



## Strategies for creating knowledge graphs to depict a multi-perspective Queer communities representation

Louann Coste, Flora Helmers, Hamamache Kheddouci, Léo Le Nestour,  
Mahsa Niazi, Genoveva Vargas-Solar

### ► To cite this version:

Louann Coste, Flora Helmers, Hamamache Kheddouci, Léo Le Nestour, Mahsa Niazi, et al.. Strategies for creating knowledge graphs to depict a multi-perspective Queer communities representation. Workshops of the EDBT/ICDT 2023 Joint Conference, Mar 2023, Ioannina, Greece. hal-04272088

**HAL Id: hal-04272088**

**<https://hal.science/hal-04272088>**

Submitted on 6 Nov 2023

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# Strategies for creating knowledge graphs to depict a multi-perspective Queer communities representation

Louann Coste<sup>1</sup>, Flora Helmers<sup>1</sup>, Hammamache Kheddouci<sup>2</sup>, Léo Le Nestour<sup>1</sup>, Mahsa Niazi<sup>1</sup> and Genoveva Vargas-Solar<sup>2</sup>

<sup>1</sup>Ecole Normale Supérieure de Lyon, 15 parvis René Descartes, 69342, Lyon

<sup>2</sup>CNRS, Univ Lyon, INSA Lyon, UCBL, LIRIS, UMR5205, F-69221, France

## Abstract

This paper introduces an experimental study for building knowledge graphs about queer communities with different perspectives. The paper describes a set of pipelines implementing several knowledge graphs construction strategies that lead to a complementary understanding of the queer communities. We implemented a library with the knowledge graph construction pipelines and conducted experiments to evaluate execution cost against the comprehensiveness of the content. The library includes visualization tools for observing the knowledge graphs concerning several granularities. Finally, we use storytelling to provide a first interpretation of our results regarding the profiling of queer communities.

## Keywords

Data science pipelines, graph analytics, knowledge graphs, queer history

## 1. Introduction

Heterogeneous datasets can be structured into networks. The notion of "graph" is a powerful mathematical concept for representing these networks as graphs. Graphs can be exploited through workflows applying algorithms to solve data science problems. When the graphs become too large, the processes used to explore, and analyse them become costly. Therefore, deployment strategies on target architectures can fulfil the resources' consumption requirements in an adaptable way [1].

This paper introduces an experimental approach for building knowledge graphs to reveal several perspectives of Queer History. Our work models the pipelines to explore the Wiki data knowledge graph and extract sub-graphs about queer identities combining different variables like age, artistic production, geographical position, etc. Extracted graphs are studied, classified, and profiled thoroughly to estimate the allocation of resources. The pipelines include performing unions, merge, and community discovery to build several perspectives of Queer History.

The remainder of the paper is organised as follows. Section 2 introduces work addressing the construction and

exploitation of knowledge graphs. Section 3 introduces our approach based on different pipelines for building queer history knowledge graphs. Section 4 reports our experiments. Finally, Section 5 concludes the paper and discusses future work.

## 2. Related Work

*Knowledge graphs construction and exploration.* A knowledge graph is a directed graph  $G = (N, R)$  whose nodes  $N$  represent entities and literal values (literals), and whose edges  $R$  represent relations between these entities. The graph can be divided into two parts: the data and the knowledge. The knowledge can come from several sources that describe entities (i.e., concepts) and from relations among entities that can be explicitly stated or discovered [2]. Knowledge graphs defined using the RDF model can be queried using the SPARQL query language.

Knowledge graphs are used to integrate bibliographical data and exhibit references and papers co-authorship. Examples of solutions are DIG [3], Knowledge Vault [4], CiteSeerX [5], Pujara et al. [6], the Knowledge Graph Identification [7] and NELL [8]. The authors of [9] propose a framework for incrementally building a knowledge graph of artists used to link the data from different museums. These approaches rely on a predefined schema describing the data used to build the graph. Other approaches, adopt the opposite strategy and do not use a schema for building knowledge graphs [10, 11, 12]. Knowledge graphs can also be built by consolidating data from different sources like LDIF [13]), YAGO and YAGO2 [14, 15] and Freebase [16].

*Modelling the Queer History with graphs.* Queer history was not recorded in a structured way for a long time. The

*Published in the Workshop Proceedings of the EDBT/ICDT 2023 Joint Conference (March 28-March 31, 2023, Ioannina, Greece).*

\*Genoveva Vargas-Solar.

