

PAD-Net: An Efficient Framework for Dynamic Networks

Shwai He¹ Liang Ding^{2*} Daize Dong³ Boan Liu⁴ Fuqiang Yu⁵ Dacheng Tao²

¹University of Maryland, College Park ²The University of Sydney

³University of Electronic Science and Technology of China

⁴Wuhan University ⁵Shandong University

shwaihe@umd.edu, liangding.liam@gmail.com

Abstract

Dynamic networks, e.g., Dynamic Convolution (DY-Conv) and the Mixture of Experts (MoE), have been extensively explored as they can considerably improve the model’s representation power with acceptable computational cost. The common practice in implementing dynamic networks is to convert the given static layers into fully dynamic ones where all parameters are dynamic (at least within a single layer) and vary with the input. However, such a fully dynamic setting may cause redundant parameters and high deployment costs, limiting the applicability of dynamic networks to a broader range of tasks and models. The main contributions of our work are challenging the basic commonsense in dynamic networks and proposing a partially dynamic network, namely PAD-Net, to transform the redundant dynamic parameters into static ones. Also, we further design Iterative Mode Partition to partition dynamic and static parameters efficiently. Our method is comprehensively supported by large-scale experiments with two typical advanced dynamic architectures, i.e., DY-Conv and MoE, on both image classification and GLUE benchmarks. Encouragingly, we surpass the fully dynamic networks by +0.7% top-1 acc with only 30% dynamic parameters for ResNet-50 and +1.9% average score in language understanding with only 50% dynamic parameters for BERT. Code will be released at: <https://github.com/Shwai-He/PAD-Net>.

1 Introduction

Deep neural networks have been continuously pushing the state-of-the-art performance in the tasks of computer vision (Girshick et al., 2014; Dosovitskiy et al., 2021) and natural language processing (Dai and Le, 2015; Brunet et al., 2019; Zan et al., 2022; Zhong et al., 2022b) in past years, at the cost of increasing training cost (Shen et al., 2023). However, most prevalent architectures perform infer-

ence statically where both the computational graph and network parameters are fixed once after training, which limits the representation power. Dynamic networks (Han et al., 2021), as opposed to static ones, adapt their parameters or architectures to each specific input, improving the model representation power with acceptable computational cost, e.g., Switch Transformers (Fedus et al., 2021). The common practice of implementing dynamic networks is transforming static networks (or modules) with counterpart dynamic ones. For example, Dynamic Convolution (DY-Conv) (Chen et al., 2020b) replaces traditional convolution by adopting k adaptive convolutional kernels; Mixture of Experts (MoE) (Shazeer et al., 2017) replaces a single fully connected layer with multiple feed-forward neural networks in a parallel manner.

The success of dynamic networks motivates practitioners to design dynamic networks (Fedus et al., 2021; Li et al., 2021a), which often follow a fully dynamic approach, where all parameters are dynamic (at least within a single layer) and change with the input. Previous works (Chen et al., 2020b) show that dynamic networks often outperform their static counterpart, and using more dynamic layers intriguingly leads to ever-increasing performance. For instance, dynamic convolution promotes the performance on the ImageNet when more static convolution layers turn into dynamic ones. However, these prior works do not explain the need for a fully dynamic mechanism, and it remains unclear whether to convert static parameters into dynamic and to what extent if yes.

On the other hand, such a fully dynamic manner is resource expensive and may cause redundancy, limiting the applicability of dynamic networks. For instance, the total parameters of BERT-Base equipped with MoE are ~506.3M (with 8 experts) compared to only ~108.9M for vanilla BERT-Base. In addition, an MoE layer also multiplies the computation. It seems more is better

*Corresponding author

when transforming static layers into dynamic ones, but how about the dynamic parameters within a dynamic network: Do all of them lead to the promotion? This urges us to reflect *whether there exist redundant dynamic parameters, in fully dynamic network layers?* Based on the above scrutinization, we hypothesize that less is more for dynamic parameters in fully dynamic networks.

Motivated by this hypothesis, we propose the Iterative Mode Partition (IMP) algorithm to progressively convert less important dynamic parameters into static ones for higher efficiency, while maintaining performance at a competitive level. Given a fully dynamic network initialized with all parameters in dynamic mode, we attempt to partition a subset of static parameters out from them. Specifically, we iteratively transform dynamic parameters based on their impact on loss values. If the transformation of the i -th element of dynamic parameters results in only a minimal loss difference, we safely make it static. Given a desired dynamic ratio (the proportion of dynamic parameters), we can balance the trade-off between dynamic and static parameters. Since static parameters are less costly to deploy, we prune redundant parameters after mode partition, obtaining a lightweight architecture, namely Partially Dynamic Networks (**PAD-Net**), which contains two modes of parameters (dynamic parameters that vary with inputs and static parameters that are fixed during inference).

Empirically, we extensively validate this hypothesis and our proposed PAD-Net, including GLUE benchmark (Wang et al., 2019) for MoE and visual image classification (Deng et al., 2009) for dynamic convolution. Experiment results reveal that we successfully converted redundant dynamic parameters into static ones and PAD-Net achieves the highest performance in all tasks with lightweight architectures. Given the superiority of PAD-Net in both effectiveness and efficiency, we show that less dynamic is more efficient in fully dynamic networks, successfully verifying the above hypothesis. The inspiration of partially dynamic can be extended to other dynamic networks and even inform future efficient architectures designation.

In short, our contributions are threefold:

- We hypothesize that a fully dynamic network contains partially dynamic subnetworks that maintain or exceed the representation power of the original network.
- Following our hypothesis, we propose the

novel PAD-Net framework to achieve a partially dynamic mechanism and devise an *Iterative Mode Partition* (IMP) algorithm to partition static and dynamic parameters.

- We empirically validate our hypothesis and PAD-Net on both NLP and CV tasks across two representative dynamic networks, including MoE and dynamic convolution.

2 Related Work

Dynamic Networks. The dynamic neural network is an emerging research topic in deep learning, which adapts structures or parameters to different inputs, leading to notable advantages in terms of accuracy, and computational efficiency. Han et al. (2021) classify dynamic networks into two categories: dynamic architecture networks and dynamic parameter networks. Dynamic architecture networks adaptively adjust architectures conditioned on each sample. Specifically, they adjust the network depth (Wang et al., 2018), width (Mullapudi et al., 2018), or route based on the input (Huang et al., 2018). Instead of changing the model architecture, dynamic parameter networks boost representation power by adapting parameters or activation functions to the input (Yang et al., 2019; Liu et al., 2021). Existing works often transform various types of static parameters into dynamic versions (Chen et al., 2020b). Among them, dynamic convolution and mixture-of-experts are the typical examples that aggregate multiple convolution parameters (and experts) dynamically based on the input, leading to significant improvement with negligible computational cost.

