

# FPGA Resource-aware Structured Pruning for Real-Time Neural Networks

Benjamin Ramhorst\*

Electrical and Electronic Engineering  
Imperial College London  
London, United Kingdom  
benjamin.ramhorst@inf.ethz.ch

Vladimir Lončar†

Laboratory for Nuclear Science  
Massachusetts Institute of Technology  
Cambridge, United States of America  
vloncar@mit.edu

George A. Constantinides

Electrical and Electronic Engineering  
Imperial College London  
London, United Kingdom  
g.constantinides@imperial.ac.uk

**Abstract**—Neural networks achieve state-of-the-art performance in image classification, speech recognition, scientific analysis and many more application areas. Due to the high computational complexity and memory footprint of neural networks, various compression techniques, such as pruning and quantization, have been proposed in literature. Pruning sparsifies a neural network, reducing the number of multiplications and memory. However, pruning often fails to capture properties of the underlying hardware, causing unstructured sparsity and load-balance inefficiency, thus bottlenecking resource improvements. We propose a hardware-centric formulation of pruning, by formulating it as a knapsack problem with resource-aware tensor structures. Evaluated on a range of tasks, including sub-microsecond particle classification at CERN’s Large Hadron Collider and fast image classification, the proposed method achieves reductions ranging between 55% and 92% in the DSP utilization and up to 81% in BRAM utilization.

**Index Terms**—FPGA, Deep Learning, Pruning

## I. INTRODUCTION

Modern neural networks are associated with high computational complexity and memory footprint, linked to the high dimensionality of weight tensors. With the ever-increasing need for faster computation and lower power consumption, driven by real-time systems and Internet-of-Things (IoT), neural networks must be adapted to reduce their complexity. A common method for network compression is pruning, the process of sparsifying a neural network by setting some of its weights to zero. In their seminal work, Han *et al.* [1] applied iterative, magnitude-based unstructured pruning to modern neural networks. The networks were trained with regularization, shifting weights towards zero. Weights below a certain magnitude were pruned, while the remaining weights were updated, thus capturing the loss of the pruned weights. While unstructured pruning achieves high sparsities, little improvements in resource utilization and inference latency are observed [2], due to the overhead of sparse encodings.

hls4ml [3], [4] is an open-source library for real-time inference of neural networks on FPGAs. Originally intended for sub-microsecond data filtering in high-energy physics,

This work is supported by Engineering and Physical Sciences Research Council (EPSRC) grant EP/S030069/1 and NSF Institute A3D3, NSF 2117997.

\* Currently at ETH Zürich, Switzerland

† Also at Institute of Physics Belgrade, Serbia




Fig. 1. Variations of RF and the impact on resource utilization [4].

hls4ml has become a versatile tool for real-time inference, with a full on-chip design and high parallelism, frequently leading to high resource utilization. A key variable in hls4ml is the *Reuse Factor* (RF), determining the trade-off between latency and resource utilization, as shown in Fig. 1. In hls4ml, matrix-vector multiplication is the baseline for fully connected layers, as well as convolutional layers, implemented using the *im2col* transformation.

We propose an FPGA resource-aware structured pruning algorithm, able to capture, at training time, the underlying mapping of network weights to digital signal processing blocks (DSPs) and block random access memory (BRAM). Groups of weights are iteratively removed through constrained optimization, formulated as a knapsack problem. The methodology described in this research has been integrated with hls4ml and open-sourced with an Apache 2.0 license<sup>1</sup>.

## II. RESOURCE-AWARE PRUNING

To maximize resource savings of accelerated neural networks, we consider the optimization of DSP and BRAM utilization, by pruning all the weights processed by the same multiplier or block of memory. We refer to these groups of weights as resource-aware vectors of weights. The mapping between weights and DSP blocks is deterministic from the RF. Given a weight matrix, weights processed by the same multiplier can be obtained by transposing and flattening the matrix into a vector. The vector is split into sub-vectors

<sup>1</sup>GitHub page: <https://github.com/fastmachinelearning/hls4ml/tree/hardware-aware-pruning>

of length RF, with each sub-vector representing the weights processed by the same DSP block. Weights processed by the same DSP are stored as subsequent words in the same block of memory, typically implemented as 1K x 36. Weights are quantized before deployment to a lower precision, so each block of RAM stores weights processed by several neighboring DSP blocks. As an example, groups of weights quantized to 18 bits would map to 2 DSP blocks for each block of RAM.

hls4ml supports heterogeneous configuration of layers, allowing fine-grained tuning of the hardware-configuration, by varying the RF and precision. Furthermore, Vivado aims to implement multiplications for precisions lower than 10 bits through LUTs, rather than DSPs. As a result of these factors, pruning becomes a non-trivial optimization problem, aimed at minimizing network loss and overall resource consumption. We solve this problem by transforming it to a knapsack problem, similar to the work by Shen *et al.* [5] on latency-aware pruning for GPUs. Given a set of resource-aware vectors  $\mathbb{W} = \{\mathbf{w}_1, \dots, \mathbf{w}_n\}$  and the vector-valued function  $R: \mathbb{R}^k \rightarrow \mathbb{R}^m$ , representing the resource consumption (DSP, BRAM) of a group of weights, we formulate pruning as an optimization problem (Eq. 1) of selecting a subset of weights,  $\hat{\mathbb{W}} \subseteq \mathbb{W}$ , minimizing network loss while maintaining a resource budget,  $\mathbf{c}$ . To minimize network loss, we aim to maximize the normalized<sup>2</sup> magnitude of the selected weights, similar to magnitude-based pruning [1].

