

# Late Breaking Results: Conversion of Neural Networks into Logic Flows for Edge Computing

Daniel Stein<sup>1</sup>, Shaoyi Huang<sup>2</sup>, Rolf Drechsler<sup>3</sup>, Bing Li<sup>4</sup>, Grace Li Zhang<sup>1</sup>

<sup>1</sup>TU Darmstadt, <sup>2</sup>Stevens Institute of Technology, <sup>3</sup>University of Bremen, <sup>4</sup>TU Ilmenau

Email: {daniel.stein, grace.zhang@tu-darmstadt.de},

shuang59@stevens.edu, drechsler@uni-bremen.de, bing.li@tu-ilmenau.de

**Abstract**—Neural networks have been successfully applied in various resource-constrained edge devices, where usually central processing units (CPUs) instead of graphics processing units exist due to limited power availability. State-of-the-art research still focuses on efficiently executing enormous numbers of multiply-accumulate (MAC) operations. However, CPUs themselves are not good at executing such mathematical operations on a large scale, since they are more suited to execute control flow logic, i.e., computer algorithms. To enhance the computation efficiency of neural networks on CPUs, in this paper, we propose to convert them into logic flows for execution. Specifically, neural networks are first converted into equivalent decision trees, from which decision paths with constant leaves are then selected and compressed into logic flows. Such logic flows consist of *if* and *else* structures and a reduced number of MAC operations. Experimental results demonstrate that the latency can be reduced by up to 14.9 % on a simulated RISC-V CPU without any accuracy degradation. - The code is open source at <https://github.com/TUDa-HWAI/NN2Logic>

**Index Terms**—Neural network on CPUs; Logic flows of neural networks; Edge computing

## I. Introduction

Neural networks executed on edge devices can provide real-time AI processing closer to the data source. In neural networks, a huge number of multiply-accumulate (MAC) operations need to be executed. However, edge devices usually have general-purpose central processing units (CPUs) and do not have graphics processing units (GPUs) with parallel MAC circuits due to their power limitations. CPUs are not good at executing such mathematical operations on a large scale. The reason is that a CPU normally contains a large circuit block for logic execution, while only a limited number of dedicated multipliers and adders are available for executing MAC operations. Therefore, the execution efficiency of neural networks is compromised on CPUs, leading to high latency.

Various methods have been proposed to enhance the execution efficiency of neural networks on CPUs. At algorithm level, pruning [1, 2], quantization [3, 4], knowledge distillation [5], neural architecture search [6], and dynamic decisions [7, 8] have been applied to reduce the number of MAC operations. At hardware level, state-of-the-art AI compilers [9, 10, 11, 12] exploit hardware mapping, data reuse, memory allocation, and fetching [13], instruction scheduling to avoid bottlenecks [14], optimization for memory latency hiding, loop-oriented optimizations, and parallelization.

The techniques described above are effective in many scenarios. However, they still focus on executing many MAC operations instead of examining the logic expression of neural

This work is funded by the European Union - European Research Council (ERC) Starting Grant - Project-ID 101219243. Views and opinions expressed are those of the author(s).

networks. Even for heterogeneous CPU platforms extended with a neural processing unit, it is still crucial to examine the execution of neural networks on CPUs in a logic form to take advantage of all the available computation resources. This study is still missing in the state of the art to date.

In this paper, we explore converting the execution of neural networks into logic flows to enhance the execution efficiency of neural networks on CPUs. Such logic flows consist of *if* and *else* structures and a reduced number of MAC operations, so that the computation efficiency on CPUs can be enhanced.

## II. Methodology

To convert a neural network into logic flows for efficient CPU execution, we first convert it to an equivalent decision tree with the method in [15]. To reduce the complexity of the decision tree, only training data is used for the tree construction, as described in Section II-A. Afterwards, we select those decision paths, traversing from the root of the decision tree to leaf nodes with constant classification results. We then convert the execution of such paths into logic flows, as described in Section II-B. The remaining decision paths will be executed in the original format of the neural network, leading to a hybrid execution on CPUs in Section II-C.

Fig. 1 illustrates the concept of conversion, where a fully connected neural network in Fig. 1(a) is converted into its equivalent decision tree in Fig. 1(b). We use the extraction of the logic flow of the decision path to the right-most leaf, denoted as leaf L4, as an example. This leaf has a constant classification result of class  $c_1$ . An excerpt of C code for such a hybrid execution is shown in Fig. 1(c).

### A. Conversion into Decision Trees with Training Data

To construct the decision tree of a neural network, we build the decision nodes corresponding to the decision of the activation functions. Decision nodes are then connected via branches until we reach the last layer of the neural network. For example, the input condition  $x_0$  making  $y_0$  larger than 0 with ReLU as the activation function, i.e.,  $2x_0 - 1 > 0$ , is set as the root node of the tree. The input condition making  $y_1$  larger than 0, i.e.,  $3x_1 - 2 > 0$ , is then constructed as decision nodes, which are connected with the root node via two branches representing the true ( $T$ ) and false ( $F$ ) decisions of the root node. The remaining neurons can be processed similarly to create more decision nodes. We only add branches that are traversed by samples in the training dataset (e.g., the false branch of the right-most decision of  $y_1$  in Fig. 1(b) is never visited and thus not added). When processing the neurons of the last layer, we create leaf nodes, which are the final nodes of the decision tree.

