

---

# AdaXpert: Adapting Neural Architecture for Growing Data

---

Shuaicheng Niu<sup>\*1,2</sup> Jiayang Wu<sup>\*3</sup> Guanghui Xu<sup>1</sup> Yifan Zhang<sup>4</sup>  
Yong Guo<sup>1</sup> Peilin Zhao<sup>3</sup> Peng Wang<sup>5</sup> Mingkui Tan<sup>1,6</sup>

## Abstract

In real-world applications, data often come in a growing manner, where the data volume and the number of classes may increase dynamically. This will bring a critical challenge for learning: given the increasing data volume or the number of classes, one has to instantaneously adjust the neural model capacity to obtain promising performance. Existing methods either ignore the growing nature of data or seek to independently search an optimal architecture for a given dataset, and thus are incapable of promptly adjusting the architectures for the changed data. To address this, we present a neural architecture adaptation method, namely **Adaptation eXpert** (AdaXpert), to efficiently adjust previous architectures on the growing data. Specifically, we introduce an architecture adjuster to generate a suitable architecture for each data snapshot, based on the previous architecture and the different extent between current and previous data distributions. Furthermore, we propose an adaptation condition to determine the necessity of adjustment, thereby avoiding unnecessary and time-consuming adjustments. Extensive experiments on two growth scenarios (increasing data volume and number of classes) demonstrate the effectiveness of the proposed method.

## 1. Introduction

Deep neural networks (DNNs) have achieved state-of-the-art results in many challenging tasks, including image classification (Hu et al., 2018; Lu et al., 2021), neural language pro-

<sup>\*</sup>Equal contribution. This work is done when Shuaicheng Niu works as an intern in Tencent AI Lab. <sup>1</sup>School of Software Engineering, South China University of Technology, China <sup>2</sup>Key Laboratory of Big Data and Intelligent Robot, Ministry of Education, China <sup>3</sup>Tencent AI Lab, China <sup>4</sup>National University of Singapore, Singapore <sup>5</sup>Northwestern Polytechnical University, China <sup>6</sup>Pazhou Laboratory, China. Correspondence to: Mingkui Tan <mingkuitan@scut.edu.cn>.




Figure 1. The motivation for architecture adaptation. **Left:** Performance comparisons of ResNets trained on different subsets of CIFAR100. The optimal architecture varies among different subset sizes. **Right:** Since data come in a growing manner and the data distribution may dynamically change, one should adjust the model architecture according to the shift of data distribution.

cessing (Devlin et al., 2019; Brown et al., 2020), and many other areas (Cao et al., 2019; Zhang et al., 2020b; Guo et al., 2020a; Zeng et al., 2020). One of the key factors behind the success of DNNs lies in the design of effective neural architectures, including 1) the manually designed architectures such as ResNet (He et al., 2016) and MobileNet (Howard et al., 2017); 2) the automatically designed architectures such as (Zoph et al., 2018; Cai et al., 2019; Tan et al., 2019). However, these methods often design a fixed architecture for a specific task/dataset.

In real-world applications, data often come in a growing manner. For example, intelligent edge devices (e.g., billions of mobile phones and surveillance cameras) and medical imaging devices continue collecting new data every day (Grantz et al., 2020; Liang et al., 2019). Specifically, the newly collected data have the following two types: (1) increasing data volume: the labels of new data have appeared in previous data, and the growing do not change the label space of data; (2) increasing number of classes: the newly arriving data have different labels from previous data, and thus the label space of data is growing. For both two scenarios, the data distribution may dynamically change. Since the optimal network architecture may vary under different data distributions (Zoph & Le, 2017), when applying DNNs to growing data, one can (and should) dynamically adjust the architecture for better performance (see Figure 1).

To achieve the above goal, one straightforward solution is to redesign a network architecture when new data arrive. However, the design of effective neural architectures

substantially relies on human expertise. Moreover, human design cannot fully explore the complete architecture space, resulting in sub-optimal architectures (Zoph & Le, 2017). Beyond manual design, one can also resort to automatic neural architecture search (NAS) techniques (Cai et al., 2019; Tan et al., 2019). Nevertheless, such methods design a new architecture for each instance of data growth separately and from scratch, and ignore that previous architectures are transferable, leading to inferior design efficiency. Moreover, neither manual nor automatic design considers the necessity of architecture adjustment to further improve the adaptation efficiency. Intuitively, it is unnecessary to conduct adaptation if newly arrived data are very similar to previous data.