† Authors contributed equally and are enumerated in alphabetical order.

✉ louann.coste@ens-lyon.fr (L. Coste); flora.helmerts@ens-lyon.fr (F. Helmers); hamamache.kheddouci@univ-lyon1.fr (H. Kheddouci); leo.le\protect1\_nestour@ens-lyon.fr (L. L. Nestour); mahsa.niazi@ens-lyon.fr (M. Niazi); genoveva.vargas-solar@cnrs.fr (G. Vargas-Solar)

© 2023 Copyright © 2023 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).



CEUR Workshop Proceedings (CEUR-WS.org)

most prominent project willing to integrate the history of LGBTQ+ communities is the WikiProject LGBT<sup>1</sup> for developing LGBT-themed content in Wikidata and advocating for the community of contributors developing this content. Wikidata contains self-organized archives starting in the 1970s. In Munich the Forum Queeres Archiv, has been collecting, researching, and publishing for 20 years: Shelf meters full of bequests, files, books, journals, objects. The project Queer data<sup>2</sup> transforms analog queer history into linked open data and incorporates it into an open, freely accessible, public data infrastructure. The knowledge graph contains 7415 entities from the data sources of the Forum Queeres Archiv München, with more than 38,000 connections.

*Discussion* The challenge to modelling Queer history is to use different properties to define identities that can be profoundly diverse and relations and provide a multi-perspective model. Exhibiting these perspectives can propose a more representative queer history knowledge graph. This paper shows different perspectives of queer communities by building and integrating graphs.

### 3. Building Knowledge Graphs Pipelines

The main contribution of our work is the design of three pipelines to build knowledge graphs using different techniques. The purpose is to propose a pipeline that:

- Balances the coverage of the resulting knowledge graph that is its scope concerning the queer communities (i.e., are all potential queer individuals part of the graph, are their connections able to exhibit different communities).
- Reasonably consumes resources (CPU and memory) to explore, analyse and build dense graphs. The second contribution is the implementation of these pipelines within a library<sup>3</sup> that allows testing them by calibrating specific parameters such as the number of iterations, the number of nodes to visit, etc. We also provide a ready-to-use notebook to run the experiments we have conducted (see Section 4).

#### Pipeline 1: Pure SPARQL based knowledge graph

Figure 1 illustrates the pipeline that builds a pure SPARQL-based knowledge graph. We propose the construction of two initial graphs by querying Wikidata to retrieve individuals identified as non-cisgender and non-heterosexual and their connections.

The Wikidata service offers the WikiData Query Service endpoint, which allows users to run SPARQL queries on the WikiData database. SPARQL allows creating a graph with the CONSTRUCT keyword, and the resulting graph can be downloaded as a .csv file

with subject, predicate, and object columns. This .csv file can be transformed using the Pandas library into a DataFrame or a networkx Graph using the `networkx.from_pandas_dataframe` function. The interest in producing a networkx graph is that the library provides additional methods for graph analysis.

In the example below, a graph is constructed with all triples where the subject is a non-heterosexual or non-cisgender individual. This activity corresponds to the data retrieval phases of the pipeline by interacting with the Wikidata endpoint server (see Figure 1).

```
CONSTRUCT {
  ?person ?predicate ?object .
}
WHERE {
  {
    ?person wdt:P31 wd:Q5 . #?
      personId is a human
    ?person ?predicate ?object .
  }
  {
    ?person wdt:P21 ?sexorgender .
      #?person has ?
        sexorgender
    #?sexorgender is not male ,
      female , cisgender male ,
      cisgender female , or
      cisgender person
    FILTER(?sexorgender NOT IN (
      wd:Q6581097 , wd:Q6581072 ,
      wd:Q15145778 , wd:
        Q15145779 , wd:Q1093205)) .
  } UNION {
    ?person wdt:P91 ?
      sexualorientation . #?
      person has ?
        sexualorientation
    FILTER(?sexualorientation !=
      wd:Q1035954) . #?
      sexualorientation is not
        heterosexual
  }
}
SERVICE wikibase:label { bd:
  serviceParam wikibase:language "[
    AUTO_LANGUAGE]". }
```

We create a second graph where the person is not the subject but the object. With these two graph views, we have all the connections to the elements. We retrieve the labels of the entity using the service Wikibase. Retrieving the whole graph is costly (the evaluation of this query doubles the time of the first one), so we set a timeout and limit the size of the graph.

