.edu, {joannap,waleeda}@allenai.org

## Abstract

We introduce GrapAL (Graph database of Academic Literature), a versatile tool for exploring and investigating a knowledge base of scientific literature that was semi-automatically constructed using NLP methods. GrapAL fills many informational needs expressed by researchers. At the core of GrapAL is a Neo4j graph database with an intuitive schema and a simple query language. In this paper, we describe the basic elements of GrapAL, how to use it, and several use cases such as finding experts on a given topic for peer reviewing, discovering indirect connections between biomedical entities, and computing citation-based metrics. We open source the demo code to help other researchers develop applications that build on GrapAL.<sup>1</sup>

## 1 Introduction

Researchers rely on scientific literature to perform a wide variety of tasks such as searching for papers, assessing applicants for a research position and keeping track of papers published on topics of interest. Several software tools are available to help researchers perform these tasks. For example, many biomedical researchers use PubMed to find papers relevant for their studies,<sup>2</sup> Google Scholar allows researchers to verify and curate their user profiles,<sup>3</sup> and Semantic Scholar extracts research topics, figures, and tables from papers and links them to external content such as slides, videos and GitHub repositories.<sup>4</sup> However, such tools tend to only feature the most commonly used functionalities in order to keep the interface simple for users,

\*Work done while at the Allen Institute for Artificial Intelligence.

<sup>1</sup><https://github.com/allenai/grapal-website>

<sup>2</sup><https://www.ncbi.nlm.nih.gov/pubmed/>

<sup>3</sup><https://scholar.google.com/>

<sup>4</sup><https://www.semanticscholar.org/>

ignoring the long tail of informational needs such as finding experts on a given topic, identifying potential collaborators, assessing influence between research areas, and discovering connections between biological entities.

In this paper, we address these limitations by introducing a tool that provides a flexible and efficient way to query the Semantic Scholar knowledge base, a semi-automatically constructed knowledge base of scientific literature (Ammar et al., 2018). In addition to bridging the gap between available tools and informational needs of researchers, GrapAL demonstrates how semi-automatically constructed knowledge bases can be effectively used to solve real-world problems.

GrapAL is publicly available at [grapal.allenai.org](http://grapal.allenai.org), along with documentation.<sup>5</sup> In the following section (§2), we introduce the schema and query language used in GrapAL and discuss how users can connect to the database. In §3, we show how GrapAL can be used to satisfy several compelling case studies. In §4, we discuss some of the design choices and the system architecture for GrapAL.

## 2 How to Use GrapAL

GrapAL is designed to satisfy many use cases requested by Semantic Scholar users who need to process scientific literature. To achieve this, we design GrapAL as a Neo4j property graph with an intuitive schema, making it queryable with the Cypher query language (Francis et al., 2018).

**Schema.** Fig. 1 demonstrates the schema of our graph database, which consists of 7 node types (displayed in turquoise) and 8 edge types (displayed in purple). The properties associated with

<sup>5</sup>A screencast of the tool is available at <https://www.youtube.com/watch?v=1ivX9sHw2RU&feature=youtu.be>




Figure 1: Overview of GrapAL schema. \*denotes indexed property.

Node Type	Count	Edge Type	Count
Affiliation*	16M	AFFILIATED_WITH	119M
Author	17M	APPEARS_IN	67M
Entity	493K	AUTHORS	148M
Paper	46M	CITES	693M
Relation	51	MENTIONS	400M
RelationInstance	347K	MENTIONS_RELATION	73M
Venue*	78K	WITH_ENTITY	1M
		WITH_RELATIONSHIP	350K

Table 1: Approximate node and edge cardinalities. (\*) indicates node types that are not canonicalized.

each node and edge type are listed. In order to avoid violating intellectual property of publishers, we do not include some information about papers such as the abstract and full text.

At the core of the graph is the `Paper` node. `Paper` nodes may connect to `Venue` nodes, `Author` nodes, `Affiliation` nodes, `Entity` nodes, `RelationInstance` nodes or other `Paper` nodes via `APPEARS_IN` edges, `AUTHORS` edges, `AFFILIATED_WITH` edges, `MENTIONS` edges, `MENTIONS_RELATION` edges and `CITES` edges, respectively. A `RelationInstance` node, e.g., `CAUSES[SMOKING,CANCER]`, represents an n-ary relationship of type `Relation` (via a `WITH_RELATIONSHIP` edge) between two or more `Entity` nodes (via `WITH_ENTITY` edges). Details on how we extract entities and various metadata for each paper can be found in [Ammar et al. \(2018\)](#). The only schema changes introduced in this work are including `Affiliation` and `Venue` nodes (and corresponding edge types), and optimizing for query execution time. Table 1 provides the number of instances of each node and edge type in the schema at the time of this writing.

