Knowledge Graphs for Workforce Reskilling Guidance

Diarmuid Cahalane*, Christophe Guéret*, Chan Le Van*, Patrick Moreau†, David Byrne†

Accenture Labs* and AI & Analytics†
The Dock, 7 Hanover Quay, Dublin
diarmuid.cahalane@accenture.com
https://labs.accenture.com

Abstract: Shifting supply and demand for skills mean that mid-career reskilling of employees is an important concern for many companies and employees alike [1]. We demonstrate the potential for Knowledge Graph Embeddings to capture contextual information about skills and work for the use case of inspiring and guiding employees selecting new-skilling options.

Keywords: Knowledge Graphs, Knowledge Graph Embeddings, Reskilling Guidance

1 Background

The choice of optimal skills for an employee to learn is informed by factors including expected future demand and how the skill relates to the employee's existing skill set. How closely a new skill relates to existing ones can influence the likelihood of the skill being professionally relevant and of interest to the employee but also how much effort may be required to develop it. Tools to help employees explore these relationships may lead to better informed decision making. Moreover, the task of documenting the existing skill set in the first instance, to inform recommendations, may be helped by suggesting proximal skills.

A knowledge graph (KG) represented enterprise data concerning employees' skills and industry specializations. We used AmpliGraph [2] to generate high dimensional embedding vectors for the skills and then a low-dimensional projection of those vectors to create a visual aid for employees to explore skills which were proximal in the embedding space. From a separate data source, we imported categorical forecasts of future demand for those skills and overlayed this in the visual presentation of the skill embeddings, to further inform the employees exploration of potential new skills to develop.

2 Our Approach

Two classes of nodes and three classes of edges were used to populate a KG. The two types of nodes represented employees and skills respectively. The "person_has_skill" edge type denoted which existing skills an employee held. Edges of the type "person_has_industry" linked employees to nodes representing categories of industry expertise. Ontologies were available for both skills and industry categories, so we assimilated these into the KG by way of "skill_has_skill_parent" and "industry_has_industry_partent" edge types, respectively. The KG was represented as a list of <subject, predicate, object> triples. AmpliGraph, a suite models for graph embedding which has been open-sourced by Accenture, consumed the triples to generate 200 dimensional vectors. Thus we learned vector representations for approximately 6,000 skills nodes, 1,000 industry nodes and 300,000 employees.

To help employees explore the skills nearby a chosen starting skill in the embedding space, we selected the closest neighbors of the starting skill in the embedding space (using a Euclidean distance metric). Projecting the vectors for the starting skill and its neighbors into a 2- or 3-dimensional space and visualizing the results gave a presentation that may help an employee choose skills to note as "already known" or "of interest to learn" [3]. The forecasted demand category for the skills was also displayed as part of the skill label and visually by way of a color code on the skills.

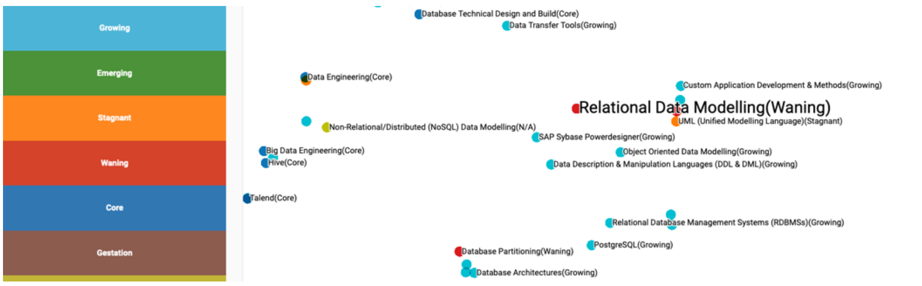

Fig. 1. A view of the UI visualizing skills embeddings and their forecasted demand

3 Conclusions

Qualitatively, the results show meaningfully related skills have nearby vectors in the embedding space. This is true even for relationships not explicit in the input data. We conclude that KG embeddings are a promising approach to help employees sift through myriad skill choices and make better informed, more confident learning choices.

4 References

1. Ovanessoff, A., Sage-Gavin, E., Morvan, L.,: It's learning. Just not as we know it. https://www.accenture.com/us-en/insights/future-workforce/transforming-learning
2. Costabello, L., Pai, S., Van, C.L., McGrath, R., McCarthy, N., Tabacof, P.: Ampli-Graph: a Library for Representation Learning on Knowledge Graphs (Mar 2019). https://doi.org/10.5281/zenodo.2595043, https://doi.org/10.5281/zenodo.2595043
3. TensorBoard Embedding Projector https://github.com/tensorflow/tensorboard