We note that we have to clean the data because they contain duplicates and "empty" nodes. Indeed, some elements have different Internationalized Resource

<sup>1</sup>[https://www.wikidata.org/wiki/Wikidata:WikiProject\\_LGBT](https://www.wikidata.org/wiki/Wikidata:WikiProject_LGBT)

<sup>2</sup><https://queerdata.forummuenchen.org/en/>

<sup>3</sup><https://github.com/FloraHelmets/QueerHistoryProject>

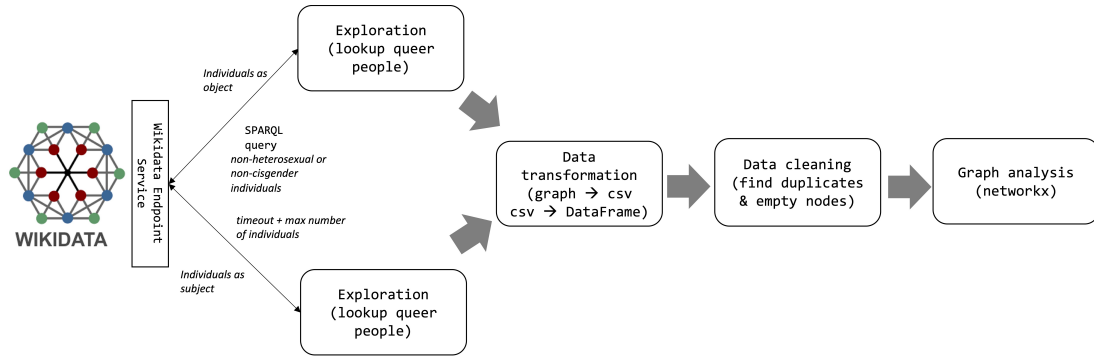


Figure 1: Pure SPARQL pipeline

Identifiers (IRIs) but have the same meaning: for instance, <http://www.wikidata.org/prop/direct/P161> and <http://www.wikidata.org/prop/statement/P161> gives access to the same property which is "cast member". However, since they have different terms, they are expressed as if they were different. We also get empty nodes which do not have labels, such as entities Q19151093<sup>4</sup> and Q19218452<sup>5</sup>. Both have different IRI and Wikipedia Identifiers but have potentially the same meaning because their properties are identical. Because we only get access to the elements located at a distance 1 from the people, we cannot correct these elements only using what we have gathered. Besides, we lose the added value of the knowledge graph, which is to automatically deduce information by using the varied distant relationships of the initial node. The principle of pipeline 1 is sound, but the execution cost is high according to the resources we have access to. Therefore we propose alternative pipelines described in the following sections.

**Pipeline 2: Merging Stars.** The construction pipeline of the knowledge graph about queer communities from Wikidata using merging stars can be summarised as follows: merge a list of graphs retrieved from Wikidata as .nt files that are akin to stars around a central entity node (see Figure 2). As shown in Figure 2, this pipeline starts with a SPARQL query to Wikidata that returns a list of Wikidata item IDs that are related to the queer community.

This list, originally in JSON, is then used to extract the relevant information necessary to query the graphs associated with each entity on Wikidata (IDs and URLs). We iteratively parse new information by using these graphs, basically creating a merged RDF graph using the RDF data of the nodes from the Wikidata entity URLs in an n-triples format (given in an .nt file format).

The merging of two knowledge graphs in .nt file format

is a process in which the data from two separate RDF graphs is combined into a single graph. The process is straightforward, as both graphs are in a standard format that can be easily integrated. The only subtleties that arise when merging two RDF graphs are related to the processing of shared blank nodes.

Shared blank nodes are nodes that appear in both RDF graphs and are used to represent anonymous resources. These shared blank nodes must be merged in an integrated RDF graph to represent the same resource. The rules for merging shared blank nodes are defined in the W3C RDF 1.1 Semantics and Abstract Syntax specification<sup>6</sup>. The merging process must ensure that the integrated graph preserves the intended meaning of the original RDF graphs, including the relationships between the nodes in the graph. Therefore, the process must consider the context in which each shared blank node is used in both graphs and combine them to preserve the RDF data's intended meaning. Our final integrated graph is then converted to a NetworkX multidigraph that proved to be the most helpful format for our analysis.