**Query Language.** Before we discuss realistic case studies in §3, we introduce the query lan-

guage used in GrapAL with a few toy examples:

First, consider the following query that matches arbitrary author nodes in GrapAL and returns the first 10:

```
// Find arbitrary authors.
MATCH (a:Author) RETURN a LIMIT 10
```

More often than not, we only want to match nodes with some desired properties. In the next example, we only match authors with first name ‘Clarence’ and last name ‘Ellis’. Note the round brackets used to specify an instance of node type `Author`, and the curly brackets used to specify its properties.

```
// Find authors by name.
MATCH (a:Author {last: "Ellis", first:
  ↳ "Clarence"})
RETURN a
```

Alternatively, we could use a `WHERE` clause to specify the desired properties of matched nodes, as demonstrated in the following example that matches papers by their title. This example also shows how to match nodes by specifying their relation to another node, e.g., authors of a paper. Note the use of square brackets to specify edges and the arrow to specify edge direction.

```
// Find authors of a specific paper.
MATCH (a:Author)-[:AUTHORS]->(p:Paper)
WHERE p.title = "One-shot learning of
  ↳ object categories"
RETURN a
```

More information about the Cypher query language can be found in [Francis et al. \(2018\)](#).

**Connecting to GrapAL.** Users can query GrapAL in a variety of methods. First, an interactive graphical interface is available

at <https://grapal.allenai.org:7473/browser/> that is suitable for interactive exploration of GrapAL with a relatively small number of results. We demonstrate how the interactive interface could be used in a screencast.<sup>6</sup>

Users can also build web applications that leverage GrapAL through the Neo4j HTTP endpoint.<sup>7</sup> As an example, we have developed a simple web-based application at <https://grapal.allenai.org/app> that can be used to load any of the case studies described in the next section.<sup>8</sup> Users can also type in arbitrary queries, share the queries with collaborators, and download the results in JSON format.

Users can also query the graph natively in their favourite programming language using one of the Neo4j language drivers. Neo4j officially supports five languages: .NET, Java, Javascript, Go and Python, but additional drivers are available.<sup>9</sup> We provide an example of using the Python driver to compute disruption scores as described in Wu et al. (2019).<sup>10</sup>

**DOI and ArXivId Compatibility.** Users can switch between Digital Object Identifiers (DOIs) or arXiv identifiers (ArXivId) and paper IDs with the Semantic Scholar API<sup>11</sup>. For example, we can look up the paper node corresponding to the DOI 10.1038/nrn3241 by first executing the HTTP query <https://api.semanticscholar.org/v1/paper/10.1038/nrn3241> that returns a JSON object with paper ID 931d6b6ee097eab80b8f89a313c8d3a6d5443cb2. Then, we execute the Cypher query:

```
// Look up paper by ID.
MATCH (p:Paper {paper_id:
  ↪ "931d6b6ee097eab80b8f89a313c8d3
  ↪ a6d5443cb2"})
RETURN p
```

In the future, we plan to add DOI properties and ArXivId properties to the knowledge base.

<sup>6</sup><https://www.youtube.com/watch?v=1ivX9sHw2RU&feature=youtu.be>

<sup>7</sup>Documentation for the API is available at <https://neo4j.com/docs/http-api/3.5/>

<sup>8</sup>For example, the following URL will load the shortest path example: <https://grapal.allenai.org/app/?example=shortest-path>

<sup>9</sup>See <https://neo4j.com/developer/language-guides/> for the complete list of Neo4j language drivers.

<sup>10</sup><https://gist.github.com/chrstnb/088f7699930ad53e757906f4d3d6c1f5>

<sup>11</sup><http://api.semanticscholar.org/>




Figure 2: Shortest path between Swabha Swayamdipta and Regina Barzilay.

### 3 Case Studies

We interviewed computer science and biomedical researchers to better understand the kinds of questions they would like to answer from a knowledge base of scientific literature. In this section, we focus on some of the more compelling use cases that were identified in the interviews, and provide example queries to address them in GrapAL.

For each example we give a link to load the query in the query loader and the full text of the query. From the query loader, users can view or save the results of a query and also copy it to be pasted into the Neo4j browser, where users can view interactive visualizations of the query results.