Network Pruning. Past works in network pruning have explored effective techniques to find efficient subnetworks (Lee et al., 2019; Evci et al., 2020; He et al., 2022) and zero out redundant parameters. According to the lottery ticket hypothesis (LTH) pioneered by Frankle and Carbin (2019), dense, randomly initialized, feed-forward networks contain the subnetwork (winning tickets) that maintains comparable test performance of the original network after training for the same iterations. This hypothesis inspires a series of follow-up works in network pruning. However, these methods always sacrifice performance because of pruned parameters. As for dynamic networks, instead of directly pruning dynamic parameters, we considered changing them to static ones. In Section 5.4, we show

our approach significantly and consistently outperforms fully dynamic networks in the GLUE benchmark (Wang et al., 2019), while the pruned model performed worse than the original network.

3 Review of Fully Dynamic Networks

Basic Concept. Dynamic networks first adjust computational parameters and then compute the input using adjusted parameters, rather than directly using intrinsic parameters to compute the input. In a fully dynamic network, all intrinsic parameters are used as dynamic factors to generate computational parameters $\hat{\Theta}$, which are dependent on two parts: the input \mathbf{x} and the intrinsic parameters Θ . Let us denote \mathcal{W} as the dynamic function, computational parameters is formulated as $\hat{\Theta} = \mathcal{W}(\mathbf{x}, \Theta)$. Given an input sample \mathbf{x} , the output of is $\mathbf{y} = \mathcal{F}(\mathbf{x}, \Theta)$ for a conventional network with static parameters and $\mathbf{y} = \mathcal{F}(\mathbf{x}, \hat{\Theta})$ for a dynamic network.

Existing dynamic networks, though using different dynamic functions, tend to follow a fully dynamic manner: Networks take all intrinsic parameters to generate the computational parameters where all elements are dynamic and vary with the input. We call such networks fully dynamic networks and, in the following, introduce instantiations coming from dynamic architecture networks, i.e., *Mixture of Experts*, and dynamic parameter networks, i.e., *Dynamic Convolution*, respectively.

Mixture of Experts. We talk about dynamic architecture networks by taking the Mixture of Experts (MoE) (Jacobs et al., 1991; Shazeer et al., 2017) as an instantiation. MoE prepares m parallel static experts with parameters $\Theta^{(i)} (i = 1, 2, \dots, m)$ and only selects n experts with the highest scores ($n \leq m$). Given a specific input, we denote $G(\mathbf{x})$ as the output scores of gating and \mathcal{T} as the indices of the selected experts. For the i -th selected expert, we denote the combination of the score $G_{\mathcal{T}_i}(\mathbf{x})$ and parameters $\Theta^{(\mathcal{T}_i)}$ as $w^{(\mathcal{T}_i)} = \{G_{\mathcal{T}_i}(\mathbf{x}), \Theta^{(\mathcal{T}_i)}\}$. The dynamic function of MoE can be represented as:

$$\mathcal{W}(\mathbf{x}, \Theta) = \{w^{(\mathcal{T}_1)}, \dots, w^{(\mathcal{T}_n)}\}, \quad (1)$$

where $w^{(\mathcal{T}_i)} = \{G_{\mathcal{T}_i}(\mathbf{x}), \Theta^{(\mathcal{T}_i)}\}$.

Dynamic Convolution. As a typical example of dynamic parameter networks, Dynamic Convolution (Chen et al., 2020b) prepares k parallel static kernels $\Theta^{(i)} (i = 1, 2, \dots, k)$ as intrinsic parameters and utilizes the linear combination of them

as the aggregated kernel. The linear scale is aggregated dynamically via a channel-wise attention block (Hu et al., 2018) denoted as Attention, so the dynamic function can be written as:

$$\mathcal{W}(\mathbf{x}, \Theta) = \sum_{i=1}^k \pi_i(\mathbf{x}) \cdot \Theta^{(i)}, \quad (2)$$

where $\pi(\mathbf{x}) = \text{Attention}(\mathbf{x})$.

Limitation Discussions. Mainstream dynamic networks usually replace static layers with fully dynamic layers, where all elements of dynamic parameters require corresponding dynamic factors co-working with input samples. However, this situation causes redundant parameters and high deployment costs, limiting the applicability of dynamic networks to a border range of resource-constrained situations and large-scale models. For this fully dynamic manner, we raise two questions: **(1) Is it necessary to pay the cost of enormous parameters and computations to aggregate dynamic parameters?** **(2) Is it necessary to make all computational parameters dynamic, to maintain the performance improvement?** We propose the Partially Dynamic Network (PAD-Net) that mixes dynamic and static parameters to answer the above questions.

4 Methodology

4.1 PAD-Net: Partially Dynamic Network

In response to the limitation of fully dynamic networks, we question whether it is necessary to make all parameters dynamic. To this end, we try to detect the less important dynamic parameters and transform them into input-agnostic static parameters. Specifically, we utilize a mask $M_i (i = 1, 2, \dots, m)$ to indicate whether the i -th element of $\hat{\Theta}$ is dynamic or static: $M_i = 1$ means the i -th element of $\hat{\Theta}$ is dynamic and vice versa. We use $\tilde{\Theta} \in \mathbb{R}^m$ to denote the dynamic parameters and $\bar{\Theta} \in \mathbb{R}^m$ to represent the static parameters, then computational parameters $\hat{\Theta}$ are reformulated as:

$$\hat{\Theta}_i = \begin{cases} \tilde{\Theta}_i = \mathcal{W}_i(\mathbf{x}, \Theta) & \text{if } M_i = 1 \\ \bar{\Theta}_i & \text{otherwise} \end{cases}, \quad (3)$$

where $\hat{\Theta}_i (i = 1, 2, \dots, m)$ represents the i -th element of $\hat{\Theta}$, and Θ denotes the dynamic factors. In our architecture, intrinsic parameters include dynamic factors Θ and static parameters $\bar{\Theta}$. Note that M partitions the computational parameters into two non-overlapping parts, forming a network with

