THE STRUCTURE OF KNOWLEDGE

Peter Bloem · Frank van Harmelen · Pieter Adriaans

The graph is one of the most powerful and versatile methods we have to store knowledge. In this project we will focus on *knowledge graphs*: graphs that encode objects in their nodes, and use labeled, directed links to encode a variety of different relations between these objects. Here are some examples:

netflix

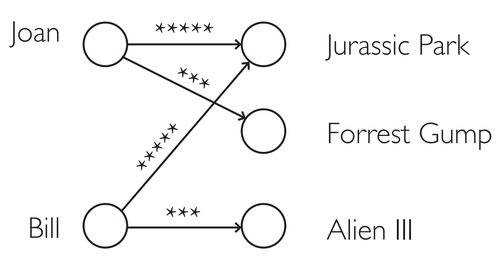

the semantic web

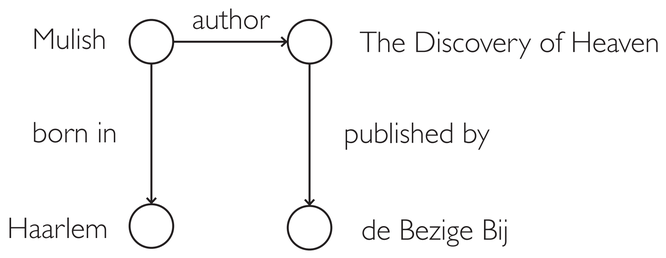

social networks

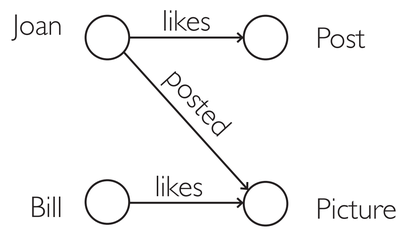

Knowledge graphs have now become so complex, that statistical methods can offer new insights beyond traditional inference. Such methods have already been succesful in *link prediction*, and *classification* tasks. However, what a statistical model, like a tensor decomposition,WW actually means when translated back to the domain of the graph, is still an open question.

network motifs

Network motifs are frequently recurring subgraphs. A large problem in graph analysis is the difficulty of cutting the full data up into largely independent chunks. Motif analysis provides one solution. The challenge in adapting motif analysis to knowledge graphs is dealing with the labels of nodes and links.

tensor decomposition

Where a normal graph can be represented as an adjacency matrix, knowledge graphs can be represented as a tensor, with the adjacency matrix for each type of link forming a slice in the tensor. A tensor decomposition is analogous to the matrix decomposition used in many recommender systems.

graph models

The more accurately we can model knowledge graphs in general, the easier it is to bring out the structure that defines a specific knowledge graph. There is some progress in this area for generic graphs, which should provide a promising starting point for the analysis of knowledge graphs.