To address the above limitations, we propose a neural architecture adaptation method, called AdaXpert (**Adaptation eXpert**), which consists of an architecture adjuster and an adaptation condition. Specifically, we first adopt the Wasserstein distance to quantitatively measure the difference between current and previous data. Then, the adjuster takes the previous architecture and data difference as inputs, and generates a suitable architecture for the current data. Next, the adjuster receives a reward for this adjustment, where the reward is designed to compromise between the accuracy and computational efficiency. The final adjusted architecture is generated by a well-trained adjuster. It is worth mentioning that our method aims to find a new optimal architecture for current data rather than simply expand the previous architecture to a larger one. The architecture adjustment is elastic, e.g., the adjustment may remove redundant layers from the previous architecture and increase the capacity (kernel size and/or number of channels) of certain layers, as illustrated in Figure 4. Moreover, we propose an adaptation condition to determine the necessity of architecture adjustment. In this way, we avoid unnecessary adjustments if the newly arrived data are highly similar to the previous one, thereby further improving the adjustment efficiency. Based on the above considerations, our AdaXpert is able to adjust an architecture automatically to obtain better performance on the current data with as minimal computational cost as possible. Our main contributions are summarized as follows:

- We propose a network adaptation method for growing data. By considering the difference between the current and previous data, our method adaptively adjusts the model architecture to achieve better performance while maintaining small computational cost.
- We propose an adaptation condition to determine the necessity of architecture adjustment. With this condition, our method avoids unnecessary adaptation for highly similar data.
- Experiments on two data growth scenarios, i.e., increasing data volume and the number of classes, demonstrate the effectiveness and superiority of our method.

## 2. Related Work

**Neural architecture search (NAS)** has attracted increasing attention to automatically design effective architectures. The classical NAS problem (Zoph & Le, 2017) exploits the paradigms of reinforcement learning (RL) to generate the model descriptions of DNNs. RL-based methods (Pham et al., 2018; Tan et al., 2019; Zoph et al., 2018) seek to learn a controller with a policy to generate architectures. Beyond RL, evolutionary-based (Real et al., 2019; Piergiovanni et al., 2019) and gradient-based (Liu et al., 2019; Xu et al., 2020) algorithms also discover new architectures with excellent performance. Recently, meta learning-based methods (Elsken et al., 2020; Wang et al., 2020) focus on the few-shot problem and automatically learn a meta-architecture that is intended to adapt to new tasks quickly. Unlike NAS that design a fixed architecture, we dynamically adjust the network architecture to handle the problem of growing data.

**Continual learning (CL)** aims to transfer the knowledge learned from previous tasks to future scenarios. To solve a new task, replay-based methods (Rebuffi et al., 2017; Chaudhry et al., 2019b; Rolnick et al., 2019) selectively store samples of previous tasks for the training on new tasks. To mitigate the catastrophic forgetting issue, regularization-based methods (Kirkpatrick et al., 2017; Liu et al., 2018; Lange et al., 2020) introduce a regularization term in the loss function, which requires the model to not change important parameters of previous tasks. Another parallel task is online learning (OL) (Hoi et al., 2018; Zhang et al., 2019), which aims to learn a well-performed model based on a sequence of training samples. With the data growth, CL aims to overcome the forgetting issue on previous data and conduct adjustments to achieve better performance on new data, and OL focuses on the parameter learning of a certain model. In this work, we seek to instantaneously adjust the architecture on the entire dataset for each time of data growth.

**Progressive neural networks.** To improve the model capacity, CL methods (Rusu et al., 2016; Chaudhry et al., 2019a; Rosenfeld & Tsotsos, 2020) propose to dynamically expand their network architectures. These methods fix the layers of previous tasks and grow branches for new tasks. Moreover, DEN (Yoon et al., 2018) first expands the architecture to a large size for a new task, and then use a pruning method to remove the unimportant weights. Recently, some works (Zhang et al., 2020a; Li et al., 2019) combine NAS techniques to design architectures for each task to achieve the goal of CL. However, these methods ignore the distribution difference between the current and previous data, and thus are hard to determine a suitable model size of the adjusted architecture. In this work, we dynamically adjust the architecture based on previous architectures and the properties of growing data. Moreover, beyond the expansion, our adjustment may remove redundant layers or add new layers.