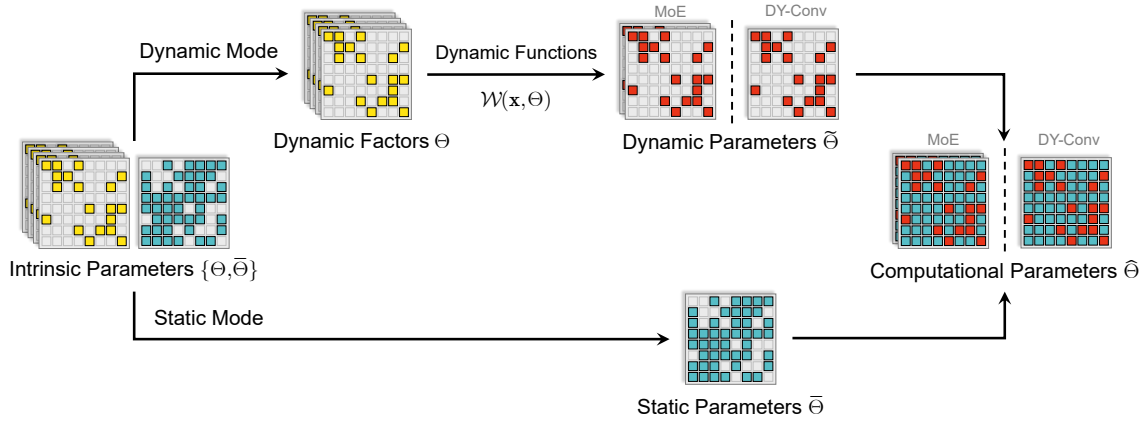


Figure 1: **The procedure of generating the computational parameters in PAD-Net, with DY-Conv and MoE as instantiations.** The intrinsic parameters include static parameters and dynamic factors. Given an input, dynamic factors activate and aggregate into dynamic parameters, which are then integrated with static parameters.

only a part of the parameters dynamic, i.e., Partially Dynamic Network (PAD-Net). Details of the procedure of generating the computational parameters from intrinsic are visualized in Figure 1.

To overcome the aforementioned challenges and limitations, we propose a novel network architecture, Partially Dynamic Network (PAD-Net). We also devise a new algorithm *Iterative Mode Partition* (IMP) to build this model efficiently.

In addition, we set two scale factors to describe the relative intensity of these subnetworks separately in terms of magnitude, namely λ_s and λ_d . With these scale factors, we factorize our method into a more general formulation:

$$\hat{\Theta}_i = \begin{cases} \lambda_d \cdot \tilde{\Theta}_i & \text{if } M_i = 1 \\ \lambda_s \cdot \bar{\Theta}_i & \text{otherwise} \end{cases}, \quad (4)$$

where we constrain $\lambda_s + \lambda_d = 2$ ($\lambda_s, \lambda_d > 0$), and Equation 3 is the special situation when both λ_s and λ_d are equal to 1. Similar to the constraint $\sum_{i=1}^k \pi_i$ in dynamic convolution (Chen et al., 2020b), this constraint compresses the parameters space and significantly simplifies the joint optimization of scale factors and the counterpart parameters.

4.2 Iterative Mode Partition

In the above section, we present the architecture of PAD-Net, which includes dynamic parameters and counterpart static parameters. Next, we further discuss our method in how to generate indicator masks to partition dynamic and static parameters. Let us first formulate this partition as an optimization problem, where our goal is to minimize loss values L . Given a dataset $\mathcal{D} = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^n$ and

a desired dynamic ratio κ of M , we briefly formulate mode partition as the following constrained optimization problem:

$$\begin{aligned} \min_M L(\hat{\Theta}, M; \mathcal{D}) &= \min_M \frac{1}{n} \sum_{i=1}^n \ell(\hat{\Theta}, M; (\mathbf{x}_i, \mathbf{y}_i)), \\ \text{s.t. } M &\in \{0, 1\}^m, \quad \|M\|_0 \leq \kappa \cdot m, \end{aligned} \quad (5)$$

where $\ell(\cdot)$ denotes the standard loss function (e.g., cross-entropy loss), $\hat{\Theta}$ is the set of computational parameters of the neural network, $\|\cdot\|_0$ is the standard L_0 norm, m is the total number of parameters. The conventional approach to optimize the above problem is adding sparsity enforcing penalty term M (Carreira-Perpinán and Idelbayev, 2018), while it often requires heavily tuned hyperparameter settings and several trials. On the other hand, LTH-based (Chen et al., 2020a; Evci et al., 2020) methods find the mask by several iterations, but it is prohibitively time-consuming. Also, considering the large-scale dynamic networks, it is unnecessary to deploy redundant parameters.

We tend to partition the two modes before training to prune redundant parameters and avoid time-consuming training iterations. Inspired by Lee et al. (2019)’s gradient-based pruning strategy, we propose an algorithm to make excessive dynamic parameters static. We resort to mini-batches of training data $\mathcal{D}_b = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^b \sim \mathcal{D}$ to detect redundant dynamic parameters. Given a dynamic parameter $\hat{\Theta}_j$ at the j -th element of $\hat{\Theta}$, we compute its importance of being dynamic based on the loss difference ΔL_j caused by making $\hat{\Theta}_j$ static (changing the value of M_j from 1 to 0):

$$\Delta L_j(M, \hat{\Theta}; \mathcal{D}_b) = L(M, \hat{\Theta}; \mathcal{D}_b) - L(M - \mathbf{t}_j, \hat{\Theta}; \mathcal{D}_b), \quad (6)$$

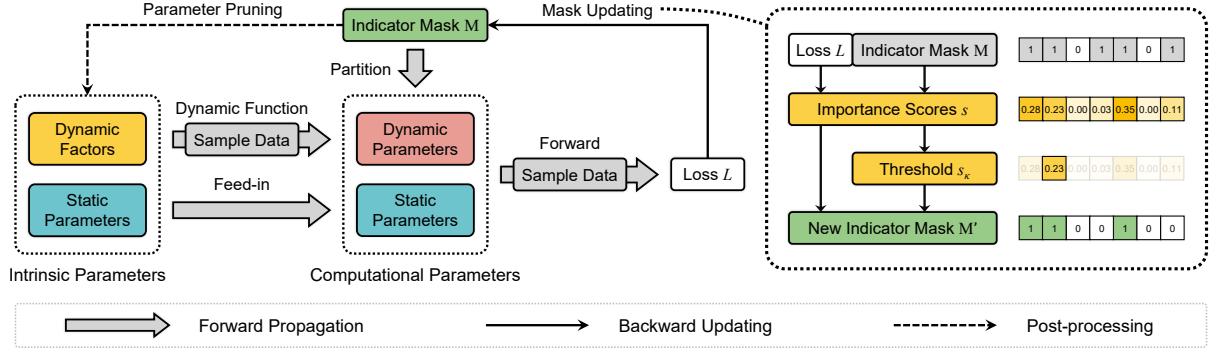


Figure 2: **Graphical illustration of Iterative Mode Partition (IMP).** *Left:* An overview of IMP, including forward propagation and backward updating. After IMP, the indicator mask prunes the redundant dynamic factors and static parameters (post-processing). *Right:* Details of mask updating.