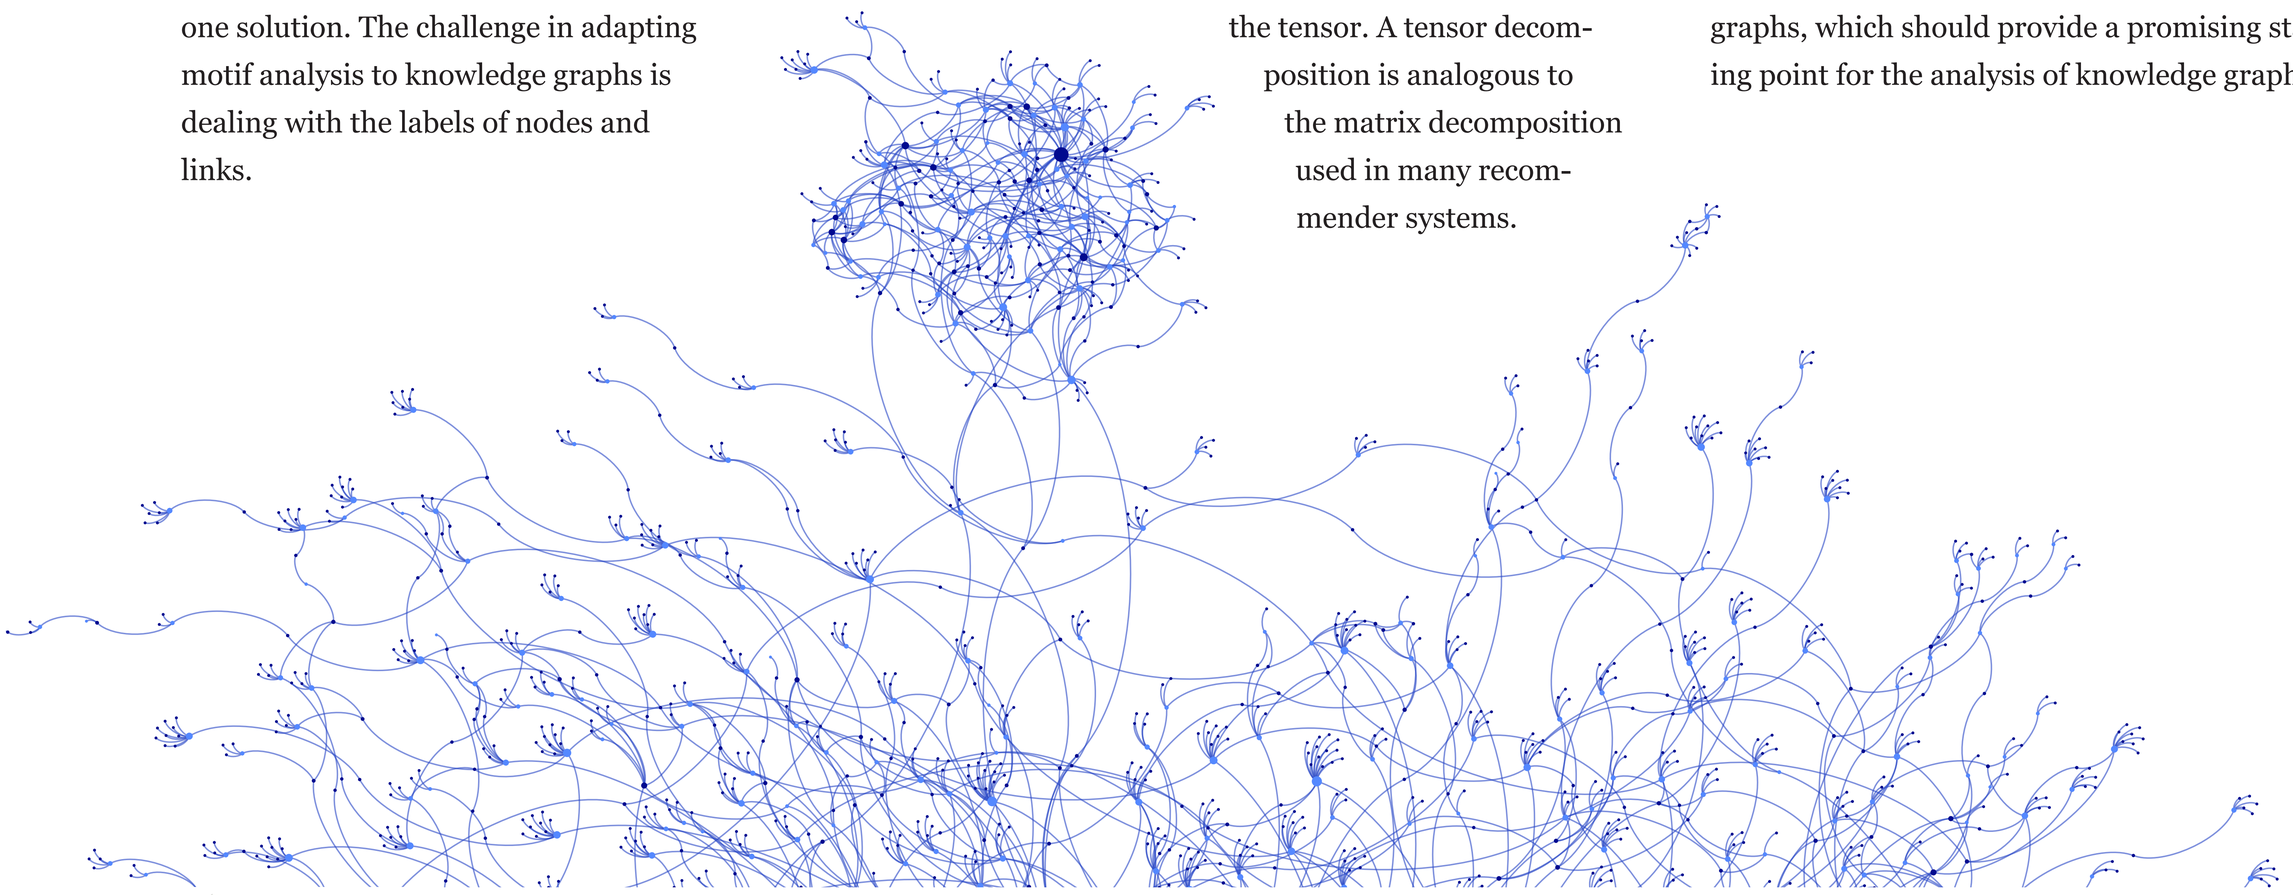

Amsterdam Data Science

This work was supported by the Amsterdam Academic Alliance Data Science (AAA-DS) Program awarded to the UvA and VU Universities.