

PAODING: A High-fidelity Data-free Pruning Toolkit for Debloating Pre-trained Neural Networks

Anonymous Authors

Abstract—We present PAODING, a toolkit to debloat pre-trained neural network models through the lens of *data-free pruning*. To preserve the model fidelity, PAODING adopts an iterative process, which dynamically measures the effect of deleting a neuron to identify candidates that have the least impact to the output layer. Our evaluation shows that PAODING can significantly reduce the model size, generalize on different datasets and models, and meanwhile preserve the model fidelity in terms of test accuracy and adversarial robustness. PAODING is publicly available on PyPI via <https://pypi.org/project/paoding-dl>.

I. INTRODUCTION

Nowadays, people tend to train large neural network models containing many hidden layers for higher accuracy and better generalization. However, the resulted models may be over-parameterized and computationally intensive, making it challenging to deploy models on devices with limited hardware capabilities, including Internet-of-Things (IoT) and mobile devices. *Neural network pruning* is a typical technique to mitigate this challenge by *debloating* a large model, which can remove the redundant parameters without significant performance loss. Most existing pruning approaches [1]–[3] are proposed based on the assumption that the pruned model has a chance to be re-trained or fine-tuned. However, the training data may not be available due to the protection of privacy or intellectual property. In this work, we develop PAODING to achieve *data-free pruning* [4]–[6]. Our pruning technique is designed from the model users’ perspective, and gets rid of the re-training or fine-tuning process that entails the original dataset. Upon receiving a pre-trained model, users can perform on-demand pruning with our approach to ease the deployment on a wide range of devices.

II. OVERVIEW OF PAODING

Due to the absence of re-training that can fix the mis-pruned neurons¹, existing pruning that adopts an aggressive cut-and-re-train strategy can no longer be applied in data-free pruning. For that reason, the design of PAODING aims to preserve model fidelity, in terms of test accuracy and robustness against undesirable inputs.

Figure 1 shows the workflow of PAODING. It begins with reading a pre-trained model and loading its parameters and configurations. Then it iterates all the convolutional (Conv2D) and FC (dense) hidden layers and performs pruning in three stages.

¹We use the term “neurons” to indicate the units/nodes within a layer. We also use “hidden units” and “channels” to more specifically refer to the neurons in fully-connected and convolutional layers.

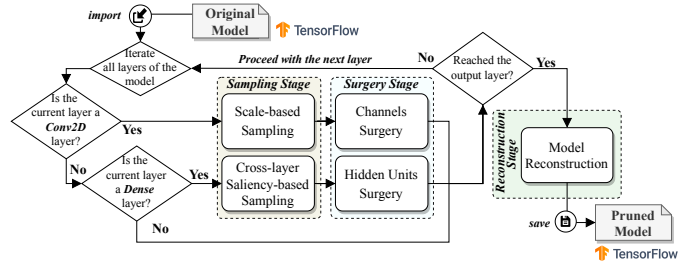


Fig. 1. The workflow of PAODING pruning

- First, it traverses all neurons of that layer and samples the pruning candidates based on specific strategies (refer to the next section) according to the layer type.
- Given a list of candidate neurons nominated, PAODING performs pruning. It adopts a *surgery* approach that deletes the candidate neurons from the model. PAODING shrinks the Conv2D layers by deleting some *channels* for every filter. For dense layers, PAODING directly cuts off the entire *hidden unit* including bias and all its connections to the previous and next layers.
- Once the pruning is accomplished, PAODING reconstructs the model with the pruned layers and thus, the pruned model has a shrunk structure.

III. SAMPLING STRATEGY

PAODING aims to minimize the impact on the model’s output for the purpose of preserving the fidelity of the original pre-trained model. To this end, we apply different strategies for the two types of layers.

Conv2D Layers. Our approach adopts a scale-based sampling strategy for Conv2D layers, for the purpose of prioritizing the least *salient* channels during the pruning. To this end, we traverse all convolutional layers and calculate the L_1 -norm of filters within each channel. The sum of L_1 -norm values of all the filters within a channel is assessed as the *channel scale*. Next, PAODING prioritizes the channels to be pruned by sorting their channel scales in ascending order until the target of pruning on that layer has been reached.

Dense Layers. We propose a *pair-wise* pruning mechanism for neurons (hidden units) of dense layers, based on an assumption that *a pruning approach that produces the least impact to the outputs best preserves the fidelity of the original model*. We use $\langle a_i^l, a_j^l \rangle$ to denote a neuron pair (indexed as i and j) in the l -th layer. These two neurons are supposed to be similar. Therefore, one of them is considered non-salient and can be replaced by the other neuron by adding all its parameters [4].




Fig. 2. Evaluation of model size, test accuracy, and robustness of models debloated by PAODING

Next, we adopt the state-of-the-art data-free pruning approach designed for Dense layers [6] to find the best candidate neuron pair to be pruned.

