

RAPID NEURAL PRUNING FOR NOVEL DATASETS WITH SET-BASED TASK-ADAPTIVE META-PRUNING

Anonymous authors

Paper under double-blind review

ABSTRACT

As deep neural networks are growing in size and being increasingly deployed to more resource-limited devices, there has been a recent surge of interest in network pruning methods, which aim to remove less important weights or activations of a given network. A common limitation of most existing pruning techniques, is that they require pre-training of the networks at least once before pruning, and thus we can benefit from reduction in memory and computation only at the inference time. However, reducing the *training* cost of neural networks with rapid structural pruning may be beneficial either to minimize monetary cost with cloud computing or to enable on-device learning on a resource-limited device. Recently introduced random-weight pruning approaches can eliminate the needs of pretraining, but they often obtain suboptimal performance over conventional pruning techniques and also does not allow for faster training since they perform unstructured pruning. To overcome their limitations, we propose *Set-based Task-Adaptive Meta Pruning (STAMP)*, which task-adaptively prunes a network pretrained on a large reference dataset by generating a pruning mask on it as a function of the target dataset. To ensure maximum performance improvements on the target task, we *meta-learn* the mask generator over different subsets of the reference dataset, such that it can generalize well to any unseen datasets within a few gradient steps of training. We validate STAMP against recent advanced pruning methods on benchmark datasets, on which it not only obtains significantly improved compression rates over the baselines at similar accuracy, but also orders of magnitude faster training speed.

1 INTRODUCTION

Deep learning has achieved remarkable progress over the last years on a variety of tasks, such as image classification (Krizhevsky et al., 2012; Wang et al., 2017; Rawat & Wang, 2017), object detection (Lin et al., 2017b; Liu et al., 2020), and semantic segmentation (Lin et al., 2017a; Huang et al., 2019). A key factor to the success of deep neural networks is their expressive power, which allows them to represent complex functions with high precision. Yet, such expressive power came at the cost of increased memory and computational requirement. Moreover, there is an increasing demand to deploy deep neural networks to resource-limited devices, which may not have sufficient memory and computing power to run the modern deep neural networks. Thus, many approaches have been proposed to reduce the size of the deep neural networks, such as network pruning, training the model with sparsity-inducing regularizations or prior (Han et al., 2015; Yoon & Hwang, 2017; Lee et al., 2019a), network distillation (Hinton et al., 2014; Hui et al.), and network quantization (Han et al., 2016; Jung et al., 2019). Arguably the most popular approach among them is network pruning, which aims to find the optimal subnetwork that is significantly smaller than the original network either by removing its weights and activations (*unstructured*) or filters and channels (*structured*). Structured pruning is often favored over unstructured pruning since GPUs can exploit its data locality to yield actual reduction of inference time, while unstructured pruning sometimes lead to longer inference time than the full networks (Wen et al., 2016).

Yet, most conventional pruning techniques have a common limitation, in that they require a network pretrained on the target dataset. With such two-stage schemes, training will inevitably take more time than training of a full network, and thus most works focus only on the efficiency at inference time. However, in many real-world scenarios, it may be desirable to obtain *training-time speedups* with pruning. For instance, if we have to train a large network for a large dataset on cloud, it may

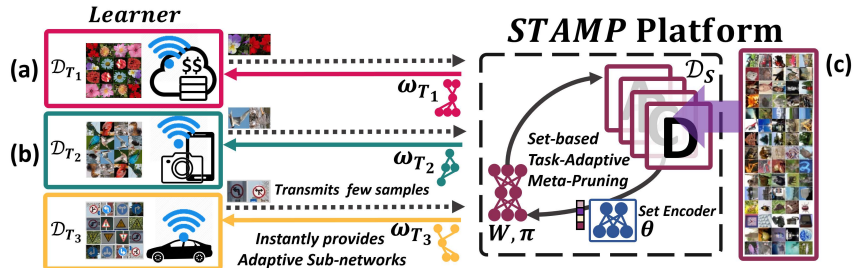


Figure 1: **Illustrations of our Set-based Task-Adaptive Meta-Pruning (STAMP):** *STAMP* meta-learns a general strategy to rapidly perform structural pruning of a reference network, for unseen tasks. If a learner gives a small fraction of information for his/her target tasks, *STAMP* almost instantly provides an optimally pruned network architecture which will train faster than the full network with minimal accuracy loss.

incur large monetary cost (Figure 1(a)). As another example, due to concerns on data privacy, we may need to train the network on resource-limited devices (Figure 1(b)). However, the device may not have enough capacity even to load the original unpruned networks on memory. Handling such diverse requirements for each end-user when pruning neural networks to their needs, is crucial for the success of an AutoML platform (Figure 1). Then how can we perform pruning without pretraining on the target tasks?

A few recently introduced methods, such as SNIP (Lee et al., 2019c) and Edge-Popup (Ramanujan et al., 2019) allow to prune randomly initialized neural networks, such that after fine-tuning, the pruned network obtains performance that is only marginally worse than that of the full network. This effectively eliminates the needs of pretraining, and SNIP further reduces pruning time by performing pruning in a single forward pass. However, they are limited in that they perform *unstructured pruning* which will not result in meaningful speedups on GPUs, either at inference or training time. Moreover, they underperform state-of-the-art structure pruning techniques with pretraining. Thus, none of the existing works can obtain structurally pruned subnetworks that provide us practical speedups both at the training and inference time, with minimal accuracy loss over the full network.

To achieve this challenging goal, we first focus on that in real-world scenarios, we may have a network pretrained on a large reference dataset (Figure 1 (c)). If we could prune such a reference pretrained network to make it obtain good performance on an *unseen* target task, it would be highly efficient since we only need to train the model once and use it for any given tasks. However, pruning a network trained on a different dataset may yield a suboptimal subnetwork for the target task. Thus, to ensure that the pruned network obtains near-optimal subnetwork for an unseen task, we propose to *meta-learn* the task-adaptive pruning mask generator as a *set function*, such that given few samples of the target dataset, it instantly generates a task-optimal subnetwork of a pretrained reference network.

We validate our *Set-based Task-Adaptive Meta Pruning (STAMP)* on multiple benchmark datasets against recently proposed structural pruning and random-weight pruning baselines. The results show that our method can rapidly prune an network to obtain a network is significantly more compact than the networks with similar accuracy using baseline pruning techniques. Further, this rapid structural pruning allows our model to significantly reduce the training cost in terms of both memory, computation, and wall-clock time, with minimal accuracy loss. Such efficiency makes *STAMP* appealing as a cheap alternative for neural architecture search in machine learning platforms (See Figure 1). The contribution of our work is threefold:

- We propose a novel **dataset-dependent structured pruning** model, which instantly generates a pruning mask for a given dataset to prune a target network.
- We suggest a **meta-learning** framework to train our set-based pruning mask generator, to obtain an approximately optimal subnetwork within few gradient steps on an unseen task.
- We validate our meta-pruning model on **three benchmark datasets** against structured and random weight pruning baselines, and show that it obtains significantly more compact subnetworks, which require only **a fraction of wall-clock time** to train the network to target accuracy, compared to baselines.

