Scaling up Logical Query Embeddings on Knowledge Graphs

Hongyu Ren*, Hanjun Dai*, Bo Dai, Xinyun Chen, Denny Zhou, Jure Leskovec, Dale Schuurmans

Stanford University, Google Brain, UC Berkeley

Multi-hop Reasoning on KGs

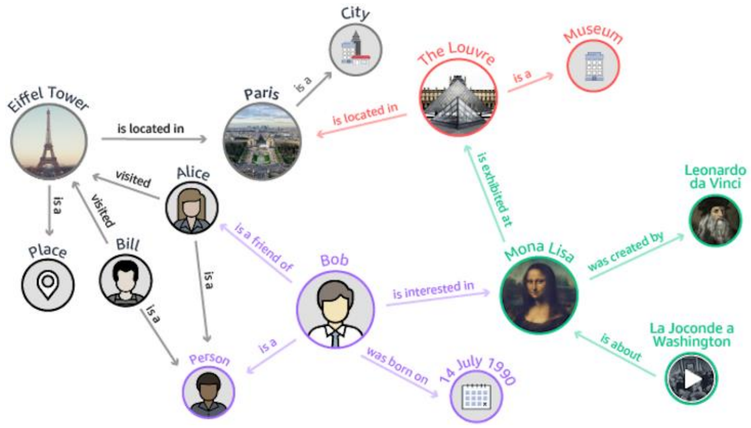

Answering complex queries on KGs

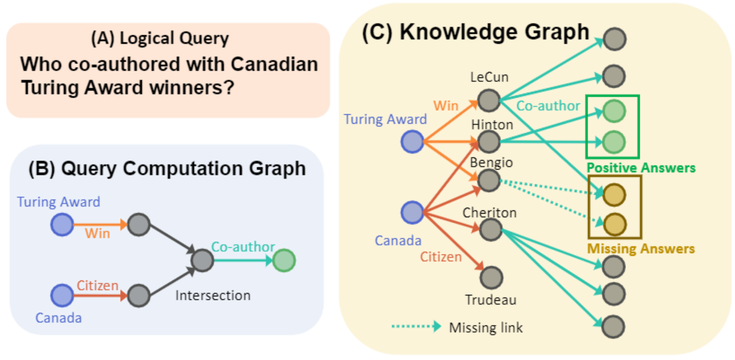

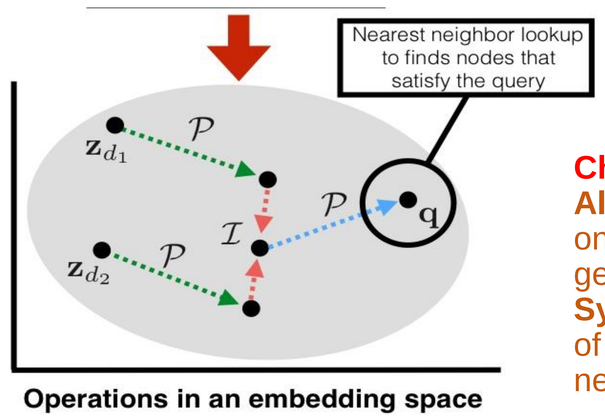

SrKG: Scalable reasoning Knowledge Graph Embeddings

Query Instantiation from Template

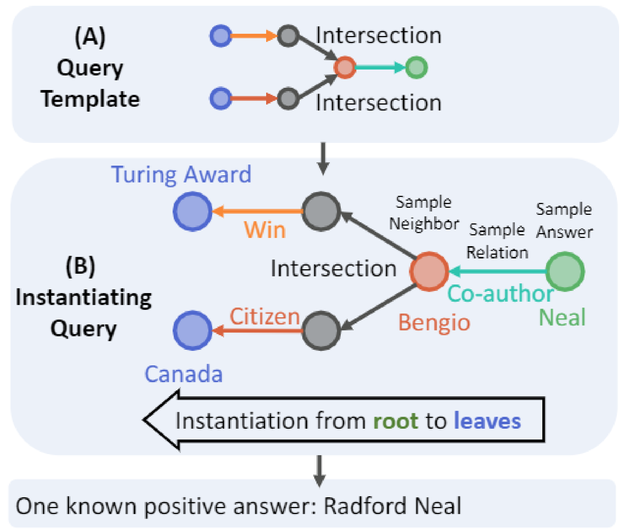

Negative Entities Sampling

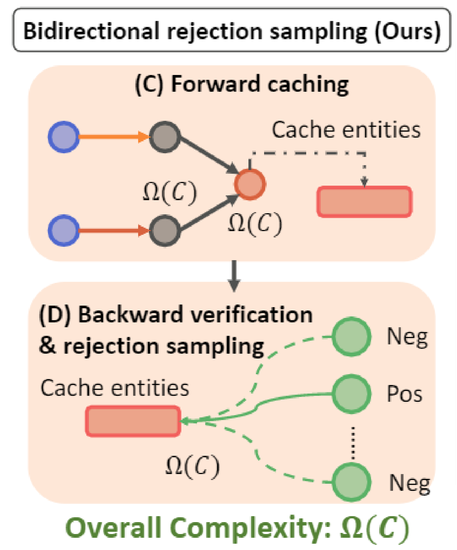

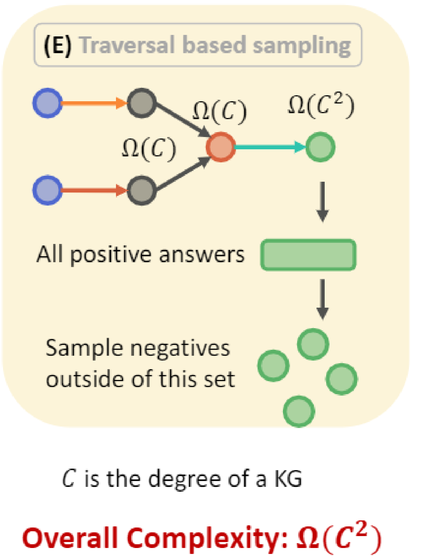

