Invertible Neural Networks for Knowledge Graph Embeddings

Sebastian Mair
sebastian.mair@liu.se

Samuel Matthiesen
samma@dtu.dk
samuel.matthiesen@liu.se

October 2, 2024

Background

Knowledge Graphs are a basis for modeling real world facts within a computational setting. Here, nodes represent entities and relations are how those entities are connected. For example, consider one fact as two entities u = "Sweden" and v = "Carl XVI Gustaf" with the relation r = `king-of`. The goal is to find embeddings for u, v and learn a function for the relation, such that $r(u) \approx v$ for a large list of facts. This is known as knowledge graph embedding (Sun et al., 2019; Nayyeri et al., 2021). Common strategies to model this as a machine learning problem is to learn functions that represent the relationships, where these take a vector representing an entity as input and output another vector representing another entity. Those functions are then trained to mimic the relations seen in the data, that is, if the entity u is connected to the entity v, then the learned function takes u as input and outputs a vector very close to v. The goal of this project is to learn graph embeddings on publicly available data and explore invertible neural networks (INNs, see, *e.g.*, Papamakarios et al. (2021)) as the relation functions.

Research questions

Due to their invertibility properties, INNs are able to naturally model symmetric relations, as the opposite direction is simply given by the inverse function.

- What are appropriate INN architectures for this setting?
- How to model relations which do not admit a symmetric property?
- For relations which admit multiple inverses, can this be modelled stochastically, *i.e.*, the output is randomly one entity or potentially others?

Eligibility requirements

- Sound knowledge of machine learning (very good grades in relevant courses)
- Very good programming skills
- Knowledge of Graph Neural Networks is beneficial
- Knowledge of frameworks such as PyTorch/JAX is beneficial

Please attach your CV and transcripts when applying.

References

Nayyeri, M., Xu, C., Hoffmann, F., Alam, M. M., Lehmann, J., and Vahdati, S. (2021). Knowledge graph representation learning using ordinary differential equations. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 9529–9548.

Papamakarios, G., Nalisnick, E., Rezende, D. J., Mohamed, S., and Lakshminarayanan, B. (2021). Normalizing flows for probabilistic modeling and inference. *Journal of Machine Learning Research*, 22(57):1–64.

Sun, Z., Deng, Z.-H., Nie, J.-Y., and Tang, J. (2019). Rotate: Knowledge graph embedding by relational rotation in complex space. In *International Conference on Learning Representations*.