






CPrune: Compiler-Informed Model Pruning for Efficient Target-Aware DNN Execution

Taeho Kim¹, Yongin Kwon², Jemin Lee², Taeho Kim², and Sangtae Ha¹

¹ University of Colorado Boulder

{taeho.kim, sangtae.ha}@colorado.edu

² Electronics and Telecommunications Research Institute

{yongin.kwon, leejaymin, taehokim}@etri.re.kr

Abstract. Mobile devices run deep learning models for various purposes, such as image classification and speech recognition. Due to the resource constraints of mobile devices, researchers have focused on either making a lightweight deep neural network (DNN) model using model pruning or generating an efficient code using compiler optimization. Surprisingly, we found that the straightforward integration between model compression and compiler auto-tuning often does not produce the most efficient model for a target device. We propose CPrune, a compiler-informed model pruning for efficient target-aware DNN execution to support an application with a required target accuracy. CPrune makes a lightweight DNN model through informed pruning based on the structural information of subgraphs built during the compiler tuning process. Our experimental results show that CPrune increases the DNN execution speed up to $2.73\times$ compared to the state-of-the-art TVM auto-tune while satisfying the accuracy requirement.

1 Introduction

Deep neural networks (DNNs) have become increasingly popular for various applications like image processing and speech recognition. DNN requires heavy computation due to large and complex neural networks, so recent approaches proposed intelligent offloading using the cloud for resource-constrained mobile devices. Unfortunately, they often suffer from network overhead and disruptions, necessitating running DNNs directly on devices. Furthermore, making DNN inferences on devices is often desirable for user experience and privacy.

Several efforts have focused on improving DNN inference on mobile devices by considering energy efficiency, accuracy, and execution time. These efforts are primarily divided into two approaches: (1) making a lightweight DNN model using model compression and (2) generating optimized code for the target device

² Yongin Kwon is the corresponding author.




Fig. 1: An experiment to find the fastest model whose accuracy is higher than 92.80% accuracy for the target device among various pruned models of VGG-16 for CIFAR-10. We use the number of processed figures per second (FPS), given that indirect metrics like FLOPs cannot replace actual execution speed. The best-pruned model does not guarantee the best after compiler optimization. For example, the best model with pruning achieves only 2174 FPS while the suboptimal one obtains a higher 2857 FPS after compiler optimization.

using a DNN compiler. The model compression, for example, optimizes the neural networks by pruning, quantization, or neural architecture search (NAS). In particular, model pruning reduces the size of a model by removing non-critical or redundant neurons from a DNN model [26], while quantization does it by reducing the precision of the datatype [23]. In addition, NAS optimizes the structure of a DNN model for maximizing performance [19]. The DNN compiler, on the other hand, takes a DNN model and generates an optimized execution code for each target device.

Efficient execution through a lightweight model is critical for many applications, such as AR/VR applications, autonomous driving [24], drone or robot control [8], medical health monitoring [16], malware detection [40], and user authentication such as Face ID [50, 27]. While one might think that the straightforward integration between model compression and compiler optimization could generate the most efficient code for a specific target device, this is not the case. Figure 1 shows experimental results of direct integration between model compression and compiler optimization. As shown, the fastest model meeting the required accuracy found by pruning is often not the best model after compiler optimization, demanding joint optimization between them.

Based on this observation, we propose CPrune, a new pruning technique for an efficient DNN execution for target devices that jointly considers model compression and compiler optimization. Instead of optimizing DNN compression and compiler optimization independently, leading to suboptimal optimization, CPrune exploits the information extracted during the compiler optimization process. In particular, CPrune uses the subgraph structures of a neural network and their execution times on the target device during compiler optimization to

create an efficient target-aware pruned model fulfilling accuracy requirements. To the best of our knowledge, this is the first work using the information collected during the compiler optimization to create the most efficient target-aware compressed model meeting accuracy requirements.

Our contributions are:

- We report from our experiments that the fastest model that meets the required accuracy found by pruning is often not the best model after compiler optimization, necessitating joint optimization.
- We propose CPrune, which incorporates the information extracted during the compiler optimization process into creating a target-oriented compressed model fulfilling accuracy requirements. This information also significantly reduces the search space for parameter tuning such that CPrune can make a compressed model substantially faster than NetAdapt [44], the state-of-the-art hardware-aware model compression framework.
- Our experimental results show that CPrune achieves target-oriented performance improvement and increases the speed up to $2.73\times$ compared to the de facto DNN compiler framework TVM auto-tune while satisfying the accuracy requirement.

While CPrune’s approach is generally applicable, we can utilize CPrune for computer vision tasks requiring fast processing but a slight reduction in accuracy is acceptable, such as object detection and image classification for autonomous driving and biosignal image processing for seizure detection on a mobile device. We implement CPrune on top of an open deep learning compiler stack Apache TVM [1] and Microsoft NNI [2]. Our source code can be found at <https://github.com/taehokim20/CPrune>.

2 Related Works

2.1 Model Compression

Model compression reduces the model size and improves speed while meeting the performance requirements: pruning compresses a model while NAS creates a much lighter model. There are also approaches using the performance measurements on real hardware as feedback for model compression.