where \mathbf{t}_j is the indicator vector of j -th element of \mathbf{M} (i.e., zeros everywhere except at the index j where it is one). We only consider transforming redundant dynamic parameters into static ones, so the loss difference ΔL_j is zero if $\hat{\Theta}_j$ is static. Note that computing ΔL_j for each dynamic parameter is prohibitively expensive, as it usually requires millions of forwarding passes over the dataset, so we resort to a simple and effective approximate alternative. Specifically, we release binary constraints of \mathbf{M} and make it differentiable and utilize the derivative of L with respect to \mathbf{M}_j to approximate ΔL_j :

$$\begin{aligned} \Delta L_j(\mathbf{M}, \hat{\Theta}; \mathcal{D}_b) &\approx g_j(\hat{\Theta}; \mathcal{D}_b) = \frac{\partial L_j(\mathbf{M}, \hat{\Theta}; \mathcal{D}_b)}{\partial \mathbf{M}} \\ &= \lim_{\delta \rightarrow 0} \frac{L_j(\mathbf{M}, \hat{\Theta}; \mathcal{D}_b) - L_j(\mathbf{M} - \delta \mathbf{t}_j, \hat{\Theta}; \mathcal{D}_b)}{\delta} \Big|_{\mathbf{t}=1}, \end{aligned} \quad (7)$$

where $g_j(\hat{\Theta}; \mathcal{D}_b)$ denotes the j -th element in derivative $g(\hat{\Theta}; \mathcal{D}_b)$. We accumulate the derivatives for all j by one forward-backward pass using automatic differentiation. Note that if the magnitude of g_j is high, it essentially means that making parameter $\hat{\Theta}_j$ static has a considerable effect on the loss, and it has to be dynamic. In contrast, the parameter should be static if the magnitude of g_j is low. Therefore, We take the normalized magnitude of the derivatives of g as the criteria:

$$s_j = \left| g_j(\hat{\Theta}; \mathcal{D}_b) \right| / \left| \sum_{k=1}^m g_k(\hat{\Theta}; \mathcal{D}_b) \right|. \quad (8)$$

Given the dynamic ratio κ , we take the s_κ (the κ -th percentile of s) as the threshold and transform the mask elements whose scores are below zero:

$$\mathbf{M} = \mathbb{1}[s - s_\kappa \geq 0], \quad (9)$$

where $\mathbb{1}[\cdot]$ is an element-wise indicator function where the output will be 1 if the condition $[\cdot]$ meets

else it will be zero. Note that the indicator mask \mathbf{M} prunes out redundant parameters in dynamic parameters $\tilde{\Theta}$ and static parameters $\bar{\Theta}$ respectively. Also, for fewer dynamic parameters to generate, we can also utilize the binary mask to prune redundant dynamic factors. Taking MoE as an example, \mathbf{M} can be directly applied to parallel experts: $\Theta^{(i)} \leftarrow \mathbf{M} \odot \Theta^{(i)}, \forall i \in \{1, 2, \dots, k\}$. In addition, we can decrease the computational cost of generating based on dynamic factors.

Inspired by the success of the iterative strategy in pruning at initialization (Verdenius et al., 2020; de Jorge et al., 2021), we start from a fully dynamic network and adopt an iterative strategy shown in Figure 2 to transform dynamic parameters into static parameters iteratively, where we increase the zero ratios of \mathbf{M} exponentially. The effectiveness of the mode partition and the iterative mode partition is experimentally verified in Section 5.3.

5 Empirical Evaluation

5.1 Implementation Details

Mixture of Experts. We use Adam (Kingma and Ba, 2015) as the optimizer with $\beta_1, \beta_2 = 0.9, 0.98$. For regularization, we set the weight decay as 0.1 and grid-search the learning rate from $\{1e-5, 5e-5, 1e-4, 5e-4\}$, where we warm up the learning rate in the first 10% steps (of the total training steps). For different data scales, we grid-search training epoch and batch size from $\{5, 10, 15, 20\}$ and $\{8, 16, 32, 64\}$, respectively. The maximum length is 128 for all tasks. Following Shazeer et al. (2017), we initialize dynamic and static parameters with pretrained parameters.

Dynamic Convolution. We use an SGD optimizer (Ruder, 2016) with 0.9 momentum, follow-

Table 1: **Comparison between PAD-Net and vanilla MoE** applied to four widely used large-scale Pretrained Language Models (PLMs). Averaged scores on all tasks are underlined. The shown results are the averaged score for 5 runs, followed by the deviation. The best results are **bold**. It shows that PAD-Net yields consistent improvements across all tasks among different MoE-equipped PLMs.

Method	BERT						ALBERT					
	#Param.	CoLA	RTE	MRPC	STS-B	Avg.	#Param.	CoLA	RTE	MRPC	STS-B	Avg.
Static	108.9M	54.6 \pm 0.4	66.4 \pm 0.7	84.6 \pm 0.3	85.8 \pm 0.3	<u>72.9</u>	11.1M	54.2 \pm 0.7	76.6 \pm 0.7	87.2 \pm 0.4	90.6 \pm 0.3	<u>77.2</u>
MoE	506.3M	58.0 \pm 0.9	69.3 \pm 1.2	85.0 \pm 0.4	87.1 \pm 0.2	<u>74.9</u>	44.2M	56.8 \pm 1.2	77.2 \pm 0.8	87.4 \pm 0.4	90.7 \pm 0.3	<u>78.0</u>
PAD-Net	308.0M	59.7 \pm 0.8	71.5 \pm 1.4	85.5 \pm 0.4	90.3 \pm 0.6	76.8	30.0M	57.4 \pm 1.4	77.6 \pm 0.5	88.4 \pm 0.3	90.9 \pm 0.2	78.6

Method	RoBERTa						ELECTRA					
	#Param.	CoLA	RTE	MRPC	STS-B	Avg.	#Param.	CoLA	RTE	MRPC	STS-B	Avg.
Static	124.1M	62.8 \pm 1.0	77.6 \pm 1.6	90.0 \pm 0.5	91.0 \pm 0.3	<u>80.4</u>	108.9M	67.3 \pm 1.5	82.6 \pm 1.7	89.0 \pm 0.5	90.6 \pm 0.1	<u>82.4</u>
MoE	521.5M	63.6 \pm 1.1	78.0 \pm 1.4	90.2 \pm 0.4	91.1 \pm 0.2	<u>80.8</u>	506.3M	67.6 \pm 1.1	83.0 \pm 1.4	89.3 \pm 0.3	90.8 \pm 0.2	<u>82.7</u>
PAD-Net	323.1M	64.2 \pm 0.8	79.4 \pm 1.2	90.7 \pm 0.3	91.4 \pm 0.3	81.4	308.0M	68.2 \pm 1.3	84.1 \pm 1.5	89.5 \pm 0.4	91.2 \pm 0.2	83.3

Table 2: **Comparison between PAD-Net and baselines for ResNet and MobileNetV2**, including CondConv and DY-Conv. The Top-1 accuracy is the averaged score for 5 runs, followed by the deviation. * indicates the dynamic model with the fewest parameters or the fewer FLOPs (the static model is not included), and the best results in accuracy are **bold**. DY-Conv and PAD-Net contain $k = 4$ kernels, while CondConv contains $k = 8$ kernels.