$$\begin{aligned} \max_{\hat{\mathbb{W}}} \quad & \sum_{\mathbf{w}_i \in \hat{\mathbb{W}}} \frac{\|\mathbf{w}_i\|_1}{\max_{j \text{ s.t. } j \cong i} \|\mathbf{w}_j\|_1} \\ \text{subject to} \quad & \sum_{\mathbf{w}_i \in \hat{\mathbb{W}}} R(\mathbf{w}_i) \preceq \mathbf{c} \\ & \hat{\mathbb{W}} \subseteq \mathbb{W} \end{aligned} \quad (1)$$

To solve the combinatorial optimization problem in Equation 1, we consider the 0-1 knapsack problem:

$$\begin{aligned} \max_{\mathbf{x}} \quad & \mathbf{v}^T \mathbf{x} \\ \text{subject to} \quad & \mathbf{U} \mathbf{x} \preceq \mathbf{c} \\ & \mathbf{x} \in \{0, 1\}^n \end{aligned} \quad (2)$$

In our case, the vector  $\mathbf{v}$  corresponds to the normalized magnitudes, while the matrix  $\mathbf{U}$  corresponds to the resource utilization. We find that the knapsack problem can be efficiently solved using branch-and-cut methods and the solution is such that:

$$\mathbf{x}_i = \begin{cases} 0 & \text{if } \mathbf{w}_i \text{ pruned} \\ 1 & \text{otherwise} \end{cases} \quad (3)$$

During training we include a resource-aware regularization loss, shifting weights sharing the same hardware resource towards zero. Pruning is implemented iteratively, such that the resource budget,  $\mathbf{c}$ , is repeatedly reduced, until the target resource consumption is achieved, or, network performance on

<sup>2</sup>Vector norms are normalized in every layer, avoiding bias to highly-dimensional layers.

TABLE I  
EFFECTS OF RESOURCE-AWARE PRUNING.

Model	Jets	SVHN	Fashion MNIST
Baseline accuracy [%]	76.61	91.95	89.72
RF	2, 4, 8, 16	3, 9, 27	Fixed per layer
DSP reduction	5.8x - 12.2x	2.2x - 3.9x	4.7x
BRAM reduction	2.3x - 5.2x	1.4x	2.1x
$\Delta$ accuracy [pp]	-0.63	+0.26	+0.02

the validation set drops below the relative tolerance. At each iteration, the remaining weights are retrained.

### III. RESULTS

We evaluate the effects of resource-aware pruning on three benchmark models: high-energy jet tagging [4], SVHN classification, as described by Aarrestad *et al.* [6] and Fashion MNIST classification [7] with a LeNet-like [8] architecture. The models are trained using Keras [9] with a TensorFlow backend [10], by minimizing categorical cross-entropy and synthesized using Vivado, Vivado HLS 2020.1 and hls4ml 0.7.1. The resource utilization is reported post place and route with the target device set to Virtex UltraScale+ XCVU9P. The results are reported in Table I. For particle and SVHN classification, we vary the RF, and report a range for DSP and BRAM reduction. Due to the size of LeNet, we manually tune the RF of every layer in Fashion MNIST classification, in order to meet Vivado partition and unroll limits. We also report the difference in accuracy between the pruned model with the lowest accuracy and the baseline model.

### REFERENCES

- [1] S. Han, J. Pool, J. Tran, and W. J. Dally, "Learning both weights and connections for efficient neural network," in *NIPS*, 2015.
- [2] T. Hoefler, D. Alistarh, T. Ben-Nun, N. Dryden, and A. Peste, "Sparsity in deep learning: Pruning and growth for efficient inference and training in neural networks," *J. Mach. Learn. Res.*, vol. 22, pp. 241:1–241:124, 2021.
- [3] FastML Team, "fastmachinelearning/hls4ml," 2023. [Online]. Available: <https://github.com/fastmachinelearning/hls4ml>
- [4] J. Duarte, S. Han, P. Harris, S. Jindariani, E. Kreinar, B. Kreis *et al.*, "Fast inference of deep neural networks in FPGAs for particle physics," *Journal of Instrumentation*, vol. 13, no. 07, p. P07027, Jul 2018.
- [5] M. Shen, H. Yin, P. Molchanov, L. Mao, J. Liu, and J. M. Alvarez, "Structural pruning via latency-saliency knapsack," in *Advances in Neural Information Processing Systems*, vol. 35, 2022, pp. 12 894–12 908.
- [6] T. Aarrestad, V. Loncar, N. Ghielmetti, M. Pierini, S. Summers, J. Ngadiuba *et al.*, "Fast convolutional neural networks on FPGAs with hls4ml," *Machine Learning: Science and Technology*, vol. 2, no. 4, p. 045015, Jul 2021.
- [7] H. Xiao, K. Rasul, and R. Vollgraf. (2017) Fashion-MNIST: a Novel Image Dataset for Benchmarking Machine Learning Algorithms.
- [8] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proc. IEEE*, vol. 86, pp. 2278–2324, 1998.
- [9] F. Chollet *et al.*, "Keras," <https://keras.io>, 2015.
- [10] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro *et al.*, "TensorFlow: Large-scale machine learning on heterogeneous systems," 2015, software available from tensorflow.org. [Online]. Available: <https://www.tensorflow.org/>