Pruning. Pruning removes non-critical, redundant neurons from a DNN model [26]. Figure 2 shows an example of a convolution operation in DNN. c_i is the number of input channels, each of which comprises height (h_i) and width (w_i) of the input. After convolving c_i input channels with filters of layer l_i , the input becomes abstracted to a feature map with output channels c_{i+1} , where c_{i+1} is determined by the number of filters in layer l_i . Structured pruning [14,15,21,39,48] removes filters or channels according to various algorithms to judge the redundancy of weights. For example, removing a few input channels c_i reduces the overall sizes of the input and filter shape, and pruning a filter reduces the sizes because it eliminates a channel in the next layer (both cases are highlighted as




Fig. 2: Structured pruning removes channels or filters to reduce the model’s overall size. In contrast, non-structured pruning selectively prunes less important weights to minimize the sizes without impacting accuracy loss.

shaded regions in Figure 2). On the other hand, non-structured pruning [11,47] selects individual weights to prune. As a result, they achieve very high compression without much accuracy loss even though it is hard to gain as much as speedups of structured ones [9,10,28].

NAS. NAS optimizes the structure of a DNN model for maximizing performance [7]. It uses various search strategies to explore the search space of neural architectures, such as random search, Bayesian optimization [35,41], evolutionary method [30,6], reinforcement learning (RL) [51], and gradient-based methods [7,25].

Hardware-aware model compression. Recently, several approaches have utilized performance measurements on real hardware to make a compressed model [22,37,42,44,45]. Especially, NetAdapt [44] points out that indirect metrics (e.g., the number of MACs or weights) may not necessarily reduce the direct metrics (e.g., execution time and energy consumption) [43,46], so they incorporate performance metrics of target hardware into its pruning adaptation algorithm.

Note that our approach CPrune is different from NetAdapt in that it jointly considers model compression and compiler optimization.

2.2 DNN Compilers with Auto-Tuning

Various DNN compilers generate code for DNN models on different hardware architectures [4,20,33]. For example, TVM [4] translates the input DNN models into the intermediate representation (IR) during the compilation process [32]. After performing hardware-independent optimizations with the IR, it translates high-level IR into low-level IR, which provides interfaces to tune the computation and memory access. In addition, auto-tuning approaches such as AutoTVM [5] and AutoScheduler [49] optimize various low-level hardware-dependent parameters since the optimal parameters vary depending on the DNN operations and the target hardware.

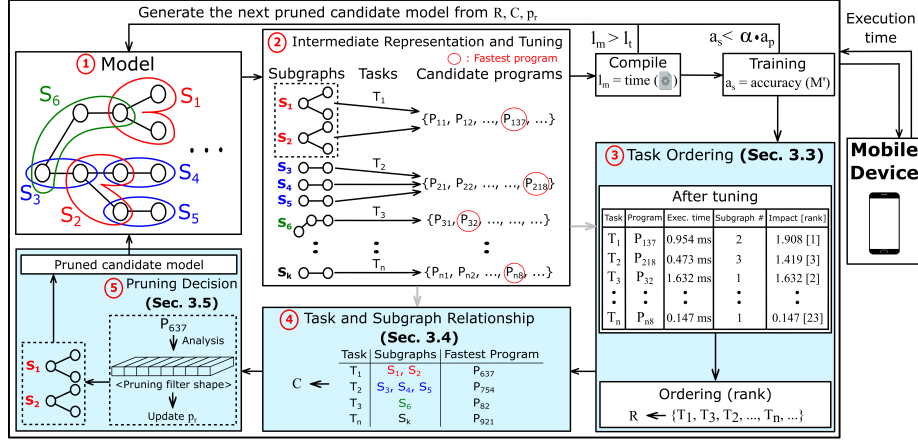


Fig. 3: CPrune overview. CPrune generates the fastest DNN execution model on the target device by using the structural information extracted during compiler optimization and the on-device performance of a DNN model.

3 CPrune: Compiler-Informed Model Pruning

Our motivation starts from a question: *Can an optimally pruned model guarantee the best performance after compiler optimization for a specific device?* To check this question, we prune VGG-16 [36] randomly and create 20 different models. We use the CIFAR-10 [17] dataset, and the original model’s accuracy is 93.29%. We check if the model showing the highest number of processed figures per second (FPS) with an accuracy higher than 92.80% would show the highest FPS after compiler optimization. We use FPS, given that indirect metrics like FLOPs cannot replace actual execution speed. For this experiment, we run and measure the FPS on a PC with an NVIDIA RTX-3080 GPU.

As shown in Figure 1, when we measure the FPS of each pruned model, the fastest model that meets the accuracy requirement is model A (92.85%, 205 FPS). However, after compiler optimization on the pruned models, the fastest model is model B (92.86%, 2857 FPS). It is different from what we expect: the most efficient model after pruning should be the most efficient one after compiler optimization. We also find no strong correlation between the performance of pruned models before and after compiler optimization, emphasizing the need for joint optimization between model compression and compiler optimization.

Inspired by this finding, we design a new pruning scheme, CPrune, that generates an efficient DNN execution model by pruning a model based on the extracted information during compiler optimization for the target device. Below we present the details of our approach.

3.1 CPrune Overview

The goal of CPrune is to find and prune neurons that primarily impact the execution of the DNN model on the target device while meeting the accuracy requirements. This informed pruning effectively prevents the issue of the best pruning model not being the best on the target device after compilation.