Depth	Model	Params	FLOPs	Top-1(w/dev)	Width	Model	Params	FLOPs	Top-1(w/dev)
ResNet-10	Static	5.2M	0.89G	63.1 \pm 0.4	$\times 0.5$	Static	2.0M	97.0M	65.7 \pm 0.3
	CondConv	36.7M	0.92G	66.9 \pm 0.2		CondConv	15.5M	113.0M	68.8 \pm 0.2
	DY-Conv	18.6M	0.91G	67.4 \pm 0.3		DY-Conv	4.0M	101.4M	69.6 \pm 0.1
	PAD-Net	*6.9M	*0.90G	68.1 \pm 0.2		PAD-Net	*2.7M	*98.3M	70.4 \pm 0.2
ResNet-18	Static	11.1M	1.81G	70.6 \pm 0.3	$\times 0.75$	Static	2.6M	209.1M	69.2 \pm 0.4
	CondConv	81.4M	1.89G	71.9 \pm 0.2		CondConv	17.5M	233.9M	72.1 \pm 0.3
	DY-Conv	42.7M	1.86G	72.4 \pm 0.3		DY-Conv	8.0M	220.1M	72.6 \pm 0.1
	PAD-Net	*15.1M	*1.83G	73.0 \pm 0.3		PAD-Net	*5.2M	*212.4M	73.5 \pm 0.2
ResNet-50	Static	23.5M	3.86G	76.2 \pm 0.2	$\times 1.0$	Static	3.5M	300.8M	72.1 \pm 0.3
	CondConv	129.9M	3.98G	76.9 \pm 0.3		CondConv	27.5M	329.0M	74.4 \pm 0.2
	DY-Conv	100.9M	3.97G	77.2 \pm 0.2		DY-Conv	11.1M	312.9M	74.8 \pm 0.2
	PAD-Net	*33.8M	*3.90G	77.9 \pm 0.2		PAD-Net	*6.1M	*304.4M	75.3 \pm 0.1

(a) ResNet

(b) MobileNetV2

ing cosine learning rate scheduling and warmup strategy. The learning rate rises to the maximum linearly in the first ten epochs and schedules to arrive at zero within a single cosine cycle. We follow Chen et al. (2020b)’s temperature annealing strategy to avoid the unstable output values of the softmax function in the first epochs. We train ResNet for 100 epochs with the max learning rate of 0.1. We train the MobilenetV2 for 300 epochs with the max learning rate of 0.05. The weight decay is $1e-4$ for ResNet and $4e-5$ for MobilenetV2. The training batch size is 256 for all models.

5.2 Main Results

Natural Language Understanding. We evaluate the performance of PAD-Net for MoE on various datasets from the General Language Understanding Evaluation (GLUE) benchmark (Wang et al., 2019).

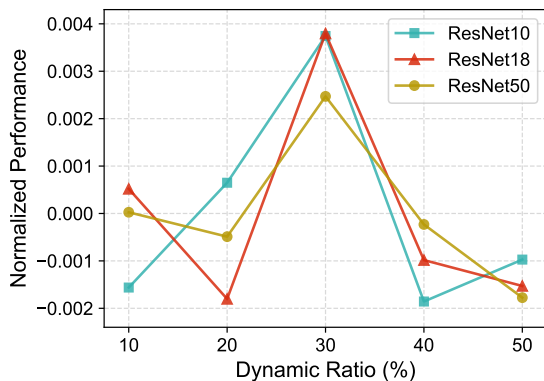
Like previous works (Lee et al., 2020; Dodge et al., 2020; Zhong et al., 2022a), we fine-tune pretrained models, e.g., BERT (Devlin et al., 2019), ALBERT (Lan et al., 2020), RoBERTa (Liu et al., 2019), ELECTRA (Clark et al., 2020) on the training set and directly report results on validation set using the last checkpoint, since the test results are only accessible by the leaderboard with submission limitation.

Following Shazeer et al. (2017); Gao et al. (2022), we replace feed-forward layers with MoE layers where we prepare 8 experts and select the top-2 experts for each input. We set the dynamic ratio $\kappa = 50\%$ because it is close to the optimal value. Table 1 shows that PAD-Net outperforms MoE on the GLUE benchmark with a 0.95% average increase for four backbones. Specifically, PAD-Net improves BERT by 1.9% and RoBERTa by 0.6%

Table 3: **Ablation study for dynamic ratio on MoE integrated with PAD-Net.** Averaged scores on all tasks are underlined. The shown results are the averaged score for 5 runs. The best results are **bold**. Methods under the dashline are our proposed PAD-Net, where κ denotes the dynamic ratio.

Method	#Param.	CoLA	SST-2	MRPC	STS-B	QQP	MNLI	QNLI	RTE	Avg.
BERT	108.9M	54.6	91.4	84.6	85.8	90.6	83.7	90.4	66.4	<u>81.2</u>
w/ MoE	506.3M	58.0	91.7	85.0	87.1	90.8	83.8	90.8	69.3	<u>82.1</u>
$\kappa = 70\%$	387.3M	58.5	92.4	85.5	89.6	90.9	83.9	90.9	70.6	<u>82.8</u>
$\kappa = 50\%$	308.0M	59.7	92.2	85.4	90.3	90.9	84.2	91.0	71.5	83.2
$\kappa = 30\%$	228.6M	59.0	92.0	85.3	89.4	91.0	84.0	90.9	71.2	<u>82.9</u>
$\kappa = 10\%$	149.3M	57.5	91.1	85.4	88.3	90.4	83.6	90.6	70.2	<u>82.1</u>
RoBERTa	124.1M	62.1	94.0	89.6	90.6	91.0	86.9	91.8	77.4	<u>85.4</u>
w/ MoE	521.5M	63.6	94.8	90.2	91.1	91.7	87.7	92.9	78.0	<u>86.3</u>
$\kappa = 70\%$	402.5M	64.6	95.0	91.0	91.0	91.8	87.7	92.9	78.2	<u>86.5</u>
$\kappa = 50\%$	323.1M	64.4	95.2	90.7	91.4	91.9	88.0	93.0	79.4	86.8
$\kappa = 30\%$	243.8M	63.4	94.6	90.5	91.2	91.4	87.8	93.2	78.8	<u>86.4</u>
$\kappa = 10\%$	164.5M	63.9	94.4	90.4	90.8	90.9	87.4	92.6	78.2	<u>86.1</u>

Figure 3: **Normalized performance of ResNet with different dynamic ratios**, which is evaluated by $\frac{x-\bar{x}}{\bar{x}}$ where \bar{x} is the mean accuracy across experiments.



on average. Equipped with PAD-Net, MoE reduces both parameters and computation significantly, and we provide a theoretical analysis of the reduced computation in Section 5.4.