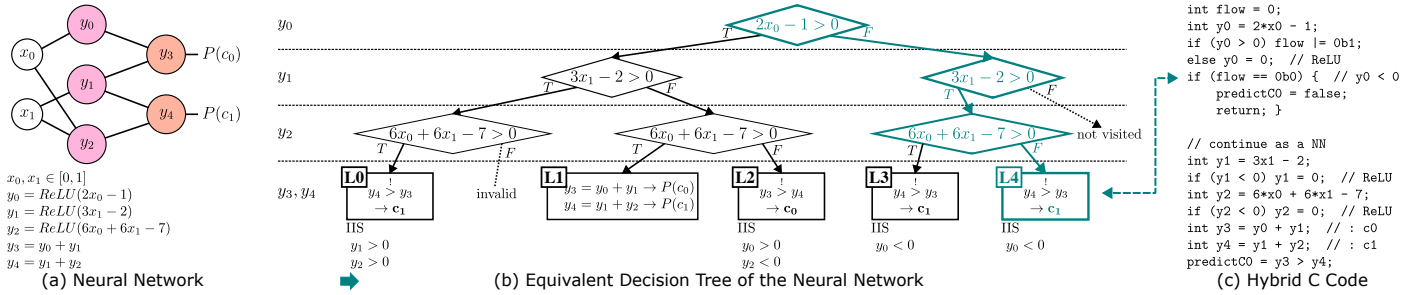


Fig. 1. Conversion of a neural network into an equivalent decision tree. (a) A neural network for binary classification (classes  $c_0$  and  $c_1$ ). (b) The decision tree of the neural network in (a) constructed using training data. (c) C Code for hybrid execution of the neural network and the logic flow of leaf L4.

TABLE I  
THE RUNTIME RESULTS OF OUR HYBRID EXECUTION.

Dataset	Network	Acc	Reference (Cycles)			Hybrid (Cycles)			Latency Comparison				
			Min	Avg	Max	Min	Avg	Max	Min	Avg	Max	Layer 2 + 3	Exit by Tree
MNIST*	25 - 30 - 2	0.937	40326	40367.08	40436	38387	40095.32	40496	4.8 %	0.7 %	-0.1 %	19.1 %	15.11 %
Occupancy I	10 - 25 - 2	0.862	1650	1665.07	1691	789	1416.63	1712	52.2 %	14.9 %	-1.2 %	39.3 %	34.1 %
Occupancy II	10 - 20 - 2	0.983	1292	1319.98	1329	963	1250.74	1337	25.5 %	5.2 %	-0.6 %	25.1 %	12.7 %

## B. Selection of Decision Paths for Extracting Logic Flows

After the construction of the decision tree, we will select those decision paths with constant leaves for logic flow extraction. A constant leaf in the case of a classification task means that one of the classes always appears, and the remaining classes never appear. For example, in Fig. 1, the leftmost leaf, denoted as leaf L0, has the constant classification result of class  $c_1$ , indicating that all the training samples lead to the same classification. To determine whether a decision path leads to a constant leaf for every possible input, we formulate a Mixed Integer Programming problem for every class on this path and try to solve it with Gurobi [16]. For example, for leaf L4 we model two MIP problems, one for the case  $y_4 > y_3$  (class  $c_1$ ) and the other for  $y_3 > y_4$  (class  $c_0$ ). As there is no feasible solution for the case  $y_3 > y_4$ , this leaf will always yield  $y_4$  larger than  $y_3$ , which means it always predicts class  $c_1$ . Afterwards, we use Gurobi to determine the Irreducible Infeasible Subsystem (IIS) of constant leaves, which is a set of constraints that cause the problem to be infeasible, and if any of the constraints would be removed, the problem becomes solvable [17]. For the example of leaf L4 in Fig. 1, the IIS is  $y_0 < 0$ , meaning class  $c_1$  is always predicted because of the constraint  $y_0 < 0$ . We then use only this constraint to describe the logic flow of this decision path, since the other decision nodes are not required.

## C. Hybrid Execution with Selected Decision Paths

A selected number of compressed paths with constant leaves are converted into logic flows. For the remaining paths, they are executed in the original format of the neural network, leading to a hybrid execution. Fig. 1(c) shows the C code for the hybrid execution when only the decision path leading to leaf L4 is converted into a logic flow. All the neurons corresponding to constraints in the IIS of leaf L4 are computed first. After all the required ReLU decisions are made, a tracking variable is checked to see if all the conditions of the logic flow are matched. If this is the case, the output is determined, and the computation has finished. After all logic flows have

been processed, the remaining code is added, which continues computation as a regular neural network.