2 RELATED WORK

Neural network pruning. There has been a surge of interest on weight pruning schemes for deep neural networks to promote memory/computationally efficient models. Unstructured pruning methods prunes the weight of the network without consideration of its structure. Some of them have been shown to obtain extremely sparse networks that match the accuracy of full network, such as iterative magnitude pruning (Han et al., 2015) which repeats between training and finetuning to recover from the damage from pruning. Lottery Ticket Hypothesis (LTH) (Frankle & Carbin, 2019) discusses the existence of a subnetwork which matches the accuracy and training time of a full network, referred as the *winning ticket*, and show that they can be found with iterative magnitude pruning. Lee et al. (2019c) propose a simple pruning method which can identify a similar subnetwork without pretraining in single forward step. Though SNIP does not strictly find a *winning ticket*, it is highly efficient and opens possibility to further research on rapid pruning without pretraining. Edge-Popup (Ramanujan et al., 2019) finds optimal subsets from random weights, without any pretraining, which is also simple. However, SNIP is faster than Edge-popup in searching a pruned network.

Although unstructured pruning methods find extremely sparse subnetworks and gets simpler, due to poor data locality, it is difficult to reduce the network inference time on general-purpose hardware. Due to this limitation, recent works (Liu et al., 2017; 2019b; He et al., 2017; Guo et al., 2020; Luo et al., 2017; Zhuang et al., 2018) target to prune groups of weights (e.g., channels or neurons), to achieve actual reduction in the model size. Such structured pruning methods are useful in a resource-limited environment with compressed architectures to practically reduce the memory requirement and the running time and at inference time. Wen et al. (2016) introduce a structured sparsity regularization method to prune neurons using (2,1)-norm during training. Yoon & Hwang (2017) propose to combine group sparsity with exclusive sparsity regularization. VIBNet (Dai et al., 2018) utilizes the variational information bottleneck principle to compress neural networks. They compel the networks to minimize the neuron redundancy across adjacent layers with binary mask vectors. BBDropout (Lee et al., 2019a) learns a structured dropout function sampled from the Bernoulli distribution where the probability is given from the beta distribution with learnable parameters. Further, they introduce a data dependent BBDropout which generates a pruning mask as a function of the given data instance.

Meta learning. Meta-learning, which learns over a distribution of task, have shown its efficiency in handling unseen tasks for various tasks, such as few-shot learning and sample-efficient reinforcement learning. The most popular meta-learning methods are gradient-based approaches (Finn et al., 2017; Nichol & Schulman, 2018), which aim to find an initialization that can rapidly adapt to new tasks. BASE (Shaw et al., 2019) learns through MAML algorithm to rapidly search for optimal neural architecture, and thus significantly reduce the search cost over state-of-art neural architecture search (NAS) methods (Liu et al., 2019a; Xie et al., 2018). BASE learns a general prior through meta learning to perform fast adaptation for unseen tasks. On the other hand, our method learns the good initialization as a function of a set, such that it can rapidly *adapt* to the given target task. MetaPruning (Liu et al., 2019b) trains a hypernetwork that can generate sparse weights for any possible structures (i.e. the number of channels) of a network architecture. However, the hypernetwork does not generalize across tasks and thus the method requires to train one hypernetwork for each task.

3 RAPID NEURAL PRUNING FOR UNSEEN TASKS WITH STAMP

We introduce a novel structural pruning method for deep neural networks, *Set-based Task-Adaptive Meta-Pruning* (STAMP), which rapidly searches and prunes uninformative units/filters of the initial neural network trained on some other reference datasets. In Section 3.1, we define an optimization problem for deep neural networks with pruning masks. In Section 3.2, we describe our set-based structural pruning method which efficiently reduces the model size in a few gradient steps while avoiding accuracy degradation. Finally, in Section 3.3, we describe our full meta-learning framework to train the pruning mask generator that generalizes to unseen tasks.

3.1 PROBLEM DEFINITION

Suppose that we have a neural network $f(\mathbf{x}; \mathbf{W})$, which is a function of the dataset $\mathcal{D} = \{\mathbf{x}_i, \mathbf{y}_i\}_{i=1}^N$ parameterized by a set of model weights $\mathbf{W} = \{\mathbf{W}_l\}_{l=1}^L$, where $\mathbf{x}_i \in \mathbb{R}^{X_d}$ and l is a layer. Further




Figure 2: **Set-based Task-Adaptive Pruning:** We sample the subset \tilde{X} from \mathcal{D} and train the model while simultaneously optimizing set-based binary masks through a set encoding function and mask generative functions.

suppose that the network has maximum desired cost κ (e.g., FLOPs, Memory, the number of parameters, and training/inference time), which depends on the hardware capability and applications. By denoting the total cost of the model as \mathcal{C} , we formulate the problem of searching for a network that minimizes the task loss while satisfying the total cost \mathcal{C} as an optimization problem, as follows:

$$\underset{\mathbf{W}}{\text{minimize}} \frac{1}{N} \sum_{i=1}^N \mathcal{L}(f(\mathbf{x}_i; \mathbf{W}), \mathbf{y}_i) + \mathcal{R}(\mathbf{W}), \quad \text{s.t. } |\mathcal{C}| \leq \kappa \quad (1)$$

where \mathcal{R} is an arbitrary regularization term. To obtain an optimal model with the desired cost, we basically follow popular pruning strategy that adopts sparsity-inducing masking parameters for deep neural networks. We reformulate the problem as obtaining compressed weights ω_l with the corresponding binary masks \mathbf{m}_l at layer l , $\omega_l = \mathbf{m}_l \odot \mathbf{W}_l$, where $\mathbf{m}_l \in \{0, 1\}^{l_l \times C_l}$. This will result in *unstructured pruning*, which will prune individual weight elements. However, we may allow the model to compress its size by *structured pruning*, to yield actual wall-clock time speedup in training/inference time. We focus on generating structural pruning masks where the compressed weights will be expressed as $\omega_l = \mathbf{m}_l \otimes \mathbf{W}_l$, where $\mathbf{m}_l \in \{0, 1\}^{C_l}$. Then, the objective function is defined to minimize a following loss function: $\mathcal{L}(h(\mathcal{D}; \omega)) + \mathcal{R}(\omega)$ where $\omega = \{\omega_l\}_{l=1}^L$.

