

On Managing Dynamic Knowledge Graphs

Location PSL Research University - LAMSADE Université Paris-Dauphine, France
Supervisors Khalid Belhajjame and Maude Manouvrier
Contact Khalid.Belhajjame@dauphine.fr and Maude.Manouvrier@dauphine.fr
Candidate profile Master in Data or Computer Science or equivalent.
Solid skills in databases and knowledge graphs are required.
A good knowledge in algorithmic, programming and machine learning is appreciated.

Context

Knowledge graphs are gaining ground as a means of encapsulating and sharing domain knowledge. Large companies, such as Amazon, Bosch, Google, Microsoft and Zalando, have already adopted knowledge graphs to represent and store their knowledge bases. In addition to enabling the sharing, querying and retrieval of facts of interest to a business or community, knowledge graphs have recently gained recognition and are becoming the backbone of cognitive artificial intelligence. Gartner predicts that the application of knowledge graphs and graph mining will grow by 100% per year to enable more complex and adaptive data science¹.

Objectives

In the context of this thesis, we will focus on RDF knowledge graphs, probably the most widely used class of knowledge graphs. A number of problems arise when managing these knowledge graphs, ranging from their construction to their exploration and exploitation. We will mainly focus on the management of dynamic knowledge graphs. Indeed, knowledge is intrinsically dynamic: the sources that feed the knowledge graph can undergo changes that have an impact on the knowledge graph itself. Moreover, new promising sources can be added to the list of sources used to enrich the knowledge graph, and other sources that are no longer relevant can be dropped, which in turn has an impact on the facts (nodes and relations) composing the knowledge graph. The general objective of the thesis is therefore: **To design new solutions to assist knowledge graph providers and users to better handle the effects of dynamic knowledge graphs.**

To achieve the above goal, a number of tasks will be undertaken, from state of the art review to design and implementation of algorithmic solutions to:

1. Characterize the changes a knowledge graph, may undergo,
2. Identify the maintenance actions that can be undertaken to smoothly manage these changes, and
3. Assess and manage the impact on the applications that utilize the knowledge graph.

References

- [1] Grigoris Antoniou, Sotiris Batsakis, Raghava Mutharaju, Jeff Z. Pan, Guilin Qi, Ilias Tachmazidis, Jacopo Urbani, and Zhangquan Zhou. A survey of large-scale reasoning on the web of data. *Knowledge Eng. Review*, 33:e21, 2018.
- [2] Garima Gaur, Arnab Bhattacharya, and Srikanta Bedathur. How and why is an answer (still) correct? maintaining provenance in dynamic knowledge graphs. *CIKM '20: The 29th ACM Int. Conf. on Information and Know. Management, Virtual Event, Ireland, Oct. 19-23, 2020*, pp. 405–414.
- [3] François Goasdoué, Ioana Manolescu, and Alexandra Roatis. Efficient query answering against dynamic RDF databases. In *Joint 2013 EDBT/ICDT Conf., EDBT '13 Proc., Genoa, Italy, March 18-22, 2013*, pp. 299–310.
- [4] M. Lissandrini, T. Bach Pedersen, K. Hose, D. Mottin. Knowledge graph exploration: where are we and where are we going? *SIGWEB Newsl.* 4:1-4:8, 2020.

¹blog.biostrand.be/en/knowledge-graphs-and-the-power-of-context