Visual Image Classification. We also report the superiority of PAD-Net in visual image classification. In Table 2, we compare PAD-Net with static convolution (Krizhevsky et al., 2012), Cond-Conv (Yang et al., 2019) and Dynamic Convolution (Chen et al., 2020b) on ImageNet (Deng et al., 2009) classification for ResNet (He et al., 2016) and MobileNetV2 (Sandler et al., 2018) in the same experimental setting with previous works, by adjusting all convolution layers except the first layer. Before training, we first partition two modes of parameters with a given dynamic ratio κ using ten batches of examples.

PAD-Net improves accuracy with significantly

Table 4: **Ablation study of scale factors**, where ‘‘Option’’ refers to the setting of scale factors.

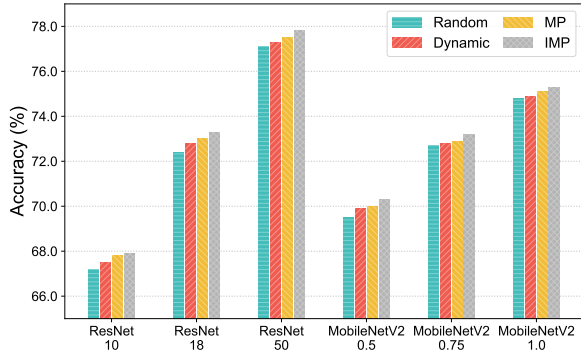
Model	Option	RTE	STS-B
BERT-base	–	69.6	87.4
	λ_s	70.7	88.1
	λ_d	70.9	89.6
	λ_s, λ_d	71.3	89.8
	$\lambda_s + \lambda_d = 2$	71.5	90.3
Model	Option	CIFAR-10	ImageNet
ResNet-50	–	93.9	77.1
	λ_s	94.3	77.2
	λ_d	94.5	77.4
	λ_s, λ_d	96.0	77.6
	$\lambda_s + \lambda_d = 2$	96.6	77.8

lighter architecture and fewer FLOPs (Floating Point Operations). For instance, PAD-Net outperforms DY-Conv by 0.7% top-1 accuracy with 33.9% parameters and 0.1G fewer FLOPs in ResNet-50.

5.3 Ablation Study

Effect of Dynamic Ratio. Inspired by Wettig et al. (2022), we investigate the impact of different dynamic ratios κ , and the results are shown in Table 3 For MoE and Figure 3 for DY-Conv. Because PAD-Net with low dynamic ratios significantly outperforms fully dynamic networks, we only consider ratios of less than 70%, allowing for more sparsity and efficiency. We empirically find that $\kappa = 50\%$ is nearly the optimal ratio for MoE to achieve the highest performance, while the best performance of DY-Conv is achieved when $\kappa = 30\%$. We believe that different dynamic functions contribute

Figure 4: **Comparison of different partition methods**, including random partition “Random”, mode partition “MP”, and iterative mode partition “IMP”. We also report dynamic convolution “Dyngmiac” as a baseline.

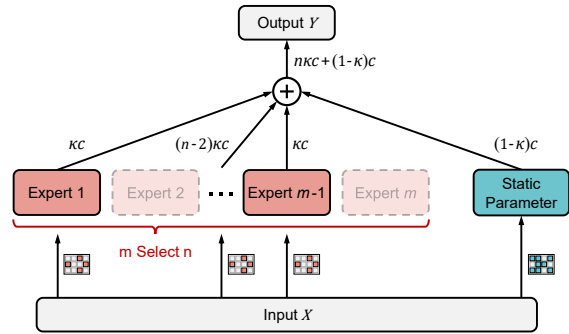


to different optimal dynamic ratios, and an efficient way, e.g., hyper-parameter optimization or meta-learning, to search them will be necessarily explored in the future.

Effect of Scale Factors. We also conduct an ablation study on the proposed scale factors and verify their necessity. Table 4 summarizes the impact of scale factors on different architectures. We initially tried to gain scale factors from a SENet structure (Hu et al., 2018), while it did not contribute to the improvement of performance. So we just set scale factors as trainable parameters to avoid redundant parameters and operations. Besides the setting “ $\lambda_s + \lambda_d = 2$ ” in Equation 4, we consider other situations: only using one factor (“ λ_s ” and “ λ_d ”), and no scale factors used (“-”). We conduct experiments on CIFAR-10 (Krizhevsky, 2009) and ImageNet for ResNet-50, RTE, and STS-B for BERT. λ_s and λ_d enhance performance substantially, and their coexistence leads to further improvement. To explore the impact of the summation constraint, we release it and denote this setting as “ λ_s, λ_d ”. Clearly, without summation constraint, the performance of ResNet-50 and BERT decreases significantly, i.e., -0.4% and -0.35% on average.

Effectiveness of Iterative Mode Partition. We compare different partition strategies in Figure 4. Compared to fully dynamic networks, accuracy degrades when we partition two modes randomly, which means this naive partition method mistakes some important dynamic parameters. In contrast, mode partition contributes to a better combination of dynamic and static parameters, improving the accuracy. IMP shows its effectiveness by achieving the best performance.

Figure 5: **Visualization description of the computation cost for PAD-Net on MoE.** Given a specific input X , we denote the computation cost for selected experts and static parameters.



5.4 Detailed Analysis

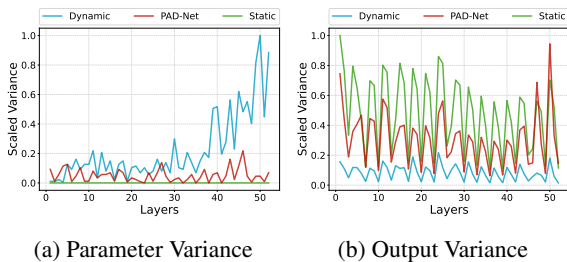
Reduced Computation. We show the computation cost of PAD-Net in Figure 5. Compared to vanilla MoE, PAD-Net reduces the computation between selected experts and the input, $y_{T_i} = E_{T_i}(x)$, where E_{T_i} denotes the i -th selected experts. Because the two methods share the same gating mechanism, we temporarily ignore its computation for simplicity. We denote the computation of the i -th expert as C_{T_i} where $C_{T_1} = \dots = C_{T_n} = c$, and the total computation of multi-experts is nc if we select n experts within m ones. In PAD-Net, given the dynamic ratio κ , it is reduced to $n\kappa c$. Together with the computation $(1-\kappa)c$, the computation of a PAD-Net layer is $n\kappa c + (1-\kappa)c$. Integrated with PAD-Net, an MoE layer can reduce the computation by $(n-1)(1-\kappa)c$. When κ is low enough, the computation of PAD-Net can be close to static networks. For DY-Conv, the reduced computation lies in the linear combination of parallel kernels $\sum_{i=1}^k \pi_i(\mathbf{x}) \cdot \Theta^{(i)}$, which is sparse in PAD-Net. In short, the degree of reduced computation depends on the specific dynamic function.