Figure 3 depicts how CPrune leverages the information collected during compiler optimization to prune neurons from the DNN model effectively. The shaded boxes indicate what CPrune adds to the existing compiler optimization framework. A DNN model consists of many convolutional layers, each representing a subgraph (①). A subgraph is assigned to a task, and two different subgraphs can point to the same task. Then, a DNN compiler creates numerous intermediate representations for each task and selects the fastest program on the target device (②). After compiling the aggregate of the fastest programs comprising a DNN model, CPrune checks if the model meets the minimum execution and accuracy requirements. After that, CPrune sorts tasks based on their execution times in descending order to find the most efficient model with further pruning. Since pruning a task with a longer execution time could significantly reduce a model’s execution time, CPrune selects the most time-consuming task as a candidate for further pruning (③). CPrune now needs to know which subgraph(s) is associated with this task as pruning candidates. CPrune also store the fastest program for each task. For this purpose, CPrune builds a table keeping the relationship among tasks, subgraphs, and programs (④). Finally, CPrune prunes subgraphs of the selected task while ensuring their code structures follow the structure of the fastest program of that task (⑤). Since the computation structure impacts the execution time, preserving the same computation structure after pruning is critical for efficient pruning. This process continues to make the most efficient pruned DNN model satisfying the accuracy requirement.

In the following subsections, we describe the details of an algorithm (Section 3.2), task ordering (Section 3.3), task and subgraph relationship table (Section 3.4) and pruning decision (Section 3.5).

3.2 CPrune Algorithm

CPrune prunes a DNN model gradually and iteratively. In each iteration, CPrune analyzes subgraphs of the model and creates relationships between subgraphs and each task. Then, it takes a pre-trained model M and the minimum accuracy requirement a_g and returns the efficient target-aware DNN executable program, as shown in Algorithm 1. The computational complexity of Algorithm 1 and a table that summarizes variables used are in the Supplementary Materials.

Initialization step (Line 1). It initializes tuning-related parameters, including the pruning rate of each subgraph p_r , the target execution time of the following iteration l_t , and the short-term accuracy of the previous best model a_p in Algorithm 1. It also initializes a task/subgraph table C , storing the relationship among tasks, subgraphs, and fastest programs and the list of tasks prioritized by tuning R .

Algorithm 1 CPrune Algorithm**Input:** Pre-trained model M and a_g **Output:** An efficient DNN executable file of M

```

1: Tune  $M$  and initialize  $p_r, l_t, a_p, C$  and  $R$ 
2: while  $a_p > a_g$  and  $R \neq \{\}$  do
3:   for  $r$  in  $R$  do
4:     For  $r$ , obtain associated subgraphs ( $S$ ) and the fastest program ( $P$ ) from  $C$ 
5:     Update pruning rate ( $p_r$ ) by analyzing the structure of  $P$ 
6:     Create a pruned candidate model ( $M'$ ) by pruning  $S$  with pruning rate  $p_r$ 
7:     Extract tasks from  $M'$  and create a candidate task/subgraph table ( $C'$ )
8:     Tune these tasks and create a candidate list of prioritized tasks ( $R'$ )
9:     Compile the aggregate of tuned tasks and measure  $l_m$ 
10:    Go line 3 if  $l_m \geq l_t$ 
11:    Short-term train  $M'$  and measure  $a_s$ 
12:    Remove  $r$  from  $R$  and go line 3 if  $a_s < \alpha \cdot a_p$ 
13:     $M \leftarrow M', R \leftarrow R', C \leftarrow C'$ , update  $p_r, l_t = \beta \cdot l_m, a_p = a_s$ 
14:    break
15:  end for
16: end while
17: Final long-term training, tuning and compilation for  $M$ 

```

Main step: pruning based on the compiler optimization (Line 2-16). In each pruning iteration (the while loop in Algorithm 1) of this step, CPrune generates a pruned candidate model (① in Figure 3) and moves to the intermediate representation (IR) and tuning stage (②). When a task in order of R is picked up (③), CPrune extracts associated subgraphs and the fastest program of the task from C (④). It decides the number of filters to prune by analyzing the arrangement of filters of the program and prunes filters of the subgraphs to create a pruned candidate model M' (⑤, details in Section 3.5). CPrune uses the relationship between subgraphs of the model M' and each task and creates a task/subgraph table C' . It also puts the fastest program of each task in C' (details in Section 3.4). Furthermore, the IR and tuning process maintains a candidate list of prioritized tasks R' for the next iteration (details in Section 3.3). If the measured execution time of M' (l_m) is less than l_t , it trains M' shortly and measures the short-term accuracy a_s . When a_s meets the requirement of the current iteration, it updates parameters and goes to the next iteration. α is the ratio to represent the minimum allowable accuracy after pruning, and β is the ratio to define the target execution time of the next pruning iteration.

If l_m is greater than or equal to l_t , the system selects the next task from R to create the next pruned candidate model. Once a_s is less than the target accuracy of the current iteration ($\alpha \cdot a_p$), CPrune removes the current task from R and does not consider this task as a candidate for pruning in the rest iterations.

