Flexible and Extensible Competency Management with Knowledge Graphs

Nicolas Heist[1,2†][0000−0002−4354−9138], Peter Haase[1], and Simon Scerri[1]

[1] metaphacts GmbH, Walldorf, Germany
{nh,ph,simon.scerri}@metaphacts.com
[2] Data and Web Science Group, University of Mannheim, Germany
nico@informatik.uni-mannheim.de

Abstract. Especially in the diverse and fast-paced field of Artificial Intelligence it is imperative to have a clear picture of relevant competencies and how they are distributed within or over organisations. For this purpose, we have developed a generic competency ontology that can be used to describe competencies of people and organisational structures in the Artificial Intelligence domain. The ontology is embedded in an application to create, manage, and utilize a Competency Knowledge Graph. In our presentation we show concrete application scenarios, advantages, and challenges.

Competency Management With a growing team, organisation, or consortium size, it becomes increasingly difficult for individuals to get an overview of their business partners' or co-workers' competencies. This is problematic in various industry settings, but especially in the field of Artificial Intelligence, which is strongly driven by innovation and where skills and know-how are quick to change and evolve. In medium and large-sized companies, where many projects are conducted and where AI technologies are increasingly involved, individuals often need to quickly identify experts of specific topics in their extended network, as well as existing alternative solutions and general know-how produced by their peers.

The value of choosing a Knowledge Graph paradigm for extensive information retrieval is widely documented [2]. We therefore target a Competency Knowledge Graph whose setup, population and upkeep can be achieved with minimal effort, while also exploiting existing open data to further augment it at virtually no cost.

Our implemented solution demonstrates how the following common business scenarios can be supported:

– Browsing through AI competencies, their interrelations and relevant publications in order to resolve an information need. For instance, given a concrete business problem, one can find the competencies relevant to solve it.

† This work was created during an internship at metaphacts GmbH.

– Discover persons or organisations having relevant competencies; suggest related competencies to identify knowledge gaps; find completed projects to foster reuse (based on competencies defined using the ontology).

The Competency Ontology The competency ontology provides the three generic classes *Competency*, *CompetencyHolder*, and *CompetencyTarget*. These are supposed to be extended by sub-classes in order to describe a concrete scenario. A *Competency* is a particular skill or piece of know-how that one might possess or apply. To be able to describe competencies in the AI domain, we incorporate the data from AI-KG [1], which is a knowledge graph automatically extracted from the Microsoft Academic Graph.[3] The graph describes 850K research entities (i.e. specific *Tasks*, *Methods*, *Metrics*, and *Materials* in the AI domain) and their relations. We include those as sub-classes of *Comptency* in our ontology. *CompetencyHolder*s are entities that possess or apply competencies. In our case, those are persons and organisations. To model them, we re-use definitions from existing ontologies like FOAF[4] and ORG[5] as sub-classes. *CompetencyTarget*s serve to describe the context in which concrete competencies are applied in. In a business setting, these are concepts like projects, processes, or applications. Here, we again rely on well-established ontologies like FOAF.

Application Design We use metaphactory [3] for the design of the competency management application as it provides several advantages; primarily its ability to guide the input of competency descriptions based on the underlying ontology through intuitive components like ontology-driven semantic forms and search. Furthermore, metaphactory's low-code approach makes it possible to rapidly develop custom dashboards. In general, a knowledge graph-based approach to define competencies reduces effort with respect to data integration, scalability and flexibility. Apart from the re-use of established ontologies to cover the modeling requirements, external data sources can also be tapped using federated queries. For instance, we display information about publications of AI competencies from an external SPARQL endpoint of the Microsoft Academic Graph (as it would be too large to integrate). The approach is also easily extendable to domains outside of AI due to the open nature of the RDF data model and metaphactory's low-code dashboarding capabilities.

References

1. Dessì, D., et al.: AI-KG: an automatically generated knowledge graph of artificial intelligence. In: International Semantic Web Conference. pp. 127–143. Springer (2020)
2. Galkin, M., Auer, S., Vidal, M.E., Scerri, S.: Enterprise knowledge graphs: A semantic approach for knowledge management in the next generation of enterprise information systems. In: ICEIS (2). pp. 88–98 (2017)
3. Haase, P., Herzig, D.M., Kozlov, A., Nikolov, A., Trame, J.: metaphactory: A platform for knowledge graph management. Semantic Web **10**(6), 1109–1125 (2019)

[3] https://www.microsoft.com/en-us/research/project/microsoft-academic-graph/
[4] http://www.foaf-project.org
[5] https://www.w3.org/TR/vocab-org/