3.2 RAPID NEURAL PRUNING WITH SET-ENCODED REPRESENTATION

To obtain an optimal pruned structure for the target task, we need to exploit the knowledge of the given task. Conventional pruning schemes search for the desired subnetworks through full mini-batch training where all of the instances are trained through numerous iterations, incurring excessive training cost \mathcal{C} as the data size gets bigger. To bypass this time-consuming search, and rapidly obtain the task-adaptive pruning masks, we adopt two learnable functions: a **set encoding function** $e(\mathcal{D}; \theta)$ generates a set encoded output and a **mask generative function** $g(\cdot; \pi)$ obtains a binary mask vector \mathbf{m} , parameterized by θ and π , respectively. That is, at each layer l , through two different functions, the model generates the task-adaptive mask vector \mathbf{m}_l given the dataset-level encoded representation from a set encoding function. To reduce an burden for encoding the entire dataset, we use a sampled subset $\tilde{\mathbf{X}} \in \mathbb{R}^{B \times X_d} \sim \{\mathbf{x}_i\}_{i=1}^N$ from \mathcal{D} , where $B < N$ is the sampled batch and X_d is the input dimensionality. To this end, we formulate the objective of our set-based task-adaptive pruning as follows:

$$\underset{\mathbf{W}, \theta, \pi}{\text{minimize}} \frac{1}{N} \sum_{i=1}^N \mathcal{L}(f(\mathbf{x}_i; \omega), \mathbf{y}_i) + \mathcal{R}(\omega), \quad |\mathcal{C}| \leq \kappa, \quad (2)$$

$$\omega_l = g_l(\mathbf{o}_l; \pi_l) \otimes \mathbf{W}_l, \quad \mathbf{o}_l = h_l(\mathbf{o}_{l-1}; \mathbf{W}_l), \quad \text{s.t. } \mathbf{o}_0 = e(\tilde{\mathbf{X}}; \theta) \in \mathbb{R}^{r \times X_d} \quad (3)$$

where r is a batch dimension of the set representation. Throughout the paper, we use $r = 1$. The illustration of the set-based task-adaptive pruning model is described in Figure 2.

3.3 META-UPDATE FOR UNSEEN TARGET TASKS

Now we describe how we learn the optimal parameters θ and π for set based pruning. The simplest approach we can use is performing gradient descent through back propagation, as our model is end-to-end differentiable. However, this only results in optimized parameters for a specific task \mathcal{D}_S , which will not allow us to obtain an optimized parameters for an unseen task. As stated earlier, we want to apply our method across different task domains, such that we learn the pruning mask generator and the set encoder on \mathcal{D}_S and apply them on \mathcal{D}_T . To this end, we apply gradient-based meta learning method which obtains initialization parameters that rapidly adapts to given tasks.

Algorithm 1 Set-based Task-Adaptive Meta-Pruning (STAMP)

input Source Dataset \mathcal{D}_S , Target Dataset \mathcal{D}_T , Learnable parameters $\phi_0 = \{\mathbf{W}_0, \boldsymbol{\pi}_0, \boldsymbol{\theta}_0\}$

- 1: **function** STAMP
- 2: **for** $i = 1, 2, \dots$ **do**
- 3: **for** $n = 1, \dots, N$ **in parallel do**
- 4: Sample task $\mathcal{D}_S^{(n)} \subset \mathcal{D}_S$, and $\phi_i^{(n)} = \phi_{i-1}$
- 5: Sample batch of tasks $\tilde{\mathbf{X}}^{(n)} = \{\mathbf{x}_j^{(n)}\}_{j=1}^B, \{\mathbf{x}_j^{(n)}, \mathbf{y}_j^{(n)}\}_{j=1}^B \sim \mathcal{D}_S^{(n)}$
- 6: Compute $\mathcal{L}(\mathcal{D}_S^{(n)}; \phi_i^{(n)})$ with Eq. 2
- 7: **end for**
- 8: Update $\phi_i \leftarrow \phi_{i-1} - \eta \sum_{n=1}^N \nabla_{\phi_i^{(n)}} \mathcal{L}(\mathcal{D}_S^{(n)}; \phi_i^{(n)})$
- 9: **end for**
- 10: **end function**
- 11: Meta train ϕ with function STAMP
- 12: Prune K step to optimize $\mathcal{L}(\mathcal{D}_T; \phi)$ with Eq. 2
- 13: Finetune the pruned architecture to minimize $\mathcal{L}(\mathcal{D}_T; \omega)$

Basically, we train the parameters on multiple tasks sampled from \mathcal{D}_S by computing inner gradient in the inner loop and combining them to update the outer loop parameters. Then, the objective of the meta-train step of *STAMP* is learning good initialization of ϕ in the outer loop. We sample N tasks from \mathcal{D}_S , where each task is $\mathcal{D}_S^{(n)} \subset \mathcal{D}_S$. From each N tasks, a batch $\{\mathbf{x}_j^{(n)}, \mathbf{y}_j^{(n)}\}_{j=1}^B$ is sampled and it is divided into mini-batches to update inner gradient in respect to the loss function \mathcal{L} and the regularization terms described in Section 3.2. We note that the whole batch excluding labels $\tilde{\mathbf{X}}^{(n)} = \{\mathbf{x}_j^{(n)}\}_{j=1}^B$ is used for encoding set representation. For updating outer loop parameters $\phi_i = \{\mathbf{W}_i, \boldsymbol{\pi}_i, \boldsymbol{\theta}_i\}$ at i^{th} epoch, we only use the gradients of the last mini-batch, similarly to first-order MAML (Finn et al., 2017) to accelerate learning as below:

$$\phi_i \leftarrow \phi_{i-1} - \eta \sum_{n=1}^N \nabla_{\phi_i^{(n)}} \mathcal{L}(\mathcal{D}_S^{(n)}; \phi_i^{(n)}) \quad (4)$$

After meta learning a set of the parameters, we can adapt it to various unseen tasks \mathcal{D}_T by performing few steps of gradient updates, with the maximum steps of 1 epoch. Through the meta-learn procedure, we can speed-up the training time on the target task by starting with the pruned network architecture in the early stage. We describe our whole process in Algorithm 1.

While we can plug in various set encoding methods (Edwards & Storkey, 2016; Zaheer et al., 2017) or pruning methods to the proposed framework, *STAMP* adopts a transformer module (Lee et al., 2019b) for set encoding function $e(\cdot)$ and proposed a set-based pruning mask generator $g(\cdot)$ based on the Beta-Bernoulli dropout (Lee et al., 2019a).

4 EXPERIMENTS

We demonstrate the effectiveness of *STAMP* with two widely used network architectures, namely VGGNet-19 (Zhuang et al., 2018) and ResNet-18 (He et al., 2016), on several benchmark datasets (CIFAR-10, SVHN and downsampled Aircraft (Maji et al., 2013)). We implement the code for all the experiments in Pytorch framework, and use Titan XP GPU for measuring the wall-clock time.