Final step: final training and tuning (Line 17). If there are no more tasks in R that can be pruned while meeting the accuracy requirement, CPrune




Fig. 4: The graph analysis of CPrune for the ResNet-18 model

progresses to the final step. In this step, CPrune trains and tunes the final model to achieve optimal accuracy and execution time.

3.3 Task Ordering

CPrune needs to select the most promising task of the current model M for pruning. Therefore, CPrune sorts tasks according to the pruning impact (task’s execution time \times number of subgraphs associated with the task). The higher the pruning impact, the more likely the task satisfies the current iteration’s target execution time l_t .

The rightmost shaded box in Figure 3 shows the result of task ordering as an example. The pruning impact of T_1 is $0.954 \times 2 = 1.908$, T_2 is $0.473 \times 3 = 1.419$, T_3 is $1.632 \times 1 = 1.632$, and so on. After CPrune performs task ordering including the results of other tasks, the ordered list R' becomes $\{T_1, T_3, T_2, \dots, T_n, \dots\}$. In this example, the subgraphs associated to T_1 will be pruned first. If this choice does not satisfy the execution time and accuracy requirements, T_3 and then T_2 will be selected for the next pruning candidate.

3.4 Task and Subgraph Relationship

CPrune creates a table storing the relationship among tasks, subgraphs, and fast programs to determine which subgraphs are pruned and which program is analyzed when selecting a task in the next pruning iteration. CPrune analyzes a pruned candidate model by checking the layer connectivity, weight shape, Rectified Linear Unit (ReLU), and Batch Normalization (BN). Figure 4 shows CPrune’s graph analysis for part of the ResNet-18 model [3]. As ResNet-18 consists of multiple convolutions of the same shape, the compiler partitions the large computational graph of a DNN into multiple subgraphs, each of which can be associated with other subgraphs [32]. For example, subgraphs S_{11} and S_{14} are connected to the same task T_7 due to the same properties of BN and input shapes. On the other hand, subgraphs S_{10} , S_{12} , and S_{13} are connected to different tasks, respectively. In addition, each task is connected to its fastest program during the tuning.




Fig. 5: Sampled program comparison extracted from one subgraph among several subgraphs of ResNet-18. c_i is the number of input channels.

3.5 Pruning Decision

CPrune prunes filters while maintaining the structure of the fastest program created during the IR tuning process of compiler optimization. Figures 5 (b) and (c) compare the structures of the fastest and slowest programs generated from the same subgraph in Figure 5 (a). As we can see, the fastest program Figure 5 (b) allows a subgraph to perform parallel processing of computations most efficiently. In contrast, the slowest program Figure 5 (c) shows the most inefficient parallel processing of computations. If CPrune prunes four filters suitable for maintaining the structure of the program Figure 5 (c), a DNN compiler may generate an inefficient program for the corresponding subgraph of a DNN model. Instead, pruning the subgraph to follow the program structure of the fastest program would lead a DNN compiler to generate an efficient program. In addition, since the execution time of the convolution layer increases in a step pattern rather than linear with the number of filters [38], pruning an insufficient number of filters would only increase the tuning time without improving FPS.

In particular, CPrune splits a given model into multiple subgraphs and determines the number of filters to prune for each subgraph. We elaborate on calculating the number of filters to be pruned by taking one subgraph as an example, as shown in Figure 5 (a). We can see that the total number of filters is 512, and the kernel shape is 7×7 like Figure 5 (d) by checking the 2D convolution layer (Conv2d) and the ReLU activation function layer (T.relu) in Figure 5

(a). Iterators related to the number of filters are **ff** and **ax3**. With this information, CPrune checks how these iterators are further split or merged. For example, Figure 5 (b) shows that 512 of **ff** in Figure 5 (a) is converted to $4 \times 8 \times 16$ by the argument **ff.3**. 512 of **ax3** is also converted to $4 \times 8 \times 16$. Therefore, both iterators arrange filters like Figure 5 (e).

Finally, CPrune calculates the minimum number of filters to prune based on the arrangement shapes of filters obtained from the related iterators. If CPrune reduces the number in width, height, or depth in the arrangement of filters of Figure 5 (e), the subgraph can maintain the structure of Figure 5 (b) after pruning the filters. However, if CPrune prunes many filters at once, it cannot achieve optimal pruning due to the high difference from the performance goal in the current pruning iteration. Therefore, we prune filters by the step size to improve performance by pruning in the convolution layer [38]. Let $L_1 = \{a_1, a_2, \dots, a_m\}$ be the set of product combinations of the first iterator, and $L_2 = \{b_1, b_2, \dots, b_n\}$ be the set of product combinations of the second iterator. CPrune determines the least common multiple (LCM) for the two pruning numbers as the minimum number of pruning filters using L_1 and L_2 .

$$LCM(\min_{l \in L_1} \frac{\prod_{k=1}^m a_k}{l}, \min_{l \in L_2} \frac{\prod_{k=1}^n b_k}{l}) = LCM(\frac{\prod_{k=1}^m a_k}{\max_{l \in L_1} l}, \frac{\prod_{k=1}^n b_k}{\max_{l \in L_2} l})$$

For example, we calculate the minimum number of filters that can be pruned while maintaining the program structure in Figure 5 (b) and Figure 5 (c), respectively. In Figure 5 (b), the product combination of iterators **ff** and **ax3** is the same as $4 \times 8 \times 16$. CPrune can prune 8×16 , 4×16 , or 4×8 filters like Figure 5 (f) while maintaining the product combination. Thus, the minimum number of pruning filters of Figure 5 (b) is $LCM(\min\{8 \times 16, 4 \times 16, 4 \times 8\}, \min\{8 \times 16, 4 \times 16, 4 \times 8\}) = 32$. On the other hand, in Figure 5 (c), **ff** and **ax3** are converted to 4×128 and 512×1 , respectively. Therefore, the minimum number of pruning filters of Figure 5 (c) is $LCM(\min\{4, 128\}, \min\{512, 1\}) = 4$.

