



A Diagrammatic Approach for Visual Question Answering over Knowledge Graphs

Dmitry Mouromtsev¹, Gerhard Wohlgenannt^{1(✉)}, Peter Haase²,
Dmitry Pavlov³, Yury Emelyanov³, and Alexey Morozov¹

¹ Intern. Laboratory of Information Science and Semantic Technologies,
ITMO University, St. Petersburg, Russia
mouromtsev@mail.ifmo.ru, gwohlg@corp.ifmo.ru

² metaphacts GmbH, Walldorf, Germany
ph@metaphacts.com

³ Vismart Ltd., St. Petersburg, Russia

Abstract. In this demo we present a tool for visual question answering (QA) over the Wikidata knowledge graph based on diagrammatic representation and reasoning. The demo is built on top of the metaphacts platform with the Ontodia library embedded. In a user study, we demonstrate and evaluate the approach of diagrammatic question answering using questions from the QALD7 (Question Answering over Linked Data) benchmark for Wikidata. The effectiveness and limitations of the proposed approach are discussed in the evaluation and conclusions sections. In the demo session at ESWC, we plan to present our tool for visual QA and show its QA capabilities using incremental creation of diagrams.

Keywords: Visual data exploration · Knowledge graphs
Diagrammatic question answering · Wikidata

1 Introduction

One key challenge in the Semantic Web area is providing easy user access to the plethora of data hidden in Linked Data repositories. Direct access to the data requires an understanding of semantic query languages and the specific datasets. One way of abstracting users from the datamodels are natural language interfaces to Linked Data, which translate natural language queries into SPARQL and thereby hide the complexity [2]. In many cases, however, visual access to data is more intuitive. Many tools exist for browsing linked data, but they are not designed specifically for QA.

In this demo¹ we present a tool for visual QA over Linked Data, which uses a diagrammatic approach. The tool helps a user to navigate over knowledge graphs (KG) and to find answers to questions by visual means only; we use the term

¹ <http://wikidata.metaphacts.com>.

Diagrammatic Question Answering (DQA) to refer to the process. The benefits of visual exploration in data navigation include the quick detection of the most relevant properties and an easy understanding of the dataset characteristics. We use diagrammatic representations as an enabler of visual interaction with data and diagrammatic reasoning as way of data exploration. Diagrammatic representation and reasoning were suggested in visual language theory, which studies the cognitive and comprehensive features of this approach [5]. The tool is based on the Ontodia library², which was originally used for diagram-based collaborative ontology development [4].

We consider visual data exploration an information retrieval task, and evaluate our approach on the QA dataset from the QALD7³ challenge. In the experiments, we apply an exploratory diagramming system that uses Wikidata as KG, metaphactory⁴ as a KG platform and the Ontodia library as visual tool for data interaction.

2 System Description

The initial idea behind Ontodia was to enable a user to explore an unknown dataset. The graph nature of Linked Data inspired an incremental approach to visual exploration and a diagram-like semantic data representation. A user starts from a node of interest, and then uses the context menu to explore node properties and thereby the KG. By default all entity properties are listed alphabetically. A user can search for a specific property or topic of interest, then properties are sorted according to the query, with literal matches shown first, followed by the other properties ordered by similarity to the query. E.g., if a user is interested in the *family* relations of entity “Wolfgang Amadeus Mozart”, they can start the investigation from the diagram shown in Fig. 1. For detecting properties similar to a query, we employ Fasttext word embeddings [1] to represent Wikidata properties and user queries, for details see Wohlgenannt et al. [6].

In a nutshell, our solution for step-by-step data exploration is realized via context menus of entities. The user connects additional relations and entities to the diagram until the query is answered. The context menu displays all the object properties or connections that a chosen node has in the dataset, whereas the datatype properties are visible by clicking on the expansion icon located below the node.

The system is executed in a browser/JavaScript environment and consists of two main parts: the Ontodia library and metaphactory platform. Ontodia is responsible for most of the user experience tasks of DQA, including diagrammatic representation and diagrammatic reasoning. It is embedded into the metaphactory platform, which serves as entry point to the DQA solution – with rich search functionality and as a foundation for building Semantic Web applications.

² www.ontodia.org.

³ <https://project-hobbit.eu/challenges/qald2017>.

⁴ <http://www.metaphacts.com/product>.

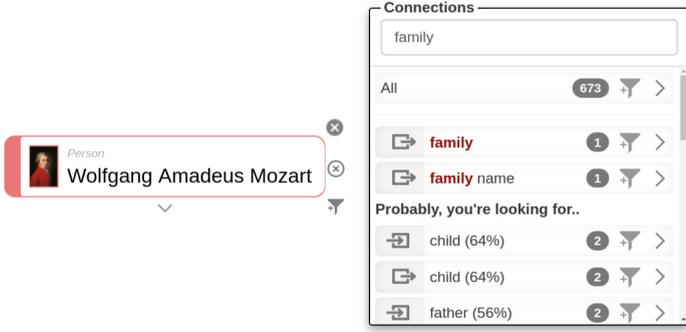


Fig. 1. Searching properties related to “family” of entity *Wolfgang Amadeus Mozart*

The demo application can be found at: <http://wikidata.metaphacts.com>. The basic QA process is exemplified by the following simple scenario:

1. Launch the demo application and enter the keyword that stands for the subject of the question.
2. Select a graph entity from the drop-down menu which best fits the subject.
3. Switch to diagrammatic representation clicking the “Show Diagram” icon.
4. Start data navigation from the root node by selecting the most relevant property in the context menu.
5. Build the graph incrementally, until arriving to the answer. Ontodia applies visual templates depending on the entity type (person, organization, location, etc.) to raise the expressivity of the diagram. Those visual templates can be freely configured by the user if needed.
6. Optionally, open the datatype properties of the selected node by clicking the expansion icon located below the node, which triggers the drown-down box.

