ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Distributed Computing

Prof. R. Wattenhofer

Distance Preserving Graph Embedding

GPS use is now prevalent. The users want to know immediately, what is the shortest path from their office to the nearest coffee shop or which road they should take if they are driving from their home in Zurich to Berlin. However, applying traditional shortest path algorithms such as Dijkstra's is slow. Because road graphs have a lot of edges and vertices. For example, New York's road network has 264 thousand nodes, while the road graph of western Europe has 18 million nodes. The system also has to answer thousands of such queries per second.

To speed up these shortest path queries one can modify a graph to introduce shortcuts that preserve distance or compute labels for each node, that encode distance to some landmark vertices. Landmarks are chosen such, that you can recover distance between any two nodes using only their labels[1].

In this project, we will instead build a deep learning model that predicts distance and next-hop to take along the shortest path between any two nodes on a road graph. The main goal is to achieve quicker query times compared to traditional approaches.

To achieve this goal we will modify existing graph embedding techniques used in machine learning such, that the resulting embeddings are more representative of the true distances on the graph.

Requirements: Interest in algorithmic problems. Knowledge in deep learning, or a solid background in machine learning. Experience with Python and TensorFlow or PyTorch is an advantage as well as basic knowledge in shortest path algorithms. We will have weekly meetings to address questions, discuss progress and think about future ideas.

Interested? Please contact us for more details!

Contact

- Karolis Martinkus: martinkus@ethz.ch, ETZ G60.1

[1] See PODC lecture notes: https://disco.ethz.ch/courses/podc/lecturenotes/chapter14.pdf