After determining the number of filters to be pruned, CPrune decides which filters to prune. It calculates the sum of each filter’s absolute weights (i.e., l_1 -norm) and prunes filters starting with the smallest sum [2,21].³ CPrune also prunes the input channels of the next layer by the determined number as described in Section 2.1 to maintain the connection consistency. In the case of the shortcut part in a ResNet model, we prune the same output channels using the dependency-aware mode supported by Microsoft NNI [2].

4 Experiments

4.1 Experimental Setup

In this section, we carry out various experiments to evaluate the performance of CPrune. We conduct experiments over different pre-trained DNN models, including ResNet-18 [12], MobileNetV2 [34], and MnasNet1.0 [37], with ImageNet [18]

³ Using other metrics can improve the performance as well.




Fig. 6: The FPS increase rate and short-term accuracy during the iterative CPrune process compared to using only TVM auto-tune. After completing the compiler-informed pruning process, CPrune performs final long-term training and tuning (Line 17 of Algorithm 1). Therefore, we marked the final results with dotted lines to distinguish them from other short-term measurements.

or CIFAR-10 [17] datasets on various resource-constrained mobile devices (Samsung Galaxy S8 (Kryo 280 CPU), S9 (Kryo 385 CPU and Mali-G72 GPU), S20+ (Kryo 585 CPU), or Google Pixel 3 XL (Kryo 385 CPU)). We also use multiple host PCs with NVIDIA GeForce RTX 1080 Ti or 2080 Ti. All pruned models are optimized by stochastic gradient descent (SGD) [31]. The number of training epochs considered varies according to the choice of the dataset and the specific phase of training (i.e., either short-term or final training). For the CIFAR-10 dataset, the short-term and final training epochs are 5 and 100, respectively. For the ImageNet dataset, the training epochs are one-fifth of CIFAR-10 due to the enormous data size. Due to space limitations, the details of finding reasonable α and β values and additional experiments about the impact of tuning and selective search are in the Supplementary Materials.

4.2 Overall Performance

This experiment shows CPrune’s FPS and short-term accuracy during the iterative pruning process compared to the case of using only a DNN compiler and its auto-tuning. For this experiment, we select TVM auto-tune [49] for the DNN compiler and use the ResNet-18 model on Kryo 385 CPU using the ImageNet dataset. Figure 6 shows the results. In each pruning iteration, CPrune modifies the current model by pruning selected subgraphs. For any given pruning iteration, if a pruned candidate model satisfies the condition of the given iteration (l_t and $\alpha \cdot a_p$), the candidate model is selected, and the iteration progresses. The final accuracy is 88.34%, and the FPS increase rate is 1.96 times faster than using TVM auto-tune alone. We emphasize that our CPrune’s FPS increase rate is around two times that obtained by only using the TVM auto-tune, while the




Fig. 7: Comparison of FPS when we execute CPrune model with the target compiler framework TVM and a target-agnostic deep learning framework TFLite.

final accuracy is within tolerable limits. Note that CPrune can stop its pruning around the 6th pruning iteration if the accuracy requirement is more than 89%. Then the FPS increase rate is 1.3 times faster than TVM auto-tune.

In a practical scenario, user applications can provide the accuracy and execution time requirements to CPrune.

4.3 Performance on Different Target Devices

This section evaluates CPrune’s performance on different target devices. For this experiment, we integrate CPrune with the TVM compiler to compare its performance on different target devices. We also convert the final pruned model to a widely used TFLite executable for performance comparison. We compare the FPS of different DNN models using the ImageNet dataset on mobile CPU (Kryo 385 or Kryo 585) and GPU (Mali-G72), as shown in Figure 7. CPrune, along with the target compiler framework, shows a significantly higher FPS than the cases of running only a DNN compiler (e.g., TVM) and a target agnostic DNN library (e.g., TFLite). Furthermore, regardless of the type of processor and model, the FPS increase rate when executing the CPrune model on a target processor is significantly higher than when we run it on other processors, as shown in Figure 8.

4.4 Comparison with Other Pruning Schemes

This section compares CPrune with other pruning schemes on different mobile platforms. We also integrate CPrune into TVM for a fair comparison. Table 1 shows the results. CPrune shows a higher FPS than the model-based pruning models (e.g., PQF [29], FPGM [13], AMC [14]). CPrune also shows similar or better performance than the hardware-aware pruning model (e.g., NetAdapt [44]) with TVM. It also shows that indirect metrics such as FLOPS and parameters do not fully reflect actual performance. While FLOPS is a suitable indirect measure of the extent of compression obtained during the pruning process, FPS suitably reflects the pruning gains in terms of execution times or speeds.




