Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Distributed Computing

Prof. R. Wattenhofer

Spectral Denoising Diffusion Probabilistic Models for Graphs

Graph generation is an interesting problem which is encountered when trying to generate novel molecules or proteins with desired properties or variations of social graphs that used to test hypotheses in social sciences. However, graph generation is very a challenging problem due to the non-uniqueness of graphs and the complex non-local dependencies between their edges. One of the most promising deep generative models lately have been the denoising diffusion probabilistic models Ho et al. [2020]. They have also been applied to small graph and molecule generation with great success, but only really work with small graphs with up to twenty nodes.

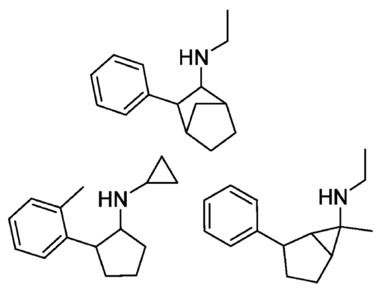

Recently we had good success improving graph generation with generative adversarial networks (GANs) by introducing spectral conditioning and generating graph spectra Martinkus et al. [2022]. In this thesis we will investigate how to include similar graph-informed biases into the denoising generative models to improve their performance on larger molecule, protein and potentially other graph generation (e.g. object meshes).

Requirements: Strong motivation, knowledge in deep learning, or a solid background in machine learning. Experience with Python and PyTorch or TensorFlow is an advantage as well as knowledge in graph theory, generative models and graph neural networks.

Interested? Please contact us for more details!

Contact

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

References

J. Ho, A. Jain, and P. Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 2020.

K. Martinkus, A. Loukas, N. Perraudin, and R. Wattenhofer. SPECTRE: Spectral conditioning helps to overcome the expressivity limits of one-shot graph generators. *Proceedings of the 39th International Conference on Machine Learning*, 2022.