14 True Facts About Knowledge Graphs

Bogdan Arsintescu, Scott Meyer, and Swee Lim
LinkedIn
San Jose, California
barsintescu@linkedin.com

LinkedIn

Abstract

Knowledge Graphs have been sprouting in the past decade driving a wedge between the traditional relational database and IR based indices. The fundamental change is modelling data as a graph, with first class relationships between entities, where edges separate them from relational databases tables and nodes strong identities separate them from word based IR indices. We present a survey of the known knowledge graphs and the application classes that they enable, including the LinkedIn Economic Graph. We discuss how graph databases differ from relational ones and IR indices and where the graph databases would benefit from IR techniques. We conclude with a set of challenges we see in our experience in scaling Knowledge graphs and graph databases in low latency on-line applications.

Bio

Bogdan Arsintescu has worked with graphs and semantic data throughout his whole career, most recently at Google Knowledge Graph leading the graph query language team, in Google Research working on semantic trajectories using location data and at LinkedIn as a manager in the graph database team. He received his Ph.D. in CS from Technical University Delft, the Netherlands and MSc in EE from Politehnica University in Bucharest, Romania.

In: L. Dietz, C. Xiong, E. Meij (eds.): Proceedings of the First Workshop on Knowledge Graphs and Semantics for Text Retrieval and Analysis (KG4IR), Tokyo, Japan, 11-Aug-2017, published at http://ceur-ws.org