Fig. 8: Comparison of FPS when we execute CPrune model on a target processor and other processors.

Table 1: Mobile CPU (Kryo 385 and 585) and GPU (Mali-G72) performance test (ResNet-18, MobileNetV2, MnasNet1.0 with ImageNet dataset)

Model	Method	FPS (Increase rate)	FLOPS	Params	Top-1 Acc	Top-5 Acc
ResNet-18 (Kryo 385)	Original (TVM)	18.86	1.81B	11.7M	69.76%	89.08%
	PQF [29]+TVM	18.73 (0.99×)	166M	8.26M	66.74%	87.16%
	FPGM [13]+TVM	22.93 (1.22×)	1.10B	7.30M	68.37%	88.43%
	NetAdapt [44]+TVM	35.17 (1.86×)	1.24B	9.59M	68.45%	88.37%
	CPrune	36.92 (1.96×)	1.17B	10.3M	68.30%	88.34%
ResNet-18 (Mali-G72)	Original	15.65	1.81B	11.7M	69.76%	89.08%
	PQF+TVM	24.14 (1.54×)	166M	8.26M	66.74%	87.16%
	FPGM+TVM	26.62 (1.70×)	1.10B	7.30M	68.37%	88.43%
	CPrune	30.02 (1.92×)	1.55B	10.4M	69.83%	89.24%
MobileNetV2 (Kryo 385)	Original	28.20	301M	3.47M	71.88%	90.29%
	AMC [14]+TVM	67.62 (2.40×)	211M	2.31M	70.85%	89.91%
	CPrune	76.92 (2.73×)	255M	3.29M	70.33%	89.57%
MobileNetV2 (Mali-G72)	Original	68.68	301M	3.47M	71.88%	90.29%
	AMC+TVM	90.58 (1.32×)	211M	2.31M	70.85%	89.91%
	CPrune	101.56 (1.48×)	281M	3.31M	71.39%	90.16%
MnasNet1.0 (Kryo 585)	Original	42.92	314M	4.35M	73.46%	91.51%
	CPrune	56.85 (1.32×)	284M	3.82M	72.90%	91.16%

4.5 Effect of Pruning on Associated Subgraphs

This experiment checks the effectiveness of pruning filters of all subgraphs related to a task as in CPrune than pruning filters of only one subgraph in each iteration. For a task associated with greater than one subgraph, we have a design choice to prune filters of a single subgraph at a time (NetAdapt [44]) or prune filters of all the subgraphs associated with the task (like CPrune). Associated subgraphs pruning can shorten the CPrune process in proportional to the number of related subgraphs in the model. We observe that the associated subgraphs pruning consumes relatively less time in the Main step of CPrune than the single subgraph pruning, as shown in Figure 9 (a). In addition, the associated subgraphs pruning improves the FPS by more than 13 FPS compared to the

Table 2: Mobile CPU performance test (ResNet-18 with CIFAR-10 dataset)

Model	Method	FPS (Increase rate)	FLOPS	Params	Top-1 Acc
ResNet-18 (Kryo 280)	Original (TVM)	33.82	555M	11.2M	94.37%
	CPrune	109.45 (3.24 \times)	161M	2.62M	93.74%
ResNet-18 (Kryo 585)	Original	40.50	555M	11.2M	94.37%
	CPrune	93.63 (2.31 \times)	297M	3.54M	94.14%
	CPrune (w/o tuning)	57.77 (1.43 \times)	390M	5.08M	94.51%
	CPrune (single sub-graph pruning)	79.62 (1.97 \times)	294M	4.55M	94.27%




Fig. 9: The effect of pruning all the subgraphs associated with the task (ResNet-18, Kryo 585, CIFAR-10). (a) is the relative time cost comparison with pruning a single subgraph at a time

single subgraph pruning strategy without significantly reducing the accuracy, as shown in Figure 9 (b) and Table 2.

5 Conclusion

Existing methods generate a compressed model by focusing on a model itself for fast DNN model execution on resource-constrained target devices. However, we have confirmed that knowing which model is best from pruning is impossible without considering compiler optimization. Therefore, we propose CPrune, a compiler-informed model pruning for efficient target-aware DNN execution. CPrune ensures the actual performance with pruning on the target device by using the relationship between the task and subgraphs during the compiler optimization. CPrune generates a target-aware DNN execution model by pruning a model based on on-device performance and compiler optimization. We verified that the pruned model generated by CPrune improves a DNN model’s execution speed significantly while meeting the accuracy requirement.

Acknowledgments: This work was supported by the Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No. 2018-0-00769, Neuromorphic Computing Software Platform for Artificial Intelligence Systems).