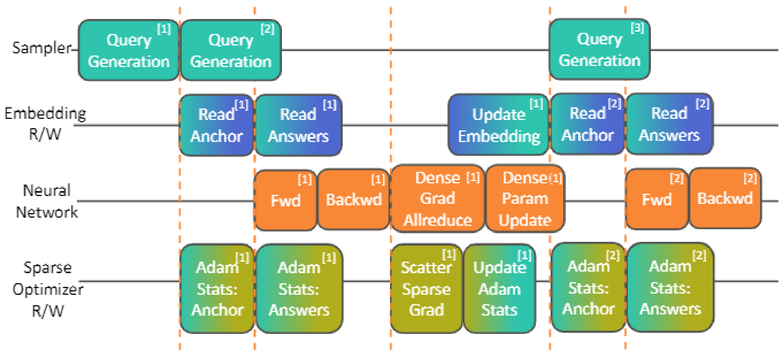

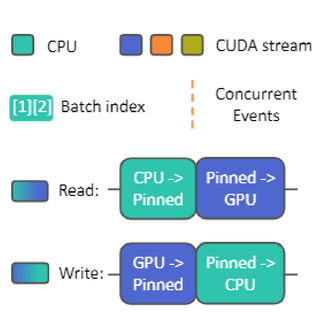

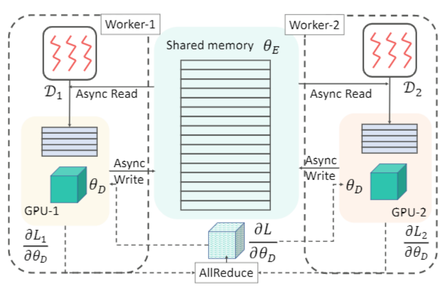

Challenges to scale to large KGs:

Algorithmic: needs an efficient training data (query, one positive answer, several negative answers) generation algorithm

System: (1) sparse + dense parameters (2) scheduling of training data generation, embedding r/w, neural network feed forward, optimization

SrKG performance:

(1) For the first time, scale reasoning algorithms to **KGs with more than 86m nodes**

(2) On small graphs, **+99.5% speed, -33.6% GPU memory**

(3) Achieve (almost) graph-size agnostic speed and GPU usage