Baselines. We validate our *STAMP* against recent structured pruning methods as well as unstructured random weight pruning methods. We also report the results on the variant of *STAMP* that only searches for the structure and randomly reinitializes the weights (*STAMP-Structure*). Baselines we use for comparative study are as follows: **1) MetaPruning** (Liu et al., 2019b): Structured pruning method which learns hypernetworks to generate pruned weights at each layer, and searches for the optimal pruned structure using an evolutionary algorithm. **2) BBDropout** (Lee et al., 2019a): Beta-Bernoulli Dropout which performs structured pruning of the channels by sampling sparse masks on them. **3) Random Pruning**: Randomly pruning of channels. We sample the random structure (i.e. the number of channels for each layer) under the given FLOP constraints in the same manner as in MetaPruning. **4) Edge-Popup** (Ramanujan et al., 2019): Unstructured pruning method that searches for the best performing sub-network of a network with random weights. **5) SNIP** (Lee et al.,

Table 1: Experiment results of CIFAR-10, SVHN, and Aircraft as a target dataset on VGGNet. **Training Time** consists of time to search for pruned network, and finetuning. **Expense** is computed by multiplying the training time by 146 ¢, which is cost of using GPU (Tesla P100) on Google Cloud. The methods are sub-divided into the full network without pruning, unstructured pruning methods, structured pruning methods and STAMP. **P** denotes the ratio of non-zero parameters. We report the mean \pm std values over three runs of experiments.

Target Task	Methods	Pretrained	Accuracy (%)	P (%)	FLOPs	Training Time	Expense
CIFAR-10	Full Network	CIFAR-100	93.72 \pm 0.07	100	x1.00	0.78 h	113 ¢
	SNIP	None	92.85 \pm 0.24	4.17	x1.00	0.83 h	121 ¢
	SNIP (P)	CIFAR-100	92.98 \pm 0.22	4.17	x1.00	0.83 h	121 ¢
	Random Pruning	None	92.01 \pm 0.29	32.20	x3.33	0.43 h	62 ¢
	MetaPruning	CIFAR-100	92.12 \pm 0.47	21.84	x3.58	4.99 h	728 ¢
	BBDropout	CIFAR-100	92.97 \pm 0.10	3.99	x3.42	2.07 h	302 ¢
	STAMP-Structure	CIFAR-100	92.69 \pm 0.13	4.43	x3.48	0.44 h	64 ¢
	STAMP	CIFAR-100	93.49 \pm 0.04	4.16	x3.56	0.44 h	64 ¢
SVHN	Full Network	CIFAR-100	95.99 \pm 0.07	100	x1.00	1.21 h	176 ¢
	SNIP	None	95.52 \pm 0.10	3.08	x1.00	1.22 h	178 ¢
	SNIP (P)	CIFAR-100	95.56 \pm 0.09	3.08	x1.00	1.22 h	178 ¢
	Random Pruning	None	95.56 \pm 0.12	28.95	x3.40	0.62 h	90 ¢
	MetaPruning	CIFAR-100	95.50 \pm 0.07	22.04	x3.64	2.08 h	303 ¢
	BBDropout	CIFAR-100	95.98 \pm 0.19	2.15	x9.67	3.05 h	445 ¢
	STAMP-Structure	CIFAR-100	95.39 \pm 0.15	3.08	x4.60	0.58 h	84 ¢
	STAMP	CIFAR-100	95.82 \pm 0.16	2.87	x5.10	0.58 h	84 ¢
Aircraft	Full Network	CIFAR-100	47.15 \pm 0.36	100	x1.00	1.90 m	4 ¢
	SNIP	None	57.58 \pm 0.00	4.07	x1.00	2.85 m	6 ¢
	SNIP (P)	CIFAR-100	57.30 \pm 0.00	4.07	x1.00	2.85 m	6 ¢
	Random Pruning	None	47.53 \pm 0.02	27.55	x3.48	0.83 m	2 ¢
	MetaPruning	CIFAR-100	50.36 \pm 0.01	24.10	x3.42	1.09 h	159 ¢
	BBDropout	CIFAR-100	42.43 \pm 0.04	13.65	x3.72	5.06 m	12 ¢
	STAMP-Structure	CIFAR-100	52.30 \pm 0.01	3.63	x4.22	1.15 m	2 ¢
	STAMP	CIFAR-100	60.36 \pm 0.01	3.39	x4.48	1.15 m	2 ¢

2019c): One-shot unstructured pruning on random weights. We also report the results of a variant of SNIP starting from pretrained weights (SNIP (P)). For finetuning, we follow the standard setting from Zhuang et al. (2018) and perform mini-batch SGD for 200 epochs where the batch size is 128.

Networks and datasets As for the base networks, we use a modified version of **VGGNet-19** with 16 convolution layers and a single fully connected layer, and **ResNet-18** with an additional 1×1 convolution layer on the shortcut operation to resolve the dimensionality difference between pruned units/filters. We use VGG-19 and ResNet-18 pretrained on CIFAR-100 (source dataset) for the baselines, and use CIFAR-10, SVHN, and Aircraft (Maji et al., 2013) as the target tasks for evaluation of the pruning performance. CIFAR-10 and SVHN have 10 classes and Aircraft has 30 classes respectively. All images are downsampled to 32x32 pixels.

Meta-training We meta-train our pruning mask generator on **CIFAR-100 dataset**. During meta-training time, we divide CIFAR-100 into 10 tasks (subsets), each of which contains 10 disjoint classes, and sampled 64 instances per class. We used total of 640 instances as the input to the set function to generate a set representation for each task. We also used the sampled instances for model training, by dividing it into 5 batches (128 instances for each). We used first-order MAML with Adam optimizer for both inner and outer parameter updates. For more details on training of the baseline methods and meta-training for STAMP, such as learning rate scheduling, please see Section C of the Appendix.


4.1 EXPERIMENTAL RESULTS

We report the results of pruning VGGNet-19 on CIFAR-10, SVHN, and Aircraft dataset in Table 1, and ResNet-18 on CIFAR-10 in Table 2. We compare the accuracy as well as wall-clock training and inference time for all models at similar compression rate (Parameter Used or FLOPs).

Accuracy over memory efficiency and FLOPs. We first compare the accuracy over the parameter usage and theoretical computation cost, FLOPs. In Table 1 and Table 2, SNIP with either random networks (SNIP) or the pretrained reference network (SNIP(P)) significantly reduce the number of activated parameters with a marginal drop of the accuracy on the target dataset (or even improvement

Table 2: Experiment results of CIFAR-10 on ResNet18. Other details are same with the Table 1. More results are in the Table 3 of the Appendix.

