Chair of Communication Networks
TUM Department of Electrical and Computer Engineering
Technical University of Munich

Bachelor's Thesis

Learning to forward based on the euclidean embedding of graphs

Routing on unstructured graphs is a problem in large data center topologies. Unstructured graphs do not allow an efficient encoding of the Forwarding Information Base with Longest Prefix Matching. An open problem is thus an effective way to forward packets in such networks.

A possible solution could be the spatial navigation of the unstructured graphs with Deep Neural Networks (DNNs). Recent work [1-4] shows that DNNs can learn to navigate in continuous euclidean spaces. The goal of this thesis is to investigate if DNNs are also capable of learning forwarding behavior on the euclidean embedding of graphs. In the euclidean embedding, nodes are mapped to two dimensional points, i.e., coordinates. The coordinates can be encoded as part of the IP-addresses, e.g., using half precision floating points. The task of the DNN is to predict the outgoing ports a switch should forward a packet on such that the packet travels on one of the shortest paths between the switch and the destination. The DNN makes the prediction using three inputs: The coordinates of the switch, the headings of the ports of the switch, and the coordinates of the destination.

The thesis should evaluate if it is possible to learn the shortest path forwarding behavior on small (tens of nodes) and arbitrarily structured graphs based on the graphs' embedding in a two-dimensional euclidean space. If successful, the thesis should further investigate if the forwarding behavior generalizes to unseen graphs, and whether the forwarding behavior scales to large graphs with hundreds of nodes. For all tasks, the NetworkX library [5] provides the required functionality to generate random graphs and embed the graphs in an euclidean space using a force directed algorithm.

References:

[1] C. J. Cueva and X.-X. Wei, "Emergence of grid-like representations by training recurrent neural networks to perform spatial localization," 2018. [Online]. Available: https://openreview.net/forum?id=B17JTOe0-

[2] A. Banino et al., "Vector-based navigation using grid-like representations in artificial agents," Nature, vol. 557, no. 7705, pp. 429–433, May 2018, doi: 10.1038/s41586-018-0102-6.

[3] S. Brahmbhatt and J. Hays, "DeepNav: Learning to Navigate Large Cities," CoRR, vol. abs/1701.09135, 2017, [Online]. Available: http://arxiv.org/abs/1701.09135

[4] P. Mirowski et al., "Learning to Navigate in Cities Without a Map," ArXiv e-prints, Mar. 2018.

[5] Aric A. Hagberg, Daniel A. Schult and Pieter J. Swart, "Exploring network structure, dynamics, and function using NetworkX", in Proceedings of the 7th Python in Science Conference (SciPy2008), Gäel Varoquaux, Travis Vaught, and Jarrod Millman (Eds), (Pasadena, CA USA), pp. 11–15, Aug 2008

Advisors

Patrick Krämer