Once the graph has been created, it is necessary to prune it in various ways:

- (1) Removing "dead-end" nodes (i.e., nodes that have an in-degree of less than 2) "Dead-end" nodes are often isolated and do not provide much information about the community.
- (2) Removing duplicates (i.e., nodes that represent the same object literally) and isolated nodes (i.e., nodes that have both in-degree and out-degree of 0). These nodes are not connected to any other nodes in the graph and do not provide information about the community.

Depending on the application, one can also specify more complex pruning processes with rule-based deletion and deletion before each merge. The pipeline is executed with two parameters: n the size of the list of

<sup>4</sup><http://www.wikidata.org/entity/Q19151093>

<sup>5</sup><http://www.wikidata.org/entity/Q19218452>

<sup>6</sup><https://www.w3.org/TR/rdf11-nt/>

#shared-blank-nodes-unions-and-merges

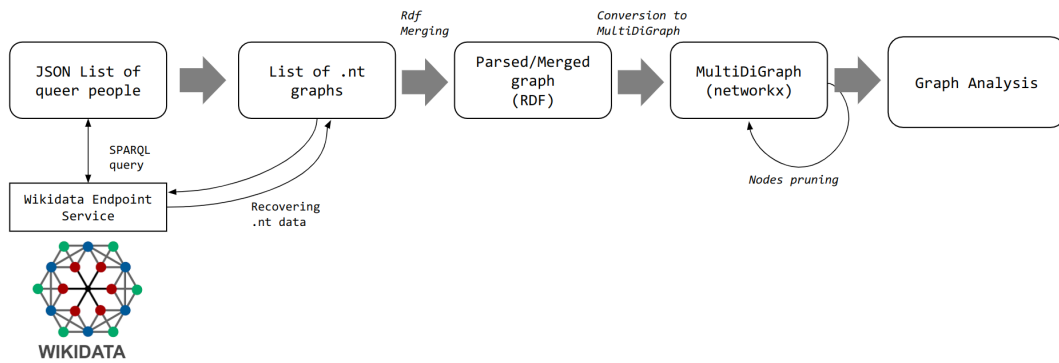


Figure 2: "Merging stars" pipeline

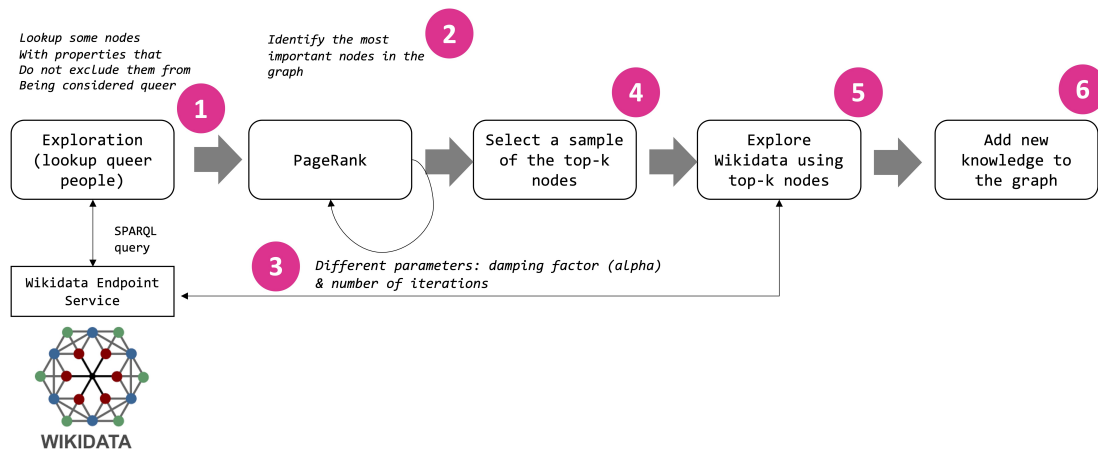


Figure 3: Crawled method pipelines

people used as starting point and a `prune_policy` dictionary describing what process for removing the nodes should be used.