Methods	Accuracy (%)	P (%)	FLOPs	Training Time	Inference time	Expense
Full Network	94.37 \pm 0.12	100	x1.00	1.08 h	1.02 sec	157 ¢
Edge-Popup	89.50 \pm 3.46	10.00	x1.00	1.38 h	2.50 sec	201 ¢
SNIP (P)	93.17 \pm 0.00	10.04	x1.00	1.71 h	1.90 sec	249 ¢
SNIP	93.11 \pm 0.00	10.04	x1.00	1.71 h	1.90 sec	249 ¢
Random Pruned	91.95 \pm 0.65	69.77	x3.65	0.58 h	0.58 sec	84 ¢
MetaPruning	91.01 \pm 0.91	66.02	x4.09	3.80 h	0.58 sec	554 ¢
BBDropout	93.47 \pm 0.14	5.94	x4.11	2.17 h	0.54 sec	316 ¢
STAMP-Structure	93.63 \pm 0.08	9.07	x4.08	0.57 h	0.54 sec	83 ¢
STAMP	93.61 \pm 0.27	9.22	x4.29	0.57 h	0.54 sec	83 ¢



(a) Acc. over # of params. (CIFAR-10) (b) Acc. over training time (CIFAR-10, SVHN, Aircraft)

Figure 3: (a): Accuracy over the ratio of used parameters for CIFAR-10 on VGGNet. Full denotes the accuracy of the VGGNet before pruning. (b): Accuracy over training time for CIFAR-10, SVHN and Aircraft.

on Aircraft). However, as the methods perform unstructured pruning, they can not reduce FLOPs which remains equal to the original full networks. On the other hand, structural pruning approaches show actual FLOPs reduction by pruning a group of weights (e.g. units/filters). Interestingly, MetaPruning, which applies a learned hypernetwork on a reference architecture and dataset to prune for the target dataset, obtains suboptimal architectures which sometimes even underperforms randomly pruned networks. This shows that the learned hypernetwork does not generalize across task domains, which is expected since it is not trained with diverse tasks. BBDropout generally outperforms other baselines (except on Aircraft) with high model compression rate, but it requires large amount of training time to train the pruning mask generator, and thus yields slower training time than the full network training. On the other hand, STAMP either outperforms or achieve comparable performance to all baselines, in terms of both accuracy and compression rate. The performance gain of STAMP is especially large on the fine-grained Aircraft dataset, where using the pretrained network as is may be harmful (see the results of full network) due to mismatch between source (CIFAR-100) and target (Aircraft) task. We further report the accuracy-sparsity trade-off for SNIP, BBDropout, and STAMP (Ours) in Figure 3 (a). Our method achieves better accuracy over similar compression rates, and shows marginal performance degeneration even with 1% of the parameters remaining. Such good performance on unseen dataset is made possible by meta-learning the pruning mask generator.

Accuracy over wall-clock time for training/inference. As described earlier, our main focus in this work is to significantly reduce the training time by obtaining a near-optimal compact deep networks for unseen targets on the fly, which is not possible with any of the existing approaches. As shown in Table 1 and Table 2, unstructured random weights pruning methods (SNIP and Edge-Popup) do not results in any speedups in training time, and increases the cost of training over the full networks. These results are consistent with the findings in Frankle & Carbin (2019), which showed that most of the subnetworks require larger number of training iterations over the full network.

While structured pruning methods yield speedups in inference time over the full networks, MetaPruning and BBDropout need $\times 2.01 - \times 36.3$ and $\times 1.72 - \times 5.96$ more training time than full networks to search pruned architectures, respectively. On the contrary, STAMP instantly obtains a good subnetwork (single or less than 10 iterations according to the pruned ratio), which trains *faster* than the full network. STAMP is remarkably efficient over other structural pruning baselines, achieving $\times 3.5 - \times 54.5$ and $\times 2.81 - \times 5.76$ speedups over MetaPruning and BBDropout, respectively, with better or comparable performance. We further report the accuracy over training time for SNIP, BBDropout, and STAMP (Ours) in Figure 3 (b). Since our philosophy is train-once, and use-everywhere, once the mask generator is meta-learned on a reference dataset, it can be applied to any number of tasks without additional cost. Thus we exclude the meta training time of STAMP (15h on VGGNet and 30h on ResNet) and MetaPruning (1.2h) per task in Table 1 and Table 2.




Figure 4: **Left:** Training time over the number of training instances. **Middle:** Accuracy over the number of training instances, **Right:** Accuracy over the number of instances. All experimental results are obtained on CIFAR-10 with VGG-19. We prune 96% ~ 97% of the parameters for STAMP and SNIP.

Data size of the target tasks. We further examine the accuracy and time-efficiency of subnetworks obtained using different pruning methods on various problem size. We previously observed that STAMP can yield larger saving in the training and inference time as the network size gets larger (ReNet-18, Table 2). Another factor that defines the problem size is the number of instances in the unseen target dataset. We used subsets of CIFAR-10 to explore the effect of the task size to training time and accuracy in Figure 4. The full dataset consists of 50K images, which corresponds to the results reported in Table 1. We observe that, as the number of instances used for training increases, STAMP obtains even larger saving in the training time, while BBDropout incurs increasingly larger time to train. Further, as the number of instances used for training becomes smaller, STAMP obtains larger gains in accuracy, even outperforming the full network, since the network will become relatively overparameterized as the number of training data becomes smaller. As another comparison with structural pruning method with learned masks, when using only 1K data instances for training, BBDropout finds the subnetwork attaining 43% of parameters of the full network with $\times 1.54$ FLOP speedup, while STAMP prunes out 95.81% of the parameters, resulting in $\times 4.17$ speedup in FLOPs. This is because BBDropout learns the pruning mask on the given target task, and thus overfits when the number of training instances is small. STAMP, on the other hand, does not overfit since it mostly relies on the meta-knowledge and takes only few gradient steps for the given task.

Pruned network structures. We further show the effect of task-adaptive pruning. Since STAMP will find different compressed subnetwork for different tasks, the pruning ratio and the pruned channels at each layer will be different for each dataset. We visualize the remaining channels of each convolution layer of VGGNet on CIFAR-10, SVHN and Aircraft in Figure 5. Note that unlike existing channel pruning methods, such as MetaPruning, we do not enforce any constraints (hyper-parameters) on how much to prune, or what layer to prune since they are automatically determined by STAMP.

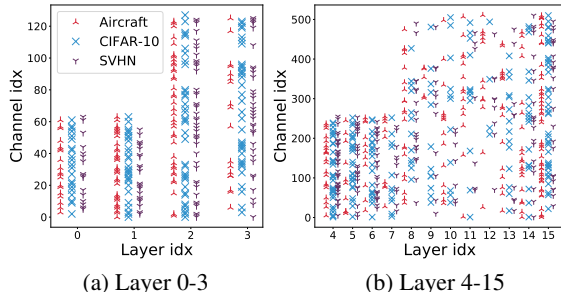


Figure 5: We denote the indices of the remaining channels at each convolution layer of VGGNet after pruning on different datasets, CIFAR-10, SVHN, and Aircraft. All of them start from same meta-learned parameters ϕ .

5 CONCLUSION

We proposed a novel set-based task-adaptive structural pruning method which instantly generates a near-optimal compact network for the given task, by performing rapid structural pruning of a global reference network trained on a large dataset. This is done by meta-learning a pruning mask generator over multiple subsets of the reference dataset as a function of a dataset, such that it can generate a pruning mask on the reference network for any unseen tasks. Our model, STAMP obtains a compact network that not only obtains good performance with large reduction in the memory and computation cost at inference time, but also enables *training* time speedup which was not possible with previous methods. Further analysis showed that STAMP obtains larger performance gains when the target dataset is small, and prunes out the channels of the same reference network differently for each dataset. We believe that both the proposal of a framework that can obtain optimal compact network for unseen tasks, and achieving training time speedup are important contributions that enhances efficiency and practicality of pruning methods.

