ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Distributed Computing

Prof. R. Wattenhofer

Fast and Efficient Language Modelling with Differentiable Logic Gate Networks

Machine learning models are often huge, which causes them to have high latency, low throughput, and high energy usage. In response to this, methods such as DiffLogic [3–5], a logic-gate-based neural network architecture designed for FPGA acceleration (see in Figure 1), offer an efficient computational alternative.

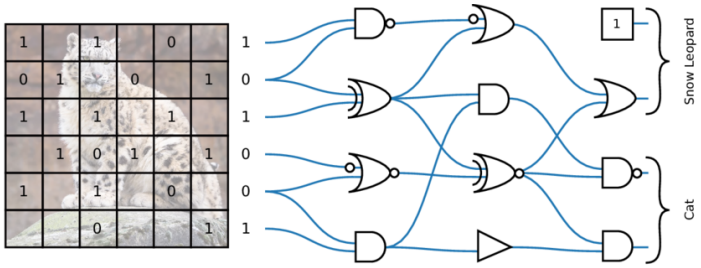

Figure 1: Overview of DiffLogic for image classification.

Most work in the area of difflogic focuses on image classification, but Bührer et al. [1] created the first language model using difflogic. These models only got BLEU scores around 20, so there is still a big gap to close before the models are competitive with modern architectures. Later work has generalized the difflogic architecture to Look Up Table (LUT) architectures called difflut (https://github.com/aplesner/difflut).

With that in mind, this project will focus on extending the initial work by Bührer et al. [1] to more modern architectures such as GRU [2], larger datasets, and newer difflogic/difflut architectural variants [5].

The project will be done fully remote. We would have weekly meetings on Zoom to go over results, discuss open questions, and resolve any potential problems. You will have a lot of possibilities to shape the project in the directions you find the most interesting.

Requirements

Solid programming skills in Python and knowledge of machine learning evaluation are required.

Contact

In a few short sentences, please describe your interest in this project and any relevant coding experience or background (e.g., projects or coursework).

- Andreas Plesner: aplesner@ethz.ch, ETZ G95
- Till Aczel: taczel@ethz.ch, ETZ G60.1

References

[1] Simon Bührer, Andreas Plesner, Till Aczel, and Roger Wattenhofer. Recurrent Deep Differentiable Logic Gate Networks, August 2025.

[2] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling, 2014. URL https://arxiv.org/abs/1412.3555.

[3] Felix Petersen, Christian Borgelt, Hilde Kuehne, and Oliver Deussen. Deep differentiable logic gate networks. *Advances in Neural Information Processing Systems*, 35:2006–2018, 2022.

[4] Felix Petersen, Hilde Kuehne, Christian Borgelt, Julian Welzel, and Stefano Ermon. Convolutional differentiable logic gate networks. *Advances in Neural Information Processing Systems*, 37:121185–121203, 2024.

[5] Lukas Rüttgers, Till Aczel, Andreas Plesner, and Roger Wattenhofer. Light Differentiable Logic Gate Networks, September 2025.