**Pipeline 3: Crawler method.** The crawler-based pipeline starts with a small number of nodes. It runs an iterative process to extract critical nodes representing properties of interest (potential common points) and use those to explore and discover more queer people and communities. Figure 3 illustrates the main activities of the pipeline. The process begins by selecting a subset of nodes using a SPARQL query (see 1 in Figure 3). The query selects distinct people and their labels who have a specific property of interest and do not have other properties that would exclude them from being considered queer. A query result is a limited number of nodes, which is used as the starting point for the iterative process.

Next, the pipeline runs the PageRank algorithm on

the initial set of nodes for ranking nodes based on the number and quality of incoming links (see 2 in Figure 3). The algorithm allows the identification of the most critical nodes in the graph, as they are more likely to represent properties of interest. The PageRank algorithm is run multiple times with different parameters, such as the damping factor ( $\alpha$ ) and the number of iterations (see 3 in Figure 3).

After the PageRank algorithm has been executed, the pipeline selects a certain number of the most critical nodes ( $k_{prop}$ ) and uses them to explore further (see 4 in Figure 3). Specifically, it runs the same SPARQL query as before, using the selected nodes as the property of interest (see 5 in Figure 3). The result is a new set of nodes connected to the previously selected nodes through the property of interest. These new nodes are added to the original graph, and the process starts over again in

iterations (see 6 in Figure 3).

The crawler-based pipeline continues to run the PageRank algorithm and explore new nodes until a certain number of iterations have been completed ( $n\_iter$ ). The resulting graph is a composite of all the nodes and edges discovered during the process, and it is more densely connected and contains a higher number of queer people than the initial set of nodes.

### 3.1. Visualising graph structures

Visualising a knowledge graph is a crucial step in understanding the structure and content of the data. To better understand the structure of the knowledge graph generated in our study, we implemented custom visualisation techniques, including different graph layouts and edge bundling.

*Nodes Layout.* The force-directed layout, described in [17] minimises overlaps in the graph, evenly distributes nodes and links and organises items so that links are of a similar length with as few crossing edges as possible. This process is done by assigning forces among the set of edges and the set of nodes based on their relative positions and then using these forces to simulate the edges and nodes' motion or minimise their energy. In the case of our pipelines based on .nt "stars" associated with entities, the force-directed layout is particularly relevant to distinguish the underlying structure obtained.

*Edge-Bundling.* We use edge bundling, a method that allows edges to curve and then groups nearby ones together to help convey structure. We use the function "Datashader.hammer\_bundle" (a variant of [18]) for the edge bundling.

## 4. Experiments and Results

*Implementation.* We have developed a library for handling and building knowledge graphs with real-world data for a comfortable experimental setting. This library, written in Python, uses the NetworkX, RDFLib and Datashader libraries and allows to evaluate the performance of different knowledge graph construction methods.

*Queer datasets.* In the conducted experiments, we use data from Wikidata, as this platform provides structured data that is well-suited for building knowledge graphs<sup>7</sup>). We used "real" data to make our simulation experiments realistic. For example, we plotted the histogram of the number of properties each node from the SPARQL query had and found that the distribution resembled a log-normal distribution. We fitted the data to a log-normal distribution and used the Kolmogorov-Smirnov test to evaluate the fitting. The test showed that the fitting

was appropriate, which allowed us to use the log-normal distribution to model the distribution of the number of out edges in our simulation experiments. This strategy allowed us to have a more accurate simulation of the real-world construction of knowledge graphs and to assess the performance of different construction methods more accurately.

Because not all properties constitute common points among people, we only take a fraction of these number of out-edges to have our distribution for the models. By running tests on Wikidata, we found that approximately 1/50 of these out-edges created an entity-linking edge. This gives us the distribution with a probability density function (PDF):

$$\frac{1}{x\sigma\sqrt{2\pi}} \exp - \frac{(\ln(x) - \mu)^2}{2\sigma^2} \quad (1)$$

with  $\mu \approx 0.495$  and  $\sigma \approx 5$ .