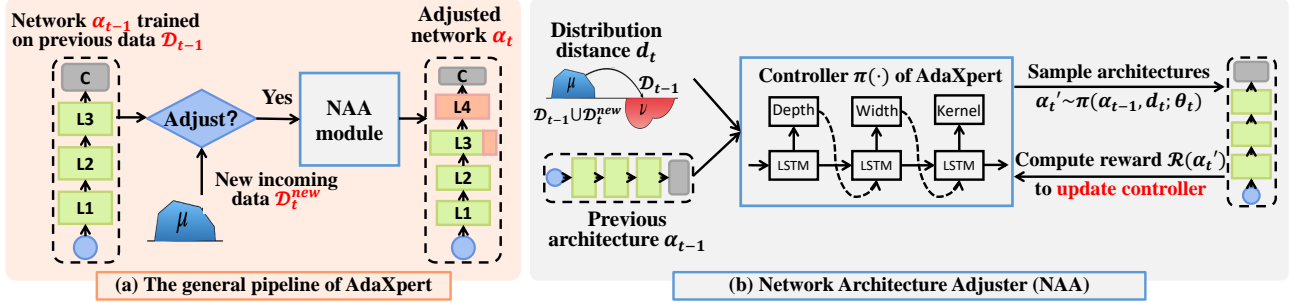


Figure 2. An illustration of our proposed AdaXpert. (a) At time-step  $t$ , given new incoming data  $\mathcal{D}_t^{new}$  and a previous model  $\alpha_{t-1}$ , we first determine whether to adjust the architecture. If necessary, the  $\alpha_{t-1}$  is fed into the NAA module for architecture adaptation. (b) Our controller takes the architecture of  $\alpha_{t-1}$  and the distribution distance between current data  $\mathcal{D}_{t-1} \cup \mathcal{D}_t^{new}$  and previous data  $\mathcal{D}_{t-1}$  as inputs, and outputs an adjusted architecture. The controller will then obtain a reward and thus can be trained via policy gradient methods. Last, we adopt the well-trained NAA to generate the final adjusted architecture  $\alpha_t$ .

### 3. Proposed Method

#### 3.1. Problem Definition

In this paper, we aim to dynamically adjust neural network architectures along with the growth of data. Formally, we denote a series of new incoming data as  $\mathcal{D}_t^{new} = \{(x_i^t, y_i^t)\}_{i=1}^{n_t}$ ,  $t=1, \dots, T$ , where  $x_i^t \in \mathcal{X}$ ,  $y_i^t \in \{1, \dots, C_t\}$ ,  $\mathcal{X}$  is the input image space,  $n_t$  and  $C_t$  are the number of images and classes, respectively. Moreover, we denote all the accumulated data at time-step  $t$  as  $\mathcal{D}_t = \mathcal{D}_1^{new} \cup \dots \cup \mathcal{D}_t^{new}$ .

For the above growing datasets  $\{\mathcal{D}_t\}_{t=1}^T$ , their corresponding data distributions may dynamically change, and thus the optimal architecture for different  $\mathcal{D}_t$  may also vary. However, existing methods usually design a fixed architecture for a specific task/dataset, while a single fixed network architecture may not be optimal for all  $\{\mathcal{D}_t\}_{t=1}^T$ . To achieve better performance, one should design different architectures for different  $\mathcal{D}_t$ , i.e., dynamically adjust the model architecture along with the growth of data.

To achieve the above goal, an intuitive way is to perform neural architecture search for each  $\mathcal{D}_t$  separately, to obtain a corresponding architecture  $\alpha_t$ . However, this method ignores that the previous  $\alpha_{t-1}$  are transferable for  $\mathcal{D}_t$ , which can improve the search efficiency of  $\alpha_t$ . Moreover, it is also important to avoid unnecessary search processes when the new incoming data are very similar to the previous data.

To address the above challenges, we propose a neural architecture adaptation method, namely **Adaptation eXpert** (AdaXpert), which aims to automatically design dynamic networks with minimal expert intervention. The AdaXpert mainly consists of two key components: (1) a reinforcement learning based neural architecture adjuster (see Sect. 3.2), which adaptively adjusts the previous model architecture  $\alpha_{t-1}$  to a new  $\alpha_t$  according to the properties of new incoming data  $\mathcal{D}_t^{new}$ ; (2) an adaptation condition that determines

the necessity of architecture adjustment (see Sect. 3.3), thereby avoiding unnecessary adjustments. The detailed pipeline of AdaXpert is summarized in Algorithm 1.

#### 3.2. Dynamic Neural Architecture Adaptation

Given a previous deep model and new incoming data, we aim to automatically adjust the model architecture to achieve better performance while retaining small model computational cost (e.g., MAdds). To this end, we devise a Network Architecture Adjuster (NAA) algorithm, which aims to conduct different adjustment strategies based on the distribution difference between  $\mathcal{D}_{t-1}$  and  $\mathcal{D}_t$ . Specifically, if new data are very similar to the previous one, we only need to conduct a slight adjustment, i.e., keeping the MAdds growth small for the adjusted architecture. Otherwise, we allow a relatively large MAdds growth for the adjustment.

In the following, we address two key problems for architecture adaptation: 1) how to measure the difference between the current data  $\mathcal{D}_t$  and the previous data  $\mathcal{D}_{t-1}$ ; and 2) how to design the architecture adjuster.