## III. Experimental Results

To verify the effectiveness of our approach, we applied it to three quantized 3-layer fully connected neural networks with three publicly available datasets: MNIST<sup>1</sup> [18] and two room occupancy detection datasets, denoted as Occupancy I [19] and Occupancy II [20], as shown in Table I. Since we generate bare-metal C code in our framework, we need a comparable reference implementation of the neural network for evaluation without the overhead of an interpreter. Therefore, we also generate the Reference code, which has the same optimizations and ordering of neurons. All implementations are evaluated on the hardware simulator of the Ibex RISC-V CPU [21].

The results are shown in Table I. None of the three neural networks exhibits an accuracy degradation for the hybrid code compared with the reference (the model accuracies are shown in column “Acc”). Considering the relative latency comparison of the reference and hybrid execution, both the minimum latency and the average latency can be reduced effectively using logic flows. For example, the minimum latency can be reduced by up to 52.2%. For the maximum latency, there is a slight increase due to additional checks when executing the sample using the original form of the neural network. The latency reduction considering only the second and the third layers can be up to 39.3%. Finally, the last column in Table I demonstrates the percentage of samples in the inference dataset exiting through logic flows of the neural network.

## IV. Conclusion

In this paper, we have proposed to convert the execution of neural networks into logic flows for edge computing. It opens up a new dimension to accelerate neural networks on edge devices with CPUs. Future work includes extending the framework into deep neural networks with multi-class classification.

<sup>1</sup>MNIST is modified to determine if a given training sample is either even or odd, denoted as MNIST\*.

## References

- [1] Lixiang Han, Zhen Xiao, and Zhenjiang Li. “Dtmm: Deploying tinyml models on extremely weak iot devices with pruning”. In: *IEEE Conference on Computer Communications*. 2024.
- [2] Mengnan Jiang et al. “Class-Aware Pruning for Efficient Neural Networks”. In: *Design, Automation and Test in Europe Conference (DATE)*. 2024.
- [3] Shaojie Zhuo et al. “An empirical study of low precision quantization for TinyML”. In: *arXiv preprint arXiv:2203.05492* (2022).
- [4] Wenhao Sun et al. “Class-based Quantization for Neural Networks”. In: *Design, Automation and Test in Europe Conference (DATE)*. 2023.
- [5] Amir Moslemi et al. “A survey on knowledge distillation: Recent advancements”. In: *Machine Learning with Applications* 18 (2024), p. 100605.
- [6] Ji Lin et al. “MCUNet: Tiny deep learning on IoT devices”. In: *Advances in neural information processing systems*. 2020.
- [7] Jifeng Guo et al. “Dynamic Neural Network Structure: A Review for its Theories and Applications”. In: *IEEE Transactions on Neural Networks and Learning Systems* 36.3 (2025), pp. 4246–4266.
- [8] Jingcun Wang, Bing Li, and Grace Li Zhang. “Early-Exit with Class Exclusion for Efficient Inference of Neural Networks”. In: *International Conference on AI Circuits and Systems (AICAS)*. 2024.
- [9] *Apache TVM*. URL: <https://tvm.apache.org/>.
- [10] *IREE: Intermediate Representation Execution Environment*. URL: <https://github.com/iree-org/iree>.
- [11] *PyTorch Documentation*. URL: <https://docs.pytorch.org/docs/stable/torch.compiler.html>.
- [12] *Multi-Level Intermediate Representation*. URL: <https://mlir.llvm.org/>.
- [13] Ji Lin et al. “Memory-efficient patch-based inference for tiny deep learning”. In: *Advances in Neural Information Processing Systems*. 2021.
- [14] Moritz Scherer et al. “DeepDeploy: Enabling Energy-Efficient Deployment of Small Language Models on Heterogeneous Microcontrollers”. In: *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD)* 43.11 (2024), pp. 4009–4020.
- [15] Çağlar AYTEKİN. “Neural Networks are Decision Trees”. In: *ArXiv* (2022).
- [16] Gurobi Optimization, LLC. *Gurobi Optimizer Reference Manual*. 2025. URL: <https://www.gurobi.com>.
- [17] John Gleeson and Jennifer Ryan. “Identifying Minimally Infeasible Subsystems of Inequalities”. In: *ORSA Journal on Computing* 2.1 (1990), pp. 61–63.
- [18] Li Deng. “The MINIST database of handwritten digit images for machine learning research”. In: *IEEE Signal Processing Magazine* 29.6 (2012), pp. 141–142.
- [19] Kamilla Heimar Andersen et al. *Room-based Indoor Environment Measurements and Occupancy Ground Truth Datasets from Five Residential Apartments in a Nordic Climate*. Aalborg University, 2023.
- [20] Adarsh Pal Singh et al. “Machine learning-based occupancy estimation using multivariate sensor nodes”. In: *IEEE Globecom Workshops (GC Wkshps)*. 2018.
- [21] lowRISC. *Ibex RISC-V Core*. URL: <https://github.com/lowRISC/ibex>.