*Understanding the queer concepts.* We defined an ontology to organise the critical elements of the Queer Community. We used preexisting standard ontologies such as foaf to describe the people and Geonames to make connections between spaces. Our ontology focuses on the European Queer scene of the 1980s. We added the correlation between gender and sex without having it explicitly said. We based our strategy on the idea of Katharina Brunner in the "Remove NA" project<sup>8</sup>. We used pyramidal relations to describe sexual orientations. Our ontology also focuses on art and collaborations, including the notion of avatars which is essential in the drag culture. The concept of influence is also present, considering elements such as songs played in clubs, characters in films or books, and access to art pieces. This version focuses on the representation of LGBTQ+ people in the culture of the 1980s. Still, it could benefit from exploring the rights of LGBTQ+ individuals from the past from a more global and non-Eurocentric perspective. One possible approach would be to examine the various terms used to define gender across different cultures, such as the third gender in the Bugis society, and to include references to the laws and changes in each country. The first we can get using the Wiki Data Query Service.

*Experiments objective.* With the experiments we aimed to directly compare the performance of the three pipelines in terms of time efficiency, size, structure, and content. To accomplish this, we measured the time each pipeline takes to create the knowledge graph, the number of nodes, and the number of edges in each built graph. Furthermore, we analysed the maximal graphs that each method/pipeline can create and compared them.

Note that the experiments were performed under the constraints of the regular account on a Google free service. We used the Google Colab environment to run our

<sup>7</sup><https://ai.stanford.edu/blog/introduction-to-knowledge-graphs/>

<sup>8</sup><https://queerdata.forummuenchen.org/en/>



experiments with an allocated virtual machine with CPU (1xsingle core hyperthreaded Xeon Processors @2.3Ghz (i.e., 1 core, 2 threads); 12,7 GB of RAM (with 0.8 GB are already taken); a storage space of 108 GB, of which only 77 GB is available to the user. Thus, the results may differ if performed under different conditions. However, this experimental setting was chosen to simulate a realistic scenario for a typical user.

#### 4.1. Structure and communities


In our experiments, we found that the structure of the created knowledge graphs varied depending on the construction method used and the distribution used. Some methods resulted in graphs with a high degree of connectivity and a dense network of edges, like the Pure SPARQL pipeline, while others produced sparser graphs with fewer connections between nodes. We also observed that some methods tended to create graphs with a more hierarchical structure, especially in trials with the crawler-based pipeline. In contrast, others generated graphs with a more flat structure.

To better understand the structure of the graphs produced through our pipelines, we used community detection algorithms to identify and analyse the communities within the graphs. In particular, we used the modularity maximisation algorithm and the Girvan-Newman algorithm [19] to identify communities based on the density of connections within communities and the sparsity of relations between communities. Considering the size of the graphs, the community algorithms were run on either small portions of the graphs or simplified versions where we linked entities according to their properties in common and then removed the redundant nodes. We could then create several dendrograms to be interpreted and used for statistical analysis. Crucially, we found that the identified communities often corresponded to semantically meaningful groups of nodes, such as nodes representing people from the same country or nodes representing concepts in the same domain.


#### 4.2. Results

The initial "raw" knowledge graph without edges built from Wikidata consists of 8171 references to queer people and 4459 predicates, and 32651 objects. The graph with "in" edges, where the LGBTQ+ people are the object of the RDF triples, contains 7281 queer people and 7281 nodes with "in" edges. We get 521 different predicates and 168 923 subjects. In the predicates, when comparing the IRI, there are interesting distinctions.

*Structural comparison of the resulting graphs.* Visual analysis on dense graphs with an actual number of nodes can be laborious, hence the need for some graphical tools. We observe structural differences on the graph generated,



**Figure 4:** Graph generated by Pure SPARQL, after removing dead ends




**Figure 5:** Graph generated by Merging stars

with examples given below in the figures 4, 5 and 6. As stated in Section 3, the pure SPARQL pipeline (pipeline 1) had to include data-cleaning activities to understand the queer communities better. Still, as shown in Figure 4 the raw graph built from exploring Wikidata with SPARQL queries without dead ends is very dense.

Pipeline 2 adopting the merging stars leads to a clean and representative view of the queer communities (in an integrated vision). As described in Section 3 in the merging stars pipeline, we defined a raw graph and then pruned it to eliminate "dead nodes", then duplicates and isolated nodes and finally, a wholly cleaned knowledge graph, thoroughly pruned, as shown in Figure 5.

The graph produced through the crawler method pipeline exploits the notion of centrality to identify the nodes representing influential queer people and then explore their connections to determine whether they belong not to a queer community. The resulting graph is shown in Figure 6. The graph exhibits highly connected nodes and communities of nodes agglutinated around.