Difference with Model Pruning. Mode partition maintains important dynamic parameters while making redundant ones static, which may be similar to network pruning. In Table 5, we compare mode partition with network pruning (Lee et al., 2019) on the GLUE benchmark for BERT-base and reveal their difference empirically. PAD-Net achieves the best performance among all tasks listed, with 1.2% average improvements over vanilla MoE. In contrast, we discover that network pruning lowers the performance of MoE significantly by 1.1% on average. Considering maintain-

Table 5: **Empirical comparison between our PAD-Net and model pruning on the GLUE benchmark.** PAD-Net is set with $\kappa = 50\%$, and MoE-P is an MoE architecture pruned by SNIP (Lee et al., 2019). We make the parameters of PAD-Net and MoE-P consistent for a fair comparison.

Method	#Param.	CoLA	SST-2	MRPC	STS-B	QQP	MNLI	QNLI	RTE	Avg.
Static	108.9M	54.6	91.4	84.6	85.8	90.6	83.7	90.4	66.4	81.2
MoE	506.3M	58.0	91.7	85.0	87.1	90.8	83.8	90.4	69.3	82.0
MoE-P	308.0M	55.6	91.6	84.7	85.8	90.8	82.4	90.2	65.7	80.9
PAD-Net		59.7	92.2	85.4	90.3	90.9	84.2	91.0	71.5	83.2

Figure 6: **Dynamic property calculation.** We plot layer-wise curves of parameter variance and output variance for ResNet-50.



ing the performance of a fully dynamic network, it is preferable to convert unimportant dynamic parameters into static ones than to prune them.

Dynamic Property. Dynamic property refers to variant numerical characteristics of a dynamic network caused by different input samples. The ideal dynamic network maintains two capacities: assigning specific parameters for the input and making counterpart output discriminative. Inspired by Li et al. (2021b), we take two levels of variance as metrics (parameter variance and output variance) to measure the dynamic property and show the result in Figure 6. Static convolution, dynamic convolution, and PAD-Net ($\kappa = 30\%$) show different properties given the same samples from ImageNet. We see that dynamic convolution retains a high degree of parameter variance while it has the lowest output variance. Static convolution performs the opposite. The outputs of PAD-Net are discriminative, which may contribute to its superiority in performance.

6 Conclusion and Future Work

In this work, we first reveal parameter redundancy and high deployment costs of fully dynamic networks. To resolve these problems, we proposed the partially dynamic network (PAD-Net) to advance both performance and efficiency. PAD-Net

demonstrated its superiority on MoE and DY-Conv frameworks. Extensive experiments on both NLP and CV tasks empirically show its effectiveness and efficiency against fully dynamic networks, significantly improving performance with much fewer dynamic parameters and less computation. Our proposed method could be extensively integrated with other mainstream architectures and inspire future work in efficient neural network designation and other fields.

Acknowledgements

We are grateful to the anonymous ACL reviewers and the area chair for their insightful comments and suggestions.

7 Limitations

Despite the progress we made, there still exist limitations in our work. On the one hand, we only investigated some classic dynamic networks and found that the proposed method contribute to the best performance in selected criteria. However, other advanced partition methods that further improve the performance and efficiency may exist, which deserve exploration in future work. On the other hand, since we only consider MoE and DY-Conv in limited tasks, it would be valuable to consider other architectures (e.g., Switch Transformer (Fedus et al., 2021)), machine learning methods (e.g., reinforcement learning (Li et al., 2022)) and tasks (e.g., machine translation (Ding et al., 2020, 2021)).

Ethics Statement

We take ethical considerations seriously and strictly adhere to the ACL Ethics Policy. This paper focuses on the higher efficiency of dynamic networks, e.g., the mixture of experts. Both the datasets and models used in this paper are publicly available and have been widely adopted by researchers. We ensure that the findings and conclusions of this paper are reported accurately and objectively.

References

- Marc-Etienne Brunet, Colleen Alkalay-Houlihan, Ashton Anderson, and Richard Zemel. 2019. [Understanding the origins of bias in word embeddings](#). In *ICML*.
- Miguel A Carreira-Perpinán and Yerlan Idelbayev. 2018. [“learning-compression” algorithms for neural net pruning](#). In *CVPR*.
- Tianlong Chen, Jonathan Frankle, Shiyu Chang, Sijia Liu, Yang Zhang, Zhangyang Wang, and Michael Carbin. 2020a. [The lottery ticket hypothesis for pre-trained bert networks](#). In *NeurIPS*.
- Yinpeng Chen, Xiyang Dai, Mengchen Liu, Dongdong Chen, Lu Yuan, and Zicheng Liu. 2020b. [Dynamic convolution: Attention over convolution kernels](#). In *CVPR*.
- Kevin Clark, Minh-Thang Luong, Quoc V Le, and Christopher D Manning. 2020. [Electra: Pre-training text encoders as discriminators rather than generators](#). In *ICLR*.
- Andrew M Dai and Quoc V Le. 2015. [Semi-supervised sequence learning](#). In *NeurIPS*.
- Pau de Jorge, Amartya Sanyal, Harkirat S Behl, Philip HS Torr, Gregory Rogez, and Puneet K Dokia. 2021. [Progressive skeletonization: Trimming more fat from a network at initialization](#). *ICLR*.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. [Imagenet: A large-scale hierarchical image database](#). In *CVPR*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [Bert: Pre-training of deep bidirectional transformers for language understanding](#). In *NAACL*.
- Liang Ding, Longyue Wang, and Dacheng Tao. 2020. [Self-attention with cross-lingual position representation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1679–1685, Online. Association for Computational Linguistics.
- Liang Ding, Di Wu, and Dacheng Tao. 2021. [Improving neural machine translation by bidirectional training](#). *EMNLP*.
- Jesse Dodge, Gabriel Ilharco, Roy Schwartz, Ali Farhadi, Hannaneh Hajishirzi, and Noah Smith. 2020. [Fine-tuning pretrained language models: Weight initializations, data orders, and early stopping](#). *arXiv*.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. 2021. [An image is worth 16x16 words: Transformers for image recognition at scale](#). In *ICLR*.
- Utku Evci, Trevor Gale, Jacob Menick, Pablo Samuel Castro, and Erich Elsen. 2020. [Rigging the lottery: Making all tickets winners](#). In *ICML*.
- William Fedus, Barret Zoph, and Noam Shazeer. 2021. [Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity](#). *arXiv*.
- Jonathan Frankle and Michael Carbin. 2019. [The lottery ticket hypothesis: Finding sparse, trainable neural networks](#). *ICLR*.
- Ze-Feng Gao, Peiyu Liu, Wayne Xin Zhao, Zhong-Yi Lu, and Ji-Rong Wen. 2022. [Parameter-efficient mixture-of-experts architecture for pre-trained language models](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 3263–3273, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. 2014. [Rich feature hierarchies for accurate object detection and semantic segmentation](#). In *CVPR*.
- Yizeng Han, Gao Huang, Shiji Song, Le Yang, Honghui Wang, and Yulin Wang. 2021. [Dynamic neural networks: A survey](#). *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. [Deep residual learning for image recognition](#). In *CVPR*.
- Shwai He, Liang Ding, Daize Dong, Miao Zhang, and Dacheng Tao. 2022. [Sparseadapter: An easy approach for improving the parameter-efficiency of adapters](#). In *EMNLP*.
- Jie Hu, Li Shen, and Gang Sun. 2018. [Squeeze-and-excitation networks](#). In *CVPR*.
- Gao Huang, Shichen Liu, Laurens Van der Maaten, and Kilian Q Weinberger. 2018. [Condensenet: An efficient densenet using learned group convolutions](#). In *CVPR*.
- Robert A. Jacobs, Michael I. Jordan, Steven J. Nowlan, and Geoffrey E. Hinton. 1991. [Adaptive mixtures of local experts](#). *Neural Computation*.
- Diederik P Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *ICLR*.
- Alex Krizhevsky. 2009. [Learning multiple layers of features from tiny images](#). Technical report.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. [Imagenet classification with deep convolutional neural networks](#). In *NeurIPS*.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. [Albert: A lite bert for self-supervised learning of language representations](#). In *ICLR*.