REFERENCES

- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- Bin Dai, Chen Zhu, and David Wipf. Compressing neural networks using the variational information bottleneck. 2018.
- Harrison Edwards and Amos Storkey. Towards a neural statistician. *arXiv preprint arXiv:1606.02185*, 2016.
- Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2017.
- Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2019.
- Jinyang Guo, Wanli Ouyang, and Dong Xu. Channel pruning guided by classification loss and feature importance. *arXiv preprint arXiv:2003.06757*, 2020.
- Song Han, Jeff Pool, John Tran, and William Dally. Learning both weights and connections for efficient neural network. In *Advances in neural information processing systems*, pp. 1135–1143, 2015.
- Song Han, Huizi Mao, and William J Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2016.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Yihui He, Xiangyu Zhang, and Jian Sun. Channel pruning for accelerating very deep neural networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 1389–1397, 2017.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. In *NIPS Deep Learning Workshop*, 2014.
- Zilong Huang, Xinggang Wang, Lichao Huang, Chang Huang, Yunchao Wei, and Wenyu Liu. Ccnet: Criss-cross attention for semantic segmentation. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2019.
- Zheng Hui, Xiumei Wang, and Xinbo Gao. Fast and accurate single image super-resolution via information distillation network. In *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Sangil Jung, Changyong Son, Seohyung Lee, Jinwoo Son, Jae-Joon Han, Youngjun Kwak, Sung Ju Hwang, and Changkyu Choi. Learning to quantize deep networks by optimizing quantization intervals with task loss. In *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pp. 1097–1105, 2012.
- Juho Lee, Saehoon Kim, Jaehong Yoon, Hae Beom Lee, Eunho Yang, and Sung Ju Hwang. Adaptive network sparsification via dependent variational beta-bernoulli dropout. 2019a. URL [booktitle=https://openreview.net/forum?id=rylfl6VFDH](https://openreview.net/forum?id=rylfl6VFDH).
- Juho Lee, Yoonho Lee, Jungtaek Kim, Adam R Kosiorek, Seungjin Choi, and Yee Whye Teh. Set transformer. *Proceedings of the International Conference on Machine Learning (ICML)*, 2019b.

- Namhoon Lee, Thalaisyasingam Ajanthan, and Philip HS Torr. Snip: Single-shot network pruning based on connection sensitivity. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2019c.
- Guosheng Lin, Anton Milan, Chunhua Shen, and Ian Reid. Refinenet: Multi-path refinement networks for high-resolution semantic segmentation. In *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017a.
- Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2117–2125, 2017b.
- Hanxiao Liu, Karen Simonyan, and Yiming Yang. Darts: Differentiable architecture search. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2019a.
- Li Liu, Wanli Ouyang, Xiaogang Wang, Paul Fieguth, Jie Chen, Xinwang Liu, and Matti Pietikäinen. Deep learning for generic object detection: A survey. *International journal of computer vision*, 128(2):261–318, 2020.
- Zechun Liu, Haoyuan Mu, Xiangyu Zhang, Zichao Guo, Xin Yang, Kwang-Ting Cheng, and Jian Sun. Metapruning: Meta learning for automatic neural network channel pruning. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2019b.
- Zhuang Liu, Jianguo Li, Zhiqiang Shen, Gao Huang, Shoumeng Yan, and Changshui Zhang. Learning efficient convolutional networks through network slimming. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2017.
- Jian-Hao Luo, Jianxin Wu, and Weiyao Lin. Thinet: A filter level pruning method for deep neural network compression. In *Proceedings of the IEEE international conference on computer vision*, pp. 5058–5066, 2017.
- Subhransu Maji, Esa Rahtu, Juho Kannala, Matthew Blaschko, and Andrea Vedaldi. Fine-grained visual classification of aircraft. *arXiv preprint arXiv:1306.5151*, 2013.
- Eric Nalisnick and Padhraic Smyth. Stick-breaking variational autoencoders. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2017.
- Alex Nichol and John Schulman. Reptile: a scalable metalearning algorithm. *arXiv preprint arXiv:1803.02999*, 2:2, 2018.
- Vivek Ramanujan, Mitchell Wortsman, Aniruddha Kembhavi, Ali Farhadi, and Mohammad Rastegari. What’s hidden in a randomly weighted neural network? In *arXiv preprint arXiv:1911.13299*, 2019.
- Waseem Rawat and Zenghui Wang. Deep convolutional neural networks for image classification: A comprehensive review. *Neural computation*, 29(9):2352–2449, 2017.
- Albert Shaw, Wei Wei, Weiyang Liu, Le Song, and Bo Dai. Meta architecture search. In *Advances in Neural Information Processing Systems (NIPS)*, 2019.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems (NIPS)*, 2017.
- Fei Wang, Mengqing Jiang, Chen Qian, Shuo Yang, Cheng Li, Honggang Zhang, Xiaogang Wang, and Xiaoou Tang. Residual attention network for image classification. In *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- Wei Wen, Chunpeng Wu, Yandan Wang, Yiran Chen, and Hai Li. Learning structured sparsity in deep neural networks. In *Advances in Neural Information Processing Systems (NIPS)*, 2016.
- Sirui Xie, Hehui Zheng, Chunxiao Liu, and Liang Lin. Snas: stochastic neural architecture search. *arXiv preprint arXiv:1812.09926*, 2018.
- Jaehong Yoon and Sung Ju Hwang. Combined group and exclusive sparsity for deep neural networks. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2017.

Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Russ R Salakhutdinov, and Alexander J Smola. Deep sets. In *Advances in Neural Information Processing Systems (NIPS)*, 2017.

Zhuangwei Zhuang, Mingkui Tan, Bohan Zhuang, Jing Liu, Yong Guo, Qingyao Wu, Junzhou Huang, and Jinhui Zhu. Discrimination-aware channel pruning for deep neural networks. In *Advances in Neural Information Processing Systems*, pp. 875–886, 2018.

A APPENDIX

Organization. The supplementary file is organized as follows: We first describe each component of our Set-based Task-Adaptive Meta-Pruning (STAMP) in detail, including the set encoding function and the set-based structural pruning method (mask generative function) in Section B. Then, in Section C, we provide the detailed experimental settings and additional results on SVHN using ResNet-18 as the backbone network.

B STRUCTURAL BINARY MASK GENERATION WITH A SET-ENCODED REPRESENTATION

We now describe how we obtain the set representation with $e(\cdot; \theta)$ and learn structural pruning masks with the set-based mask generative function $g(\cdot; \pi)$ introduced in Section 3.2.