Overall, our experiments suggest that the structure



**Figure 6:** Graph generated by the crawler method pipeline

of the created knowledge graphs is influenced by the construction method used and that different methods may be more suitable for various applications depending on the desired configuration of the resulting graph, with the crawler-based method offering the most freedom in choosing the behaviour of the exploration and allowing to obtain the most salient communities. For example, the behaviour of exploration through common points allowed us to discover key communities naturally, e.g. [Q160456, Q239533, Q461657, Q2041541, Q3611679, Q6312325, Q104358, Q187814, Q346309, Q93349, Q6897510, Q463319, Q267951, Q2849427, Q7933271, Q259507, Q461758, Q3942181] corresponding to a group of Black American queer Women. Additionally, it suggests that community detection algorithms can provide valuable insights into the created graphs' structure and organisation and help identify semantically meaningful groups of nodes within the graph.

*Performance evaluation of the pipelines.* To evaluate the performance of the three pipelines, we compare the time taken, the number of nodes, and the number of edges in the knowledge graph created by each pipeline for a fixed number of people ( $n = 30$ ). Hereafter, Pipelines 1, 2 and 3 represent, respectively, the Pure SPARQL, Merging stars and Crawler pipelines and Table 1 shows the results.

Pipeline	Time	Nodes	Edges
Pipeline 1	1223 ms	6659	7217
Pipeline 2	36.3 s	8956	19172
Pipeline 3	486 s	10321	20112

**Table 1**

Comparison of the variables in the knowledge graph created by each pipeline for  $n = 50$ .

Additionally, we compare the maximal knowledge graph that each pipeline can create, where we do not impose any limit and try to get as many people as possi-

ble included. Table 2 shows the results.

Pipeline	Size	Time (s)
Pipeline 1	887 048 triples	38 294 ms
Pipeline 2	1 038 197 triples	52 min
Pipeline 3	>300 000 triples	NA

**Table 2**

Comparison of the maximal knowledge graph that each pipeline can create.

The results of the comparison between the three pipelines show that Pipeline 1 is the fastest in creating a knowledge graph. Still, Pipeline 2 has more freedom in choosing which nodes are kept and, as such, could be sped up (for example, by pruning more useless nodes before each merging) while maintaining good quality. In the same way, Pipeline 3 offers maximal freedom in creating the knowledge graph but is, at the same time, the slowest of the three.

These results suggest that while time efficiency is critical, having more precise processing of the nodes in the knowledge graph is also crucial for providing a comprehensive representation of the information that can be visualised and used more efficiently after. Pipeline 3 may be the most suitable for applications requiring a directed and precise information model. In contrast, Pipeline 1 may be ideal for applications with real-time constraints and frequent data updates.

Overall, the results highlight the trade-off between time efficiency and the comprehensiveness of the knowledge graph and suggest that the choice of a pipeline may depend on the specific requirements and constraints of the application.

*Interpretation* A queer community may consist of people who identify as lesbian, gay, bisexual, transgender. Accordingly, this can result in creating graphs with various densities and branching. Having several groups with dense connections or adding gender communities besides sexual communities, we get more sparse graphs. In this context, storytelling is a process of confronting information about the people of LGBTQ+ and discussing connections discovered in graphs which have not been made or were ignored. This process can rely on statistical analysis of the graphs obtained and data encountered when building our pipelines. Graphs brought interesting perspectives on how the existing state of the data, with the aforementioned log-normal distribution indicating a wealth of data for a selected set of individuals in the long-tail but also a vast underacknowledged majority.

## 5. Conclusions and Future Work

Through knowledge graphs, this paper introduced our experimental approach to building queer communities'



history from different perspectives. We propose and exhibit the pipelines for building the knowledge graphs, including the sets of SPARQL queries for exploring Wikidata. Our experiments have shown that we can use these models to choose the method, reduce the time it takes to build the graph, and improve the quality of the resulting graph. These findings have important implications for constructing and using knowledge graphs in various applications.

There are two future directions regarding these tasks, first is deciding which are the most suitable graph operations and analytics algorithms for answering particular research questions. The second is interpreting results. Therefore, we will apply storytelling techniques to produce dashboards that provide plots and multimedia content to propose an interpretation of the represented knowledge.

## Acknowledgments

This project has been partially funded by the project GALILEAN of the intergroup collaboration program of the LIRIS lab.