**Shortest Path.** Consider a researcher **a** seeking an introduction or an endorsement to work with another researcher **b**. By finding the shortest path between the two researchers in GrapAL, researcher **a** can identify common collaborators connecting the two. The following query, for instance, matches a path connecting Swabha Swayamdipta and Regina Barzilay using authorship edges only, and returns a path that connects them via Luke Zettlemoyer who co-authored papers with both researchers (see Fig. 2).<sup>12</sup>

```
// Find shortest path between two
  ↪ researchers by name.
MATCH p=shortestPath((a:Author)-
  [:AUTHORS*0..6]-(b:Author))
WHERE a.first = "Swabha"
  AND a.last = "Swayamdipta"
  AND b.first = "Regina"
  AND b.last = "Barzilay"
RETURN p
```

In this example, we constrain the number and type of edges in the graph to a maximum of six AUTHORS edges. For authors with an ambiguous name, it may be necessary to specify the author by

<sup>12</sup>This query can be loaded and modified at <https://grapal.allenai.org/app/?example=shortest-path>

their ID, which can be found by inspecting their author page URL on Semantic Scholar:<sup>13</sup>

```
// Find shortest path between two
↪ researchers, one by author ID.
MATCH p=shortestPath((a:Author)-
[:AUTHORS*0..6]-(b:Author))
WHERE a.author_id = 2705113
AND b.first = "Regina"
AND b.last = "Barzilay"
RETURN p
```

Similar queries can be used to find colleagues who published at a given venue, or currently work at a given university or research lab.

**Finding Experts.** One of the pain points in organizing a conference is identifying reviewers who are knowledgeable about the research topics discussed in submitted papers. By querying GrapAL, members of the organizing committee will be able to find more competent reviewers, while relying less on their (often biased) professional network when deciding whom to invite for peer reviewing. For example, the following query can be used to find researchers who published the most on “Relationship extraction” since 2013.<sup>14</sup>

```
// Find authors who published the most
↪ on relation extraction since 2013.
MATCH (a:Author)-[:AUTHORS]->(p:Paper),
(p)-[:MENTIONS]->
(:Entity {name: "Relationship
extraction"})
WHERE p.year > 2013
WITH a, count(p) as cp
RETURN a, cp
ORDER BY cp DESC
```

Here, we use ORDER BY cp DESC to sort the authors by the number of papers they published on this topic. In order to find the node that represents a topic of interest in GrapAL, users could use the search feature on semantic scholar and inspect the relevant topic page URL for the entity ID, or use regular expressions to query GrapAL, e.g.,<sup>15</sup>

```
// Fuzzy matching of entity names.
MATCH (e:Entity)
WHERE e.name =~ "(?i)relationship
extraction"
RETURN e
```

<sup>13</sup>E.g., Swabha Swayamdipta’s author page URL is <https://www.semanticscholar.org/author/Swabha-Swayamdipta/2705113>

<sup>14</sup>This query can be loaded and modified at <https://grapal.allenai.org/app/?example=experts>

<sup>15</sup>This query can be loaded and modified at <https://grapal.allenai.org/app/?example=canonical-entity>



Figure 3: Ten papers that mention both ‘Natural language processing’ and ‘Constraint programming.’.

**Papers at the Intersection of Entities.** Search engine results sometimes make it difficult to find papers that discuss multiple topics or fields. With GrapAL, we can return papers that discuss any number of entities of interest, e.g., “Constraint programming” and “Natural language processing”. Fig. 3 shows a visualization of the results on the Neo4j browser, limited to 10 papers.<sup>16</sup>

```
// Find papers which mention both
↪ constraint programming and natural
↪ language processing.
MATCH (p:Paper)-[:MENTIONS]->
(e1:Entity {name: "Constraint
programming"}),
(p:Paper)-[:MENTIONS]->
(e2:Entity {name: "Natural language
processing"})
RETURN p
```

**Connecting Scientific Concepts.** Some researchers wanted to explore direct and indirect connections between two scientific concepts (entities) of interest, e.g., the impact of ‘adjuvant antiestrogen therapy (Arimidex)’ on ‘estrogen receptors’. Using GrapAL, we can find how two entities are indirectly connected via coded relationships and a chain of entities in the knowledge base, which can help generate new hypotheses or quickly assess the viability of a hypothesis before conducting expensive lab experiments.<sup>17</sup>

<sup>16</sup>This query can be loaded and modified at <https://grapal.allenai.org/app/?example=intersecting-entities>