References

1. Apache tvm. <https://github.com/apache/tvm>, accessed: 2022-07-03 3
2. Microsoft nni. <https://github.com/microsoft/nni>, accessed: 2022-07-03 3, 10
3. Torchvision models. <https://pytorch.org/vision/stable/models.html>, accessed: 2022-07-03 8
4. Chen, T., Moreau, T., Jiang, Z., Zheng, L., Yan, E., Shen, H., Cowan, M., Wang, L., Hu, Y., Ceze, L., et al.: {TVM}: An automated end-to-end optimizing compiler for deep learning. In: 13th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 18). pp. 578–594 (2018) 4
5. Chen, T., Zheng, L., Yan, E., Jiang, Z., Moreau, T., Ceze, L., Guestrin, C., Krishnamurthy, A.: Learning to optimize tensor programs. arXiv preprint arXiv:1805.08166 (2018) 4
6. Chen, Y., Meng, G., Zhang, Q., Xiang, S., Huang, C., Mu, L., Wang, X.: Renas: Reinforced evolutionary neural architecture search. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4787–4796 (2019) 4
7. Elsken, T., Metzen, J.H., Hutter, F.: Neural architecture search: A survey. The Journal of Machine Learning Research 20(1), 1997–2017 (2019) 4
8. Fang, B., Zeng, X., Zhang, M.: Nstdnn: Resource-aware multi-tenant on-device deep learning for continuous mobile vision. In: Proceedings of the 24th Annual International Conference on Mobile Computing and Networking. pp. 115–127 (2018) 2
9. Gale, T., Zaharia, M., Young, C., Elsen, E.: Sparse gpu kernels for deep learning. IEEE Press (2020) 4
10. Gong, Z., Ji, H., Fletcher, C.W., Hughes, C.J., Bagsorkhi, S., Torrellas, J.: Save: Sparsity-aware vector engine for accelerating dnn training and inference on cpus. In: 2020 53rd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO). pp. 796–810 (2020) 4
11. Guo, Y., Yao, A., Chen, Y.: Dynamic network surgery for efficient dnns. In: Proceedings of the 30th International Conference on Neural Information Processing Systems. pp. 1387–1395 (2016) 4
12. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016) 10
13. He, Y., Liu, P., Wang, Z., Hu, Z., Yang, Y.: Filter pruning via geometric median for deep convolutional neural networks acceleration. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4340–4349 (2019) 12, 13
14. He, Y., Lin, J., Liu, Z., Wang, H., Li, L.J., Han, S.: Amc: Automl for model compression and acceleration on mobile devices. In: Proceedings of the European conference on computer vision (ECCV). pp. 784–800 (2018) 3, 12, 13
15. He, Y., Zhang, X., Sun, J.: Channel pruning for accelerating very deep neural networks. In: Proceedings of the IEEE international conference on computer vision. pp. 1389–1397 (2017) 3
16. Kim, T., Nguyen, P., Pham, N., Bui, N., Truong, H., Ha, S., Vu, T.: Epileptic seizure detection and experimental treatment: a review. Frontiers in Neurology 11, 701 (2020) 2
17. Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images (2009) 5, 11

18. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems* **25**, 1097–1105 (2012) [10](#)
19. Kyriakides, G., Margaritis, K.: An introduction to neural architecture search for convolutional networks. *arXiv preprint arXiv:2005.11074* (2020) [2](#)
20. Lattner, C., Amini, M., Bondhugula, U., Cohen, A., Davis, A., Pienaar, J., Riddle, R., Shpeisman, T., Vasilache, N., Zinenko, O.: Mlir: Scaling compiler infrastructure for domain specific computation. In: *2021 IEEE/ACM International Symposium on Code Generation and Optimization (CGO)*. pp. 2–14. IEEE (2021) [4](#)
21. Li, H., Kadav, A., Durdanovic, I., Samet, H., Graf, H.P.: Pruning filters for efficient convnets. *arXiv preprint arXiv:1608.08710* (2016) [3](#), [10](#)
22. Li, Z., Yuan, G., Niu, W., Zhao, P., Li, Y., Cai, Y., Shen, X., Zhan, Z., Kong, Z., Jin, Q., et al.: Npas: A compiler-aware framework of unified network pruning and architecture search for beyond real-time mobile acceleration. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 14255–14266 (2021) [4](#)
23. Liang, T., Glossner, J., Wang, L., Shi, S., Zhang, X.: Pruning and quantization for deep neural network acceleration: A survey. *Neurocomputing* **461**, 370–403 (2021) [2](#)
24. Liu, H., He, Y., Yu, F.R., James, J.: Flexi-compression: A flexible model compression method for autonomous driving. In: *Proceedings of the 11th ACM Symposium on Design and Analysis of Intelligent Vehicular Networks and Applications*. pp. 19–26 (2021) [2](#)
25. Liu, H., Simonyan, K., Yang, Y.: Darts: Differentiable architecture search. *arXiv preprint arXiv:1806.09055* (2018) [4](#)
26. Liu, J., Tripathi, S., Kurup, U., Shah, M.: Pruning algorithms to accelerate convolutional neural networks for edge applications: A survey. *arXiv preprint arXiv:2005.04275* (2020) [2](#), [3](#)
27. Lu, L., Yu, J., Chen, Y., Liu, H., Zhu, Y., Kong, L., Li, M.: Lip reading-based user authentication through acoustic sensing on smartphones. *IEEE/ACM Transactions on Networking* **27**(1), 447–460 (2019) [2](#)
28. Ma, X., Guo, F.M., Niu, W., Lin, X., Tang, J., Ma, K., Ren, B., Wang, Y.: Pconv: The missing but desirable sparsity in dnn weight pruning for real-time execution on mobile devices. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 34, pp. 5117–5124 (2020) [4](#)
29. Martinez, J., Shewakramani, J., Liu, T.W., Bârsan, I.A., Zeng, W., Urtasun, R.: Permute, quantize, and fine-tune: Efficient compression of neural networks. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 15699–15708 (2021) [12](#), [13](#)
30. Periaux, J., Gonzalez, F., Lee, D.S.C.: *Evolutionary methods*. In: *Evolutionary Optimization and Game Strategies for Advanced Multi-Disciplinary Design*, pp. 9–20. Springer (2015) [4](#)
31. Robbins, H., Monro, S.: A stochastic approximation method. *The annals of mathematical statistics* pp. 400–407 (1951) [11](#)
32. Roesch, J., Lyubomirsky, S., Weber, L., Pollock, J., Kirisame, M., Chen, T., Tatlock, Z.: Relay: A new ir for machine learning frameworks. In: *Proceedings of the 2nd ACM SIGPLAN International Workshop on Machine Learning and Programming Languages*. pp. 58–68 (2018) [4](#), [8](#)
33. Rotem, N., Fix, J., Abdulrasool, S., Catron, G., Deng, S., Dzhabarov, R., Gibson, N., Hegeman, J., Lele, M., Levenstein, R., et al.: Glow: Graph lowering compiler techniques for neural networks. *arXiv preprint arXiv:1805.00907* (2018) [4](#)

34. Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., Chen, L.C.: Mobilenetv2: Inverted residuals and linear bottlenecks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4510–4520 (2018) [10](#)
35. Shahriari, B., Swersky, K., Wang, Z., Adams, R.P., De Freitas, N.: Taking the human out of the loop: A review of bayesian optimization. Proceedings of the IEEE **104**(1), 148–175 (2015) [4](#)
36. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014) [5](#)
37. Tan, M., Chen, B., Pang, R., Vasudevan, V., Sandler, M., Howard, A., Le, Q.V.: Mnasnet: Platform-aware neural architecture search for mobile. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2820–2828 (2019) [4](#), [10](#)
38. Tang, X., Han, S., Zhang, L.L., Cao, T., Liu, Y.: To bridge neural network design and real-world performance: A behaviour study for neural networks. Proceedings of Machine Learning and Systems **3**, 21–37 (2021) [9](#), [10](#)
39. Wang, Z., Li, C., Wang, X.: Convolutional neural network pruning with structural redundancy reduction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14913–14922 (2021) [3](#)
40. Wei, L., Luo, W., Weng, J., Zhong, Y., Zhang, X., Yan, Z.: Machine learning-based malicious application detection of android. IEEE Access **5**, 25591–25601 (2017) [2](#)
41. White, C., Neiswanger, W., Savani, Y.: Bananas: Bayesian optimization with neural architectures for neural architecture search. arXiv preprint arXiv:1910.11858 **1**(2) (2019) [4](#)
42. Wu, B., Dai, X., Zhang, P., Wang, Y., Sun, F., Wu, Y., Tian, Y., Vajda, P., Jia, Y., Keutzer, K.: Fbnet: Hardware-aware efficient convnet design via differentiable neural architecture search. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10734–10742 (2019) [4](#)
43. Yang, T.J., Chen, Y.H., Sze, V.: Designing energy-efficient convolutional neural networks using energy-aware pruning. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 5687–5695 (2017) [4](#)
44. Yang, T.J., Howard, A., Chen, B., Zhang, X., Go, A., Sandler, M., Sze, V., Adam, H.: Netadapt: Platform-aware neural network adaptation for mobile applications. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 285–300 (2018) [3](#), [4](#), [12](#), [13](#)
45. Yang, T.J., Liao, Y.L., Sze, V.: Netadaptv2: Efficient neural architecture search with fast super-network training and architecture optimization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2402–2411 (2021) [4](#)
46. Yu, J., Lukefahr, A., Palframan, D., Dasika, G., Das, R., Mahlke, S.: Scalpel: Customizing dnn pruning to the underlying hardware parallelism. ACM SIGARCH Computer Architecture News **45**(2), 548–560 (2017) [4](#)
47. Zhang, T., Ye, S., Zhang, K., Tang, J., Wen, W., Fardad, M., Wang, Y.: A systematic dnn weight pruning framework using alternating direction method of multipliers. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 184–199 (2018) [4](#)
48. Zhao, C., Ni, B., Zhang, J., Zhao, Q., Zhang, W., Tian, Q.: Variational convolutional neural network pruning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2780–2789 (2019) [3](#)
49. Zheng, L., Jia, C., Sun, M., Wu, Z., Yu, C.H., Haj-Ali, A., Wang, Y., Yang, J., Zhuo, D., Sen, K., et al.: Ansor: Generating high-performance tensor programs for

- deep learning. In: 14th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 20). pp. 863–879 (2020) [4](#), [11](#)
50. Zhou, B., Lohokare, J., Gao, R., Ye, F.: Echoprint: Two-factor authentication using acoustics and vision on smartphones. In: Proceedings of the 24th Annual International Conference on Mobile Computing and Networking. pp. 321–336 (2018) [2](#)
51. Zoph, B., Le, Q.V.: Neural architecture search with reinforcement learning. arXiv preprint [arXiv:1611.01578](#) (2016) [4](#)