Deep Differentiable Logic Gate Networks: Neuron Collapse Through a Neural Architecture Search Perspective

Shakir Yousefi

Supervisors: Andreas Plesner, Till Aczél

Prof. R. Wattenhofer

March 2025

1 Introduction

Deep Differentiable Logic Gate Networks (Petersen, Borgelt, et al. 2022) (Petersen, Kuehne, et al. 2024) integrate classical Boolean logic gates into neural networks via continuous relaxation. Each neuron represents a categorical distribution over possible logic gates, which empirically collapses to a single discrete gate during inference. This *neuron collapse* enables efficient execution on CPUs and FPGAs without compromising accuracy. Empirical results for Deep Differentiable Logic Gate Networks are encouraging, yet their theoretical foundations and potential extensions warrant further investigation.

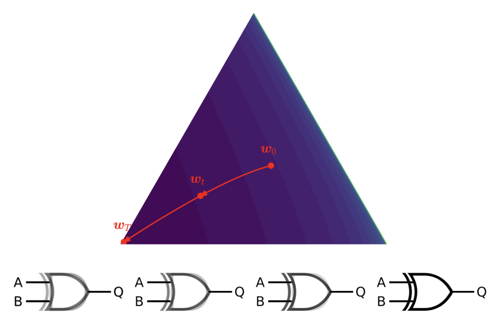

Neural Architecture Search (NAS) (Zoph and Le 2017) (Pham et al. 2018) automates the exploration of neural network architectures from a fixed set of operations. Differentiable Architecture Search (DARTS) (Liu, Simonyan, and Yang 2019) tackles this by stacking cells of directed acyclic graphs with softmax-weighted edges over candidate operations (e.g., convolutions, pooling). Like Deep Differentiable Logic Gate Networks, DARTS relies on the network converging toward discrete, prunable choices.

This project will (1) explore conditions and guarantees needed for *neuron collapse*, (2) compare how this phenomenon is related to DARTS, and (3) explore Deep Differentiable Logic Gate Networks in the framework of NAS and consider its implications and extensions.

2 Project Timeline

A tentative timeline for the project is as follows:

- **Weeks 1**: Literature review and theoretical framework development.
  - Read and understand relevant literature.
  - Clarify reasonable theoretical and practical assumptions for framework.
- **Weeks 1-3:** Explore conditions and guarantees needed for neuron collapse in Deep Differentiable Logic Gate Networks.
  - Explore interplay between softmax and gradient descent in a toy-example.
  - Conditions and restrictions imposed by logic gates.
  - Deep Differentiable Logic Gate Networks as a composition of optimization problems.
- **Weeks 3-5:** Compare Deep Differentiable Logic Gate Networks to DART.
  - Relate DART and Deep Differentiable Logic Gate Network optimization problem.
  - Compare behavior of general functions and logic gates.
- **Weeks 5-6:** Deep Differentiable Logic Gate Networks in the framework of NAS.
  - Explore if other techniques and results for NAS would be effective for Deep Differentiable Logic Gate Networks.
- **Week 7:** Writing final report and mock presentation.
- **Week 8:** Final presentation.

3 Extensions and Related Works

NAS and DARTS are well-studied, and have several results and extensions. Unifying Deep Differentiable Logic Gate Networks within the NAS framework may allow us to reinterpret or extend established results in the NAS/DARTS literature. Results on robustness of DARTS include (Zela et al. 2020) (Chen and Hsieh 2020). Recent efforts on NAS (Mellor et al. 2021) reframes DARTS as a one-shot procedure and instead proposes a zero-shot approach. Current state-of-the-art zero-shot approaches (Jiang, Wang, and Bie 2023) exploits the eigenvalues of the correlation matrix of the network.

4 Student Responsibilities and Deliverables

- Weekly meetings with the supervisors to discuss progress and challenges.
- Weekly detailed research logs.
- Submission of a final report in English.
- Delivery of a 15-minute final presentation summarizing project outcomes.

References

Chen, Xiangning and Cho-Jui Hsieh (2020). "Stabilizing Differentiable Architecture Search via Perturbation-based Regularization". In: *Proceedings of the 37th International Conference on Machine Learning*. URL: https://proceedings.mlr.press/v119/chen20f.html.

Jiang, Tangyu, Haodi Wang, and Rongfang Bie (2023). "MeCo: Zero-Shot NAS with One Data and Single Forward Pass via Minimum Eigenvalue of Correlation". In: *Thirty-seventh Conference on Neural Information Processing Systems*. URL: https://openreview.net/forum?id=KFm2lZiI7n.

Liu, Hanxiao, Karen Simonyan, and Yiming Yang (2019). "DARTS: Differentiable Architecture Search". In: *International Conference on Learning Representations*. URL: https://openreview.net/forum?id=S1eYHoC5FX.

Mellor, Joe et al. (2021). "Neural Architecture Search without Training". In: *Proceedings of the 38th International Conference on Machine Learning*. URL: https://proceedings.mlr.press/v139/mellor21a.html.

Petersen, Felix, Christian Borgelt, et al. (2022). "Deep Differentiable Logic Gate Networks". In: *Advances in Neural Information Processing Systems*. Ed. by Alice H. Oh et al. URL: https://openreview.net/forum?id=vF3WefcoePW.

Petersen, Felix, Hilde Kuehne, et al. (2024). "Convolutional Differentiable Logic Gate Networks". In: *The Thirty-eighth Annual Conference on Neural Information Processing Systems*. URL: https://openreview.net/forum?id=4bKEFyUHT4.

Pham, Hieu et al. (2018). "Efficient Neural Architecture Search via Parameters Sharing". In: *Proceedings of the 35th International Conference on Machine Learning*. URL: https://proceedings.mlr.press/v80/pham18a.html.

Zela, Arber et al. (2020). "Understanding and Robustifying Differentiable Architecture Search". In: *International Conference on Learning Representations*. URL: https://openreview.net/forum?id=H1gDNyrKDS.

Zoph, Barret and Quoc Le (2017). "Neural Architecture Search with Reinforcement Learning". In: *International Conference on Learning Representations*. URL: https://openreview.net/forum?id=r1Ue8Hcxg.