<sup>17</sup>This query can be loaded and modified at <https://grapal.allenai.org/app/?example=scientific-concepts>

```

// Find path between Estrogen Receptors
↪ and Arimidex via coded
↪ relationships.
MATCH path=shortestPath(
  (er:Entity {name: "Estrogen
    Receptors"})-
  [:WITH_ENTITY*0..15]-
  (ar:Entity {name: "Arimidex"}))
WITH nodes(path) as n
UNWIND ns as n
MATCH (n)-[:WITH_ENTITY {position: 0}]->
  (e0:Entity),
  (n)-[:WITH_ENTITY {position: 1}]->
  (e1:Entity),
  (n)-[:WITH_RELATIONSHIP]->
  (r:Relation)
RETURN e0, r, e1

```

This query returns a list of triples (e0, r, e1) that connect ‘Arimidex’ to ‘Estrogen Receptors’. The UNWIND operator allows us to examine each node on the shortest path and process it as needed.

**Citation-Based Metrics.** Citations are often used as a proxy for the impact of papers, researchers or venues. In addition to computing traditional metrics such as h-index and i10-index, GrapAL can also be used to compute more granular metrics, e.g., to estimate the rate at which papers in one conference cite papers in another conference:<sup>18</sup>

```

// Find the number of times a NAACL
↪ paper cites a CVPR paper.
MATCH (p1:Paper)-[:APPEARS_IN]->
  (naacl:Venue),
  (p2:Paper)-[:APPEARS_IN]->(cvpr:Venue),
  path=(p1)-[:CITES]->(p2))
WHERE naacl.text =~ ".*NAACL.*"
AND cvpr.text =~ ".*CVPR.*"
RETURN count(path)

```

This query returns the number of times a NAACL paper cites a CVPR paper. We use the =~ operator to match on venue names by regular expression because venues are stored as unstructured strings.

## 4 System Design

**Graph Database.** Due to the high connectivity in the data and the nature of queries GrapAL is designed for, we opted to create GrapAL using a graph-native database instead of a more conventional relational database. Unlike a relational database, a graph database provides a natural and efficient way to query and traverse multi-hop relations without using computationally expensive join operations. Several graph database systems

<sup>18</sup>This query can be loaded and modified at <https://grapal.allenai.org/app/?example=citation-metrics>

have recently become available, including AWS Neptune, Grakn.ai, dgraph and Neo4j. We decided to build GrapAL on Neo4j since it is one of the more mature platforms, has a strong community of developers, and is the most widely used graph database system as of the time of this writing.<sup>19</sup> One limitation of Neo4j is that it is not a distributed database system, but we were able to fit GrapAL on a single server.

**Building and Deploying GrapAL.** GrapAL is powered by the same data that powers the [semanticscholar.org](https://semanticscholar.org) website, as described in [Ammar et al. \(2018\)](#). We use a *staging server* to read a snapshot of the data as Spark DataFrames from AWS S3 and write CSV files that match the property schema described earlier. Due to the sheer amount of records, we process different shards of the data in parallel before aggregating all shards into one CSV file for each node and edge type of the schema. Then, we use the Neo4j CSV import function to build the database. Once we’ve built the database, we start up a Neo4j server and run a Cypher script to create indexes. The staging server is an EC2 machine with instance type r5.24xlarge. This process takes around 6 hours and the resulting database is roughly 80 GB (including indexes).

Once the data is imported, the database files are copied over to a *production server* that serves the dataset publicly and has lower processor and memory requirements compared to the staging server. The staging server is an EC2 machine with instance type r4.16xlarge. We plan to rebuild GrapAL at a monthly cadence with new snapshots of the data.

## 5 Related Work

Related APIs are available to help researchers navigate scientific literature. [Singh et al. \(2018\)](#) provides an API to interact with the ACL anthology. However, it is limited to the areas of computational linguistics and natural language processing, and it uses a predefined list of query templates with placeholders for authors, papers and venues. Springer Nature SciGraph<sup>20</sup> provides an API for accessing publication metadata from the Springer Nature corpus, but it is limited to papers and books

<sup>19</sup><https://db-engines.com/en/ranking/graph+dbms>

<sup>20</sup><https://scigraph.springernature.com/explorer>

published by Springer Nature. The Microsoft Academic Graph (Shen et al., 2018) is similarly an API for examining academic literature. As a relational database, it is hard to query with complex, multi-hop relations as discussed in §4. This work is also related to a line of NLP work focusing on scientific documents including citation prediction (e.g., Yogatama et al., 2011; Bhagavatula et al., 2018), author modeling (e.g., Sim et al., 2015), stylometry (e.g., Bergsma et al., 2012), bibliometrics (e.g., Foulds and Smyth, 2013; Weihs and Etzioni, 2017) and information extraction (e.g., Kergosien et al., 2018; Andruszkiewicz and Hazan, 2018).