- Cheolhyoung Lee, Kyunghyun Cho, and Wanmo Kang. 2020. [Mixout: Effective regularization to finetune large-scale pretrained language models](#). *ICLR*.
- Namhoon Lee, Thalaiyasingam Ajanthan, and Philip HS Torr. 2019. [Snip: Single-shot network pruning based on connection sensitivity](#). In *ICLR*.
- Chao Li, Aojun Zhou, and Anbang Yao. 2021a. [Omni-dimensional dynamic convolution](#). In *ICLR*.
- Qian Li, Hao Peng, Jianxin Li, Jia Wu, Yuanxing Ning, Lihong Wang, Philip S. Yu, and Zheng Wang. 2022. [Reinforcement learning-based dialogue guided event extraction to exploit argument relations](#). *IEEE ACM Trans. Audio Speech Lang. Process.*
- Yunsheng Li, Yinpeng Chen, Xiyang Dai, Mengchen Liu, Dongdong Chen, Ye Yu, Lu Yuan, Zicheng Liu, and Others. 2021b. [Revisiting dynamic convolution via matrix decomposition](#). In *ICLR*.
- Chuan Liu, Yi Gao, and Jiancheng Lv. 2021. [Dynamic normalization](#). *arXiv*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Others. 2019. [Roberta: A robustly optimized bert pretraining approach](#). *arXiv*.
- Ravi Teja Mullapudi, William R Mark, Noam Shazeer, and Kayvon Fatahalian. 2018. [Hydranets: Specialized dynamic architectures for efficient inference](#). In *CVPR*.
- Sebastian Ruder. 2016. [An overview of gradient descent optimization algorithms](#). *arXiv*.
- Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. 2018. [Mobilenetv2: Inverted residuals and linear bottlenecks](#). In *CVPR*.
- Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc V. Le, Geoffrey E. Hinton, and Jeff Dean. 2017. [Outrageously large neural networks: The sparsely-gated mixture-of-experts layer](#). In *ICLR*.
- Li Shen, Yan Sun, Zhiyuan Yu, Liang Ding, Xinmei Tian, and Dacheng Tao. 2023. [On efficient training of large-scale deep learning models: A literature review](#). *arXiv*.
- Stijn Verdenius, Maarten Stol, and Patrick Forré. 2020. [Pruning via iterative ranking of sensitivity statistics](#). *arXiv*.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. 2019. [Glue: A multi-task benchmark and analysis platform for natural language understanding](#). In *ICLR*.
- Xin Wang, Fisher Yu, Zi-Yi Dou, Trevor Darrell, and Joseph E Gonzalez. 2018. [Skipnet: Learning dynamic routing in convolutional networks](#). In *ECCV*.
- Alexander Wettig, Tianyu Gao, Zexuan Zhong, and Danqi Chen. 2022. [Should you mask 15% in masked language modeling?](#) *arXiv*.
- Brandon Yang, Gabriel Bender, Quoc V. Le, and Jiquan Ngiam. 2019. [Condconv: Conditionally parameterized convolutions for efficient inference](#). In *NeurIPS*.
- Changtong Zan, Keqin Peng, Liang Ding, Baopu Qiu, Boan Liu, Shwai He, Qingyu Lu, Zheng Zhang, Chuang Liu, Weifeng Liu, Yibing Zhan, and Dacheng Tao. 2022. [Vega-MT: The JD explore academy machine translation system for WMT22](#). In *WMT*.
- Qihuang Zhong, Liang Ding, Li Shen, Peng Mi, Juhua Liu, Bo Du, and Dacheng Tao. 2022a. [Improving sharpness-aware minimization with fisher mask for better generalization on language models](#). In *EMNLP*.
- Qihuang Zhong, Liang Ding, Yibing Zhan, Yu Qiao, Yonggang Wen, Li Shen, Juhua Liu, Baosheng Yu, Bo Du, Yixin Chen, et al. 2022b. [Toward efficient language model pretraining and downstream adaptation via self-evolution: A case study on superglue](#). *arXiv*.

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
7
- A2. Did you discuss any potential risks of your work?
Both the methods we cite and the improvement we make involve no potential risk.
- A3. Do the abstract and introduction summarize the paper’s main claims?
1
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

Left blank.

- B1. Did you cite the creators of artifacts you used?
Left blank.
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
Left blank.
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
Left blank.
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
Left blank.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
Left blank.
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
Left blank.

C Did you run computational experiments?

5.1 5.2

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
5.1 5.2

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

Appendix A

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

Left blank.

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

Left blank.

D Did you use human annotators (e.g., crowdworkers) or research with human participants?

Left blank.

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

Left blank.

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

Left blank.

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

Left blank.

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

Left blank.

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

Left blank.