B.1 SET ENCODING FUNCTION

To obtain an optimally pruned neural architecture for the target task, we need to exploit the knowledge of the given task. Conventional pruning schemes search for the desired subnetworks through full mini-batch training, which will incur excessive training cost \mathcal{C} when the data size is large. On the other hand, we rapidly and precisely obtain the desired pruned structures given the encoded representations for each dataset. This procedure of obtaining a set representation $e(\tilde{\mathbf{X}}; \theta)$ with the set encoder e parameterized by $\theta(\cdot)$ is given as follows:

$$e(\tilde{\mathbf{X}}; \theta) = f_D(f_E(\tilde{\mathbf{X}}; \theta_E); \theta_D) \quad s.t. \quad \tilde{\mathbf{X}} \in \mathbb{R}^{B \times X_d} \sim \{\mathbf{x}_i\}_{i=1}^N, \quad \mathbf{o}_0 = e(\tilde{\mathbf{X}}; \theta) \in \mathbb{R}^{r \times X_d} \quad (5)$$

where $\tilde{\mathbf{X}} \in \mathbb{R}^{B \times X_d}$ from \mathcal{D} is a sampled task set, $B \leq N$ is the sampled batch, X_d is the input dimensionality, and r is the batch dimension of the set representation ($r \ll B$). We then define the set function as a stacked function of the encoder and the decoder where $f_E(\cdot; \theta_E) \in \mathbb{R}^{B \times X_d}$ is an arbitrary encoder function parameterized by θ_E and $f_D(\cdot; \theta_D)$ is a decoder function parameterized by θ_D . The encoder encodes the sampled task set $\tilde{\mathbf{X}}$ and the decoder regenerates the encoded vector to the dataset-level representation. Throughout the paper, we use $r = 1$ as the dimension of the set representation. We adopt a transformer module Lee et al. (2019b) for set encoding, which is a learnable pooling neural network module with an attention mechanism as shown below:

$$f_E(\cdot) = \text{AE}(\text{AE}(\cdot)), \quad f_D(\cdot) = \text{rF}(\text{AE}(\text{Pool}_r(\cdot))). \quad (6)$$

where rF is a fully connected layer, AE is an attention-based block Vaswani et al. (2017), and Pool_r is a Pooling by Multihead Attention with r seed vectors Lee et al. (2019b). AE is a permutation equivariant block that constructs upon Multi-head attention (MH), introduced in Transformer Vaswani et al. (2017). In other words, AE encodes the set information and consists $f_E(\cdot)$, while $f_D(\cdot)$ also includes AE to model correlation between r vectors after pooling. AE is defined as below:

$$\text{AE}(\cdot) = \text{Norm}(H + \text{rF}(H)) \quad (7)$$

$$\text{where } H = \text{Norm}(\cdot + \text{MH}(\cdot, \cdot, \cdot)) \quad (8)$$

where Norm is layer normalization Ba et al. (2016). The encoder encodes the given dataset using the above module, and the decoder aggregates the encoded vector. The full encoding-decoding process can be described as follows:

$$e(\tilde{\mathbf{X}}; \theta) = \text{rF}(\text{AE}(\text{Pool}_r(\text{AE}(\text{AE}(\tilde{\mathbf{X}})))))) \in \mathbb{R}^{r \times X_d} \quad (9)$$

$$\text{where } \text{Pool}_r(\cdot) = \text{AE}(R, \text{rF}(\cdot)) \in \mathbb{R}^{r \times X_d} \quad (10)$$

In here, pooling is done by applying multihead attention on a learnable vector $R \in \mathbb{R}^{r \times X_d}$. We set $r = 1$ in the experiments to obtain a single set representation vector. By stacking these attention based permutation equivariant functions, we can obtain the set representation $e(\tilde{\mathbf{X}}; \theta) \in \mathbb{R}^{1 \times X_d}$ from the sampled task $\tilde{\mathbf{X}} \in \mathbb{R}^{B \times X_d}$.

B.2 MASK GENERATIVE FUNCTION

We now describe the mask generation function $\mathbf{m}_l = g_l(\mathbf{o}_l; \boldsymbol{\pi}_l)$ at layer l , from which we obtain the pruned model parameter $\boldsymbol{\omega}_l = \mathbf{m}_l \otimes \mathbf{W}_l$. Similarly as in Lee et al. (2019a), we use the following sparsity-inducing beta-Bernoulli prior to generate a binary pruning mask at each layer, which follows Bernoulli distribution, $\text{Bernoulli}(\mathbf{m}_l; \boldsymbol{\pi}_l)$, given the probability of parameterized beta distribution as follows:

$$\mathbf{m}_l | \boldsymbol{\pi}_l \sim \prod_{c=1}^{C_l} \text{Bernoulli}(m_{c,l}; \pi_{c,l}), \quad \text{where } \boldsymbol{\pi}_l \sim \prod_{c=1}^{C_l} \text{beta}(\pi_{c,l}; \alpha_l, 1) \quad (11)$$

where C_l is the number of channels in layer l . With a learnable parameter α_l for the beta distribution, the model learns the optimal binary masks from a randomly sampled value from the beta distribution, to determine which neurons/channels should be pruned. We extend this input-independent pruning method to sample binary masks based on the set representation of the target task. This *set-dependent* pruning with STAMP is different from data-dependent BBDropout Lee et al. (2019a) in that the former generates a mask per dataset while the latter generates a mask per instance, which makes it difficult to globally eliminate a channel. Furthermore, rather than searching for the compressed structure by training with mini-batch SGD at each iteration, we utilize a set representation to rapidly obtain a near-optimal subnetwork within a few gradient steps. With the set $e(\tilde{\mathbf{X}}; \boldsymbol{\theta})$ representation obtained from the given dataset \mathbf{X} , we calculate the activation $\mathbf{o}_l = h_l(\mathbf{o}_{l-1}; \mathbf{W}_l)$ for each layer l , where h_l is the function of the layer l (i.e. convolution) and $\mathbf{o}_1 = h_1(e(\tilde{\mathbf{X}}; \boldsymbol{\theta}); \mathbf{W}_1)$. We omit the layer notation l for readability in the following equations. Then, we sample a structural pruning mask vector \mathbf{m} as follows:

$$\mathbf{m} | \boldsymbol{\pi}, \mathbf{o} \sim \prod_{c=1}^C \text{Bernoulli}(m_c; \pi_c \psi(\gamma_c o_c + \beta_c, \epsilon)), \quad \text{s.t. } [o_1; \dots; o_C] = \text{Pool}(\mathbf{o}), \quad (12)$$

where γ and β are learnable scaling and shifting factors and Pool is the average pooling for \mathbf{o} which obtains a representative value for each channel. The clamping function is defined as $\psi(\cdot, \epsilon) = \min(1 - \epsilon, \max(\epsilon, \cdot))$ with a small $\epsilon > 0$. Using a clamping function, the network will retain only the meaningful channels. We employ variational inference to approximate sparsity inducing posterior $p(\mathbf{W}, \boldsymbol{\pi}, \boldsymbol{\theta}, \boldsymbol{\beta} | \mathbf{X})$. The KL-divergence term for our set-based task-adaptive pruning is as follows:

$$D_{KL}[q(\mathbf{m}, \boldsymbol{\pi}, \boldsymbol{\theta} | \mathbf{X}) || p(\mathbf{m}, \boldsymbol{\pi}, \boldsymbol{\theta})] + D_{KL}[q(\boldsymbol{\beta}) || p(\boldsymbol{\beta})], \quad \text{s.t. } \boldsymbol{\beta} \sim \mathcal{N}(\mathbf{0}, b\mathbf{I}), \quad (13)$$

where b is a fixed value for a variance of the shifting factor $\boldsymbol{\beta}$ to prevent $\boldsymbol{\beta}$ from drifting away. The first term can be computed analytically to obtain a closed form solution Lee et al. (2019a); Nalisnick & Smyth (2017). Also, we can easily compute the second term, $\mathcal{R}(\boldsymbol{\omega})$ in the objective function of STAMP (Equation 2) by updating it with gradient-based methods.