## 6 Conclusion

GrapAL is a versatile tool for exploring and investigating scientific literature built on the Neo4j graph database framework. We describe the basic elements of GrapAL, how to use it, and use cases such as finding experts on a given topic for peer reviewing, discovering indirect connections between biomedical entities, and computing citation-based metrics.

Future improvements include more metadata and changes to the structure of affiliation and venue data. We intend to change the data pipeline architecture to perform event-based incremental updates rather than a regular batch build. We continue to improve the models used to populate GrapAL’s nodes and edges (e.g., author disambiguation and entity extraction and linking).

## Acknowledgments

We thank Khaled Ammar for his graph database suggestions, Michal Guerquin for his help in designing and building the pipeline, and Darrell Plessas for his technical assistance. We also thank Noah Smith and the Semantic Scholar team for their support.

## References

Waleed Ammar, Dirk Groeneveld, Chandra Bhagavatula, Iz Beltagy, Miles Crawford, Doug Downey, Jason Dunkelberger, Ahmed Elgohary, Sergey Feldman, Vu Ha, Rodney Kinney, Sebastian Kohlmeier, Kyle Lo, Tyler Murray, Hsu-Han Ooi, Matthew E. Peters, Joanna Power, Sam Skjonsberg, Lucy Lu Wang, Chris Wilhelm, Zheng Yuan, Madeleine van Zuylen, and Oren Etzioni. 2018. Construction of

the literature graph in semantic scholar. In *Proc. of NAACL*.

Piotr Andruszkiewicz and Rafal Hazan. 2018. [Annotated corpus of scientific conference’s homepages for information extraction](#). In *Proc. of LREC*.

Shane Bergsma, Matt Post, and David Yarowsky. 2012. [Stylometric analysis of scientific articles](#). In *Proc. of NAACL*.

Chandra Bhagavatula, Sergey Feldman, Russell Power, and Waleed Ammar. 2018. Content-based citation recommendation. In *Proc. of NAACL*.

James Foulds and Padhraic Smyth. 2013. [Modeling scientific impact with topical influence regression](#). In *Proc. of EMNLP*.

Nadime Francis, Alastair Green, Paolo Guagliardo, Leonid Libkin, Tobias Lindaaker, Victor Marsault, Stefan Plantikow, Mats Rydberg, Petra Selmer, and Andrés Taylor. 2018. [Cypher: An evolving query language for property graphs](#). In *Proc. of SIGMOD*.

Eric Kergosien, Amin Farvardin, Maguelonne Teisseire, Marie-Noëlle Bessagnet, Joachim Schöpfel, Stéphane Chaudiron, Bernard Jacquemin, Annig Lacayrelle, Mathieu Roche, Christian Sallaberry, and Jean-Philippe Tonneau. 2018. [Automatic identification of research fields in scientific papers](#). In *Proc. of LREC*.

Zhihong Shen, Hao Ma, and Kuansan Wang. 2018. [A web-scale system for scientific knowledge exploration](#). In *Proc. of ACL*.

Yanchuan Sim, Bryan Routledge, and Noah A. Smith. 2015. [A utility model of authors in the scientific community](#). In *Proc. of EMNLP*.

Mayank Singh, Pradeep Dogga, Sohan Patro, Dhiraaj Barnwal, Ritam Dutt, Rajarshi Haldar, Pawan Goyal, and Animesh Mukherjee. 2018. [Cl scholar: The acl anthology knowledge graph miner](#). In *Proc. of NAACL*.

Luca Weihs and Oren Etzioni. 2017. [Learning to predict citation-based impact measures](#). In *Proc. of ACM/IEEE Joint Conference on Digital Libraries*, pages 49–58. IEEE Press.

Lingfei Wu, Dashun Wang, and James A. Evans. 2019. [Large teams develop and small teams disrupt science and technology](#). *Nature*, 566:378–382.

Dani Yogatama, Michael Heilman, Brendan O’Connor, Chris Dyer, Bryan R. Routledge, and Noah A. Smith. 2011. [Predicting a scientific community’s response to an article](#). In *Proc. of EMNLP*.