To simply put, the *impact* of pruning a neuron pair $\langle a_i^l, a_j^l \rangle$ is calculated based on the simulated propagation of removing the pruned neuron, i.e., a_i^l , starting from layer l until the output layer. The sampling criterion of a neuron pair is jointly determined by two metrics, namely the L_1 -norm and the Shannon’s entropy of the impact. The former calculates the sum of change of all output nodes caused by pruning a neuron pair, and the latter measures the *uniformness* of the pruning impact made on all output nodes, stipulating a higher degree of uniformness would less likely to make the pruned model misclassify. Neuron pairs with small values in both the metrics will be given priority for the pruning.

IV. EVALUATION

Implementation and Availability. We implement PAODING using Python v3.9 and evaluate it with neural network models on TensorFlow v2.4.1. It can be accessed through PyPI with the key “paoding-dl”. At the moment of writing, PAODING has over 12 thousands downloads on PyPI.

PAODING accepts any legitimate format of CNN and MLP models trained by TensorFlow and is compatible with various further optimization techniques such as quantization. It also allows the user to configure the pruning target and the batch size of pruning per step. Given a trained model, it automatically identifies the dense and Conv2D layers, prunes the neurons from them, and stops once the pruning target has been reached.

Evaluation Setup. We evaluate our approach on four neural network applications. The first model is a small MLP model for the ULB credit card fraud dataset. The remaining three models are CNNs trained with the Brain Tumor MRI, MNIST, and CIFAR-10 datasets.

We perform pruning progressively in multiple epochs up to 50% of its neurons have been pruned. At each epoch, we only prune 5% of neurons in eligible hidden layers. After each pruning epoch, we record the file size of the pruned model, test accuracy, and the robustness² against FGSM adversary.

Effectiveness of PAODING. PAODING shows effective in debloating all four tested models. As shown in Figure 2 (a), a 25% pruning can remove up to 49.4% (model #1) of parameters in our tested models. By pruning another 25% of

²Only models #3-4 are evaluated because their datasets have been widely studied on the subject of robustness.

neurons in these four models, our approach can remove 66.7% of parameters in average (63.9%-78.1%), which is equivalent to 2.5x-4.5x of shrinking.

In addition to compressing model size, PAODING also aims to preserve model fidelity in terms of prediction accuracy and robustness against undesirable inputs. According to Figure 2 (b), we find that all four models show promising accuracy even after 25% of neurons have been pruned. This, by jointly interpreting the decay of model size, implies that we can reduce around 40% of their parameters (36.5%-49.4%) without incurring a significant sacrifice of test accuracy. We also observe less than 50% of accuracy decay on all tested models even after pruning 50% of their neurons.

We also assess the robustness against adversarial input perturbations (see Figure 2 (c)). Although the tested models are not specially trained for robustness, our approach can still preserve the robustness of the pruned model and ensure the robustness declines at a controllable pace during the pruning. Both models maintain almost the same robustness until 20% of its neurons have been pruned, and preserve 50% of its original robustness even after 50% of its parameters have been pruned.

Demonstration. An online Python notebook demonstration is hosted on Google Colab. The script can be accessed through the link <https://tinyurl.com/mr47vrnr>. The demonstration begins with installing PAODING through pip repository, followed by training a sample neural network from scratch, configuring PAODING settings, and pruning the trained model.

V. CONCLUSION

In this work, we propose PAODING, a pre-trained neural network model pruning toolkit in a data-free context. We implement our pruning as a conservative progressive process to preserve the fidelity of the model after pruning. Our evaluation shows that PAODING can significantly shrink the model size by up to 4.5x by applying our pruning approach, at the cost of losing less than 50% of its original accuracy and robustness on all four tested models. Our toolkit is made publicly available and would encourage future research on exploring more optimization techniques for neural networks.

REFERENCES

- [1] N. Lee, T. Ajanthan, and P. Torr, “SNIP: Single-shot network pruning based on connection sensitivity,” in *6th International Conference on Learning Representations, ICLR*, 2018.
- [2] H. Li, A. Kadav, I. Durdanovic, H. Samet, and H. P. Graf, “Pruning filters for efficient convnets,” in *5th International Conference on Learning Representations, ICLR*, 2017.
- [3] Z. Liu, M. Sun, T. Zhou, G. Huang, and T. Darrell, “Rethinking the value of network pruning,” in *7th International Conference on Learning Representations, ICLR*, 2019.
- [4] S. Srinivas and R. V. Babu, “Data-free parameter pruning for deep neural networks,” in *Proceedings of the British Machine Vision Conference 2015*. BMVA Press, 2015, pp. 31.1–31.12.
- [5] H. Tanaka, D. Kunin, D. L. Yamins, and S. Ganguli, “Pruning neural networks without any data by iteratively conserving synaptic flow,” in *Proceedings of the 34th International Conference on Neural Information Processing Systems*, 2020, pp. 6377–6389.
- [6] M. H. Meng, G. Bai, S. G. Teo, and J. S. Dong, “Supervised robustness-preserving data-free neural network pruning,” in *The 27th International Conference on Engineering of Complex Computer Systems.*, 2023.