We can further approximate the expectation for the prediction of given dataset \mathcal{D} as follows:

$$p(\mathbf{y} | \mathbf{x}, \mathbf{W}, \boldsymbol{\theta}) \approx p(\mathbf{y} | f(\mathbf{x}; \mathbb{E}_q[\mathbf{m}] \otimes \mathbf{W}, \boldsymbol{\theta})), \quad \text{s.t. } \mathbb{E}_q[\mathbf{m}] = \mathbb{E}_q[\boldsymbol{\pi}] \cdot \psi(\gamma \text{Pool}(\mathbf{o}) + \boldsymbol{\beta}, \epsilon) \quad (14)$$

C EXPERIMENTS


C.1 EXPERIMENTAL SETTINGS

We first describe how we meta-train STAMP and set the settings for the baselines, SNIP Lee et al. (2019c) and MetaPruning Liu et al. (2019b), for the experiments in the main paper (VGGNet and ResNet-18 on benchmark datasets: CIFAR-10, SVHN, and Aircraft).

For STAMP, in **function** STAMP in Algorithm 1, we update $\phi = \{\mathbf{W}, \boldsymbol{\pi}, \boldsymbol{\theta}\}$ with the learning rate 0.001, 0.01, and 0.001 with Adam optimizer, while decreasing the learning rate by 0.1 at 50% and 80% of the total epoch, following the settings of BBDropout Lee et al. (2019a). For Algorithm 1, we select $i = 2000$, $B = 640$ for VGGNet and $i = 3000$, $B = 500$ for ResNet-18. We sampled the same number of instances per class. We further set $N = 10$ and the size of the minibatch as 32. When pruning with STAMP, we use the same learning rate as the one we use in the meta training stage for VGGNet. However, for ResNet-18, we set the learning rates as $\boldsymbol{\pi} = 0.5$ to adjust the pruning rate.

Table 3: Experiment results of SVHN on ResNet-18. **Training Time** consists of time to search for the pruned network and finetuning (20 epochs). **Expense** is computed by multiplying the training time by 1.46 €, which is the cost of using GPU (Tesla P100) on Google Cloud. The methods are sub-divided into the full network without pruning, unstructured pruning methods, structured pruning methods, and STAMP (STAMP-Structure is a variation of STAMP, which re-initializes the pruned architecture). **P** is the remaining parameter ratio. We run each experiment 3 times and report the mean \pm std values.

Methods	Accuracy (%)	P (%)	FLOPs	Training Time	Inference time	Expense
Full Network	94.57 \pm 0.01	100	x1.00	0.16 h	3.30 sec	24 €
Edge-Popup	92.61 \pm 0.01	5.00	x1.00	0.20 h	6.15 sec	29 €
SNIP (P)	95.38 \pm 0.01	6.06	x1.00	0.35 h	6.64 sec	51 €
SNIP	94.88 \pm 0.01	6.06	x1.00	0.35 h	6.64 sec	51 €
Random Pruned	94.39 \pm 0.23	72.17	x2.99	0.08 h	1.66 sec	12 €
MetaPruning	94.49 \pm 0.19	70.99	x2.83	2.41 h	1.68 sec	351 €
BBDropout	94.32 \pm 0.02	4.90	x5.25	0.31 h	1.52 sec	46 €
STAMP-Structure	95.17 \pm 0.01	4.81	x5.47	0.11 h	1.51 sec	16 €
STAMP	95.41 \pm 0.01	4.81	x5.47	0.11 h	1.51 sec	16 €



Methods	Accuracy (%)	P (%)	FLOPs	Training Time
Full Network	94.57	100	x1.00	0.16 h
BBD ($kl=15$)	94.70	11.09	x2.86	0.31 h
BBD ($kl=20$)	94.30	4.86	x5.14	0.31 h
BBD ($kl=25$)	94.12	2.71	x8.62	0.31 h
STAMP ($K=1, kl=15$)	95.44	7.83	x4.26	0.11 h
STAMP ($K=5$)	95.73	3.84	x6.39	0.19 h
STAMP ($K=10$)	95.77	1.81	x8.13	0.31 h

Figure 6: **Left:** Accuracy over the ratio of used parameters for SVHN on ResNet-18. *Full* denotes the accuracy of the ResNet-18 before pruning. **Right:** Exploring different K (the number of epochs of pruning stage) of STAMP compared with BBDropout Lee et al. (2019a) (BBD). kl is a scale factor for the regularization term.

For SNIP Lee et al. (2019c), in the ResNet-18 experiment, we do not prune the 1×1 convolution layer to match the settings for STAMP experiments. Additionally, we modify the learning rate to 0.01, since at the learning rate of 0.1, SNIP (P) and SNIP obtained lower accuracies (88.51% and 85.26% respectively). For VGGNet, we prune the weights of 16 convolution layers. For SNIP (P) we load the pretrained weights on CIFAR-100 before pruning.

For MetaPruning Liu et al. (2019b), we used the same settings for ResNet-18 and ResNet-50 experiments. For VGGNet, we prune filters of 16 convolution layers which is the same as STAMP. At the search phase, we search for the architecture under given FLOP constraints. We set the pruning ratio at each layer between 20 % to 60 %, to satisfy the FLOP constraints, which is 40 % to 80 % for the given setting. For the rest of the experimental settings, we followed the settings in Liu et al. (2019b).

C.2 EXPERIMENTAL RESULTS

We report the experimental results on SVHN with ResNet18 in Table 3, which was omitted from the main paper due to the page limit. We followed the settings of Liu et al. Liu et al. (2017) and trained on SVHN for 20 epochs. All other settings are kept the same as the experimental setting in the previous paragraph. The results show that STAMP has the best trade-off between the accuracy and the efficiency.

STAMP obtains higher accuracy over BBDropout at the same compression rate as shown in Figure 6). Further, when trained for larger number of epochs, STAMP can obtain even higher accuracy and larger compression rate over BBD as shown in Figure 6, outperforming all baselines in Table 3. Although training STAMP for longer epochs yields slightly higher training time than the time required to train the full network, STAMP still trains faster than BBDropout.