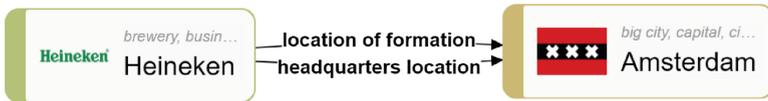


Fig. 2. Answering the question: *In what city is the Heineken brewery?*

Figure 2 shows a simple example diagram answering a query from the QALD7 challenge. The created diagram can also be saved and shared with other users.

3 Evaluation

In the evaluation, we address the research question whether diagrammatic representation and reasoning efficiently assist a user in QA and understanding knowledge from a large knowledge base.

3.1 Evaluation Setup

For the evaluation we reuse the QALD Benchmark, specifically task 4 “QA over Wikidata”. QALD was originally developed for systems that interpret natural language queries. We adopted the benchmark to evaluate our approach of DQA, where the answers are not produced by a system, but through data exploration performed by human users. We designed four unique questionnaires with nine questions each⁵.

The questions are grouped to the principles of question classification in QA [3]. Table 1 presents the main dimensions for classification, the question and answer type, and examples. We selected an equal number of questions for each type.

Table 1. Examples of test questions, and their classification.

Question type	Answer type	Example
WHO	Person	Who was the doctoral supervisor of Albert Einstein?
WHICH	Location	In what city is the Heineken brewery?
WHAT	Location	What is the capital of Cameroon?
NAME	Title	Show me all books in Asimov’s Foundation series
HOW	Number	How many people live in Poland?

20 persons, from 8 countries, participated in the evaluation. Each participant obtained a questionnaire and an instruction sheet by email. 131 diagrams (of 140 expected) were returned by the users.

3.2 Evaluation Results and Discussion

Participants were instructed to go through the questions and to build diagrams containing the answers, if possible. We measured both precision and recall – as ratio of correct answers to given answers, and correct answers to expected gold standard answers, respectively. Some question types like WHAT and WHICH provide best results, with an F1 around 90%, for WHO and HOW questions we measure an F1 of about 75%. Only for NAME questions F1 is rather low (38%). As in the example in Table 1, this type of question usually involves the listing of many results items.

From the conducted user study we learned that the three the most common types of diagrams are: (i) two connected nodes, (ii) a complex diagram with multiple nodes and (iii) a diagram containing one node with drop-down datatype property box shown. The first type is rather simple and easy to understand. The diagram for the second type is more complex the interpret, it reflects the process

⁵ <https://github.com/ontodia-org/DQA/wiki/Questionnaires2>.

of diagrammatic reasoning. In principle, the third type is the most simple, it contains only one node. E.g., for the question “How many people live in Poland?” a user only has to find the datatype property “population” in the drop-down box. However, many participants tried to solve such tasks only with object property connections, and failed to provide a result. In future work, we will take this finding into account to improve the user interface.

In general, the results are promising, but there are cases where it is difficult to find a correct answer with DQA, for example when the answer can be obtained only with joins of queries, or when it is hard to find the initial starting concept related to question focus.

4 Conclusions

The demo presents a tool developed for visual question answering using a diagrammatic approach (DQA). We evaluate the tool with questions from the QALD7 benchmark, specifically for QA over the Wikidata. For most types of questions, DQA provides promising results, and supports users in understanding the context of an answer and the characteristics of the knowledge graph itself.

The contributions of this work include: (i) a model for diagrammatic representation of semantic data, (ii) an exploratory diagramming system which integrates the metaphactory KG platform with Ontodia, and (iii) an evaluation of the diagrammatic approach with a user-study.

In future work we plan to solve the listed restrictions, eg. developing a more advanced property search box with embedded facet filtering to tackle more complex queries, or more powerful linguistic algorithms to address the problem of finding a suitable starting concept.

Acknowledgments. This work was supported by the Government of the Russian Federation (Grant 074-U01) through the ITMO Fellowship and Professorship Program.

References

1. Bojanowski, P., Grave, E., Joulin, A., Mikolov, T.: Enriching word vectors with subword information. arXiv preprint [arXiv:1607.04606](https://arxiv.org/abs/1607.04606) (2016)
2. López, V., Unger, C., Cimiano, P., Motta, E.: Evaluating question answering over linked data. *J. Web Sem.* **21**, 3–13 (2013)
3. Moldovan, D., et al.: The structure and performance of an open-domain question answering system. In: Proceedings of the 38th Annual Meeting of ACL, ACL 2000, pp. 563–570 (2000)
4. Mouromtsev, D., Pavlov, D., Emelyanov, Y., Morozov, A., Razdyakonov, D., Parkhimovich, O.: Workflow supporting toolset for diagram-based collaborative ontology development implemented in the open budget domain. In: Sack, H., Rizzo, G., Steinmetz, N., Mladenčić, D., Auer, S., Lange, C. (eds.) ESWC 2016. LNCS, vol. 9989, pp. 178–182. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-47602-5_35

5. Narayanan, N.H., Hübscher, R.: Visual language theory: towards a human-computer interaction perspective. In: Marriott, K., Meyer, B. (eds.) *Visual Language Theory*, pp. 87–128. Springer, New York (1998)
6. Wohlgenannt, G., Klimov, N., Mouromtsev, D., Razdyakonov, D., Pavlov, D., Emelyanov, Y.: Using word embeddings for visual data exploration with Ontodia and Wikidata. In: *BLINK/NLIWoD3@ISWC*, International Semantic Web Conference (ISWC), vol. 1932 (2017). CEUR-WS.org