## References

- [1] G. Vargas-Solar, M. S. Hassan, A. Akoglu, JITA4DS: disaggregated execution of data science pipelines between the edge and the data centre, *J. Web Eng.* 21 (2022). URL: <https://doi.org/10.13052/jwe1540-9589.2111>. doi:10.13052/jwe1540-9589.2111.
- [2] M. Kejriwal, C. A. Knoblock, P. Szekely, *Knowledge graphs: Fundamentals, techniques, and applications*, MIT Press, 2021.
- [3] P. Szekely, C. A. Knoblock, J. Slepicka, A. Philpot, A. Singh, C. Yin, D. Kapoor, P. Natarajan, D. Marcu, K. Knight, et al., Building and using a knowledge graph to combat human trafficking, in: *International Semantic Web Conference*, Springer, 2015, pp. 205–221.
- [4] X. Dong, E. Gabrilovich, G. Heitz, W. Horn, N. Lao, K. Murphy, T. Strohmman, S. Sun, W. Zhang, Knowledge vault: A web-scale approach to probabilistic knowledge fusion, in: *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, 2014, pp. 601–610.
- [5] A. G. Ororbia II, J. Wu, C. L. Giles, CiteSeerx: Intelligent information extraction and knowledge creation from web-based data (????).
- [6] J. Pujara, L. Getoor, Building dynamic knowledge graphs, in: *NIPS Workshop on Automated Knowledge Base Construction*, volume 9, 2014.
- [7] J. Pujara, H. Miao, L. Getoor, W. Cohen, Knowledge graph identification, in: *International semantic web conference*, Springer, 2013, pp. 542–557.
- [8] A. Carlson, J. Betteridge, B. Kisiel, B. Settles, E. R. Hruschka, T. M. Mitchell, Toward an architecture for never-ending language learning, in: *Twenty-Fourth AAAI conference on artificial intelligence*, 2010.
- [9] G. Gawriljuk, A. Harth, C. A. Knoblock, P. Szekely, A scalable approach to incrementally building knowledge graphs, in: *International conference on theory and practice of digital libraries*, Springer, 2016, pp. 188–199.
- [10] A. Fader, S. Soderland, O. Etzioni, Identifying relations for open information extraction, in: *Proceedings of the 2011 conference on empirical methods in natural language processing*, 2011, pp. 1535–1545.
- [11] O. Etzioni, R. E. Bart, M. D. Schmitz, S. G. Doderland, et al., Open language learning for information extraction, 2014. US Patent App. 14/083,261.
- [12] J. Fan, D. Ferrucci, D. Gondek, A. Kalyanpur, Prismatic: Inducing knowledge from a large scale lexicalized relation resource, in: *Proceedings of the NAACL HLT 2010 first international workshop on formalisms and methodology for learning by reading*, Association for Computational Linguistics, 2010, pp. 122–127.
- [13] C. Bizer, C. Becker, P. N. Mendes, R. Isele, A. Matteini, A. Schultz, Ldif-a framework for large-scale linked data integration (2012).
- [14] F. M. Suchanek, G. Kasneci, G. Weikum, Yago: a core of semantic knowledge, in: *Proceedings of the 16th international conference on World Wide Web*, 2007, pp. 697–706.
- [15] J. Hoffart, F. M. Suchanek, K. Berberich, G. Weikum, Yago2: A spatially and temporally enhanced knowledge base from wikipedia, *Artificial intelligence* 194 (2013) 28–61.
- [16] K. Bollacker, C. Evans, P. Paritosh, T. Sturge, J. Taylor, Freebase: a collaboratively created graph database for structuring human knowledge, in: *Proceedings of the 2008 ACM SIGMOD international conference on Management of data*, 2008, pp. 1247–1250.
- [17] S. G. Kobourov, Spring embedders and force directed graph drawing algorithms, *CoRR* abs/1201.3011 (2012). URL: <http://arxiv.org/abs/1201.3011>. arXiv:1201.3011.
- [18] C. Hurter, O. Ersoy, A. Telea, Graph bundling by kernel density estimation, in: *Computer graphics forum*, volume 31, Wiley Online Library, 2012, pp. 865–874.
- [19] M. Girvan, M. E. Newman, Community structure in social and biological networks, *Proceedings of the national academy of sciences* 99 (2002) 7821–7826.