Reg. No: 2021/41/B/ST6/01370; Principal Investigator: prof. dr hab. Jacek Tabor

Deep Self-Organizing Neural Graphs

Neural networks and decision trees, see Figure 1, are two exceptionally effective machine learning computational models with a rich history of successful applications in artificial intelligence. In particular, they are used in classification and regression tasks. They typically come with mutually exclusive benefits and limitations:

- Neural networks outperform decision trees by jointly learning how to represent and classify data. However, they are widely opaque and suffer from a lack of transparency and interpretability.
- On the other hand, it is easy to show an explanation of a decision tree prediction. This is because it depends on a relatively short decision sequence. However, decision trees usually do not generalize as well as deep neural networks.

This motivates research into joining the positive aspects of both models.

While decision trees can increase the performance and interpretability of neural networks, they usually suffer from exponential growth with depth, repeating nodes, and suboptimal structure, often selected manually before training. Hence, more and more attention is put on combining neural networks with decision graphs instead of trees. Decision graphs present some advantages compared to decision trees. For instance, they have a flexible structure that allows for the creation of multiple paths from the root to each leaf. As a result, the nodes are reused, which solves the replication problem (see Figure 2). Moreover, they require substantially less memory while considerably improving the generalization. Nevertheless, decision graphs are not commonly used in deep learning due to the following reasons:

- Joining graphs with neural networks often requires a predefined tree/graph structure. This is in contrast to classical decision graphs, which allow for arbitrary architecture.
- The choice of the cost function, typically realized by the logistic regression, is suboptimal, comparing to other classification models, such as SVMs.
- Ensemble models can significantly improve the prediction of shallow approaches (such as decision trees) by averaging the predictions obtained from single, separately trained models. However, optimizing entire ensembles by neural networks often leads to very high computational costs.

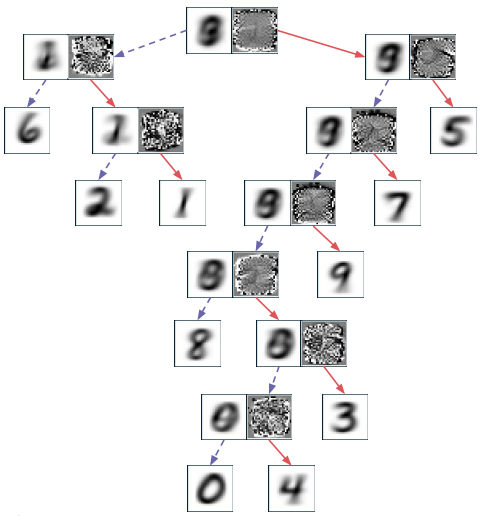

Figure 1:
Visualization of a shallow decision tree trained on the MNIST dataset. Each node represents the learned filter (right) and the "average" processed image (left) that passes through that node during the computation.

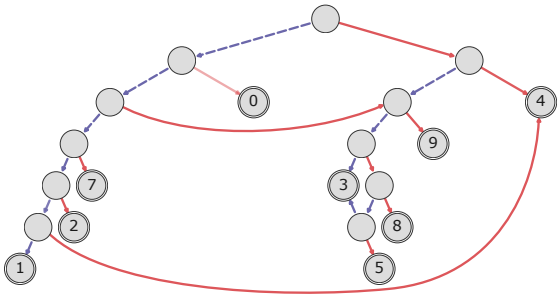

Figure 2:
Visualization of a decision graph trained on the MNIST dataset. Observe that some nodes are reused.

The goal of this project is to addresses the above shortcomings by introducing a Self-Organizing Neural Graph model — a specific type of decision graph efficiently trained by a neural network. The main aim lies in constructing a flexible graph structure that could allow for easy integration with a standard neural network structure. In particular, a Self-Organizing Neural Graph should have the ability to reconfigure its connections in a differentiable fashion so that one can obtain a final graph like the one presented in Figure 2.

Our main idea is to introduce "soft connections", where after applying the binary decision given in internal nodes, we can jump with trainable probability to all inner nodes.