**Quantitative measurement of data difference.** Different architecture adaptation strategies should be conducted based on the different extent between the current data and previous data. To quantify such difference, we compute the distribution distance between current and previous data as follows.

Formally, given the current dataset  $\mathcal{D}_t$  and the previous dataset  $\mathcal{D}_{t-1}$ , we first feed these two datasets to the previous model  $\alpha_{t-1}$  to obtain their feature embeddings  $\mathbf{M}_t \in \mathbb{R}^{m \times q}$  and  $\mathbf{M}_{t-1} \in \mathbb{R}^{n \times q}$ , respectively. Here,  $m$  and  $n$  denote the number of samples in  $\mathcal{D}_t$  and  $\mathcal{D}_{t-1}$  respectively, and  $q$  denotes the feature dimension. Then,  $\mathbf{M}_t$  and  $\mathbf{M}_{t-1}$  can be considered as two sample matrices that are sampled from two unknown distributions  $\mathbb{P}_t$  and  $\mathbb{P}_{t-1}$ . To compute the distance between  $\mathbb{P}_t$  and  $\mathbb{P}_{t-1}$ , one can use non-parametric

**Algorithm 1** The overall algorithm of AdaXpert.

**Input:** Incoming datasets  $\{\mathcal{D}_t^{new}\}_{t=1}^T$ ; well-trained model  $\alpha_1$  for  $\mathcal{D}_1^{new}$ ; supernet  $\mathcal{N}_1$  and controller  $\pi(\cdot; \theta_1)$ ; threshold  $\epsilon$ .

- 1: Let  $\mathcal{D}_1 = \mathcal{D}_1^{new}$ .
- 2: **for**  $t=2, \dots, T$  **do**
- 3:   Let  $\mathcal{D}_t = \mathcal{D}_{t-1} \cup \mathcal{D}_t^{new}$ ,  $\mathcal{N}_t = \mathcal{N}_{t-1}$  and  $\theta_t = \theta_{t-1}$ .
- 4:   Compute accuracy difference  $H_t$  using Eqn. (4).
- 5:   **if**  $H_t > \epsilon$  **then**
- 6:     Update  $\mathcal{N}_t$  and  $\pi(\cdot; \theta_t)$  on  $\mathcal{D}_t$  using Algorithm 2.
- 7:     Generate the adjusted architecture  $\alpha_t \sim \pi(\cdot; \theta_t)$ .
- 8:     Retrain the adjusted architecture  $\alpha_t$  on  $\mathcal{D}_t$ .
- 9:   **else**
- 10:    Let  $\alpha_t = \alpha_{t-1}$ .
- 11:   **end if**
- 12: **end for**

**Output:** The adjusted architectures  $\{\alpha_t\}_{t=1}^T$ .

estimation methods to compute the Kullback–Leibler (KL) divergence (Nguyen et al., 2007) or Wasserstein distance (WD) (Sriperumbudur et al., 2010).

However, the above non-parametric estimation methods may be computationally expensive. For example, computing the KL divergence needs solving a quadratic programming problem. Luckily, our preliminary studies show that the sample matrices  $\mathbf{M}_t$  and  $\mathbf{M}_{t-1}$  approximately satisfy the multivariate Gaussian distribution (more details are put in the supplementary). Therefore, in this paper, we assume that  $\mathbb{P}_t$  and  $\mathbb{P}_{t-1}$  are two multivariate Gaussian distributions, and use the Maximum Likelihood Estimation method to obtain their distribution parameters, i.e.,  $\mathbb{P}_t \sim \mathcal{N}(\mu_t, \Sigma_t)$  and  $\mathbb{P}_{t-1} \sim \mathcal{N}(\mu_{t-1}, \Sigma_{t-1})$ . Then, we compute the Wasserstein distance (Takatsu et al., 2011) as follows:

$$\mathcal{W}(\mathcal{D}_t, \mathcal{D}_{t-1}) = \|\mu_t - \mu_{t-1}\|_2^2 + \text{tr}\left(\Sigma_t + \Sigma_{t-1} - 2(\Sigma_{t-1}^{1/2} \Sigma_t \Sigma_{t-1}^{1/2})^{1/2}\right). \quad (1)$$

Here, one can also use other metrics to compute the distribution distance, such as KL and Jensen-Shannon divergence (Fuglede & Topsøe, 2004). More discussions about WD are put in supplementary. Based on  $\mathcal{W}(\mathcal{D}_t, \mathcal{D}_{t-1})$ , we devise a data difference-aware controller to conduct different architecture adjustments. Due to the highly non-convex nature of our adjustment problem, we cast it into a Markov Decision Process (MDP), and then train the controller using reinforcement learning methods.

**MDP reformulation for Neural Architecture Adjuster (NAA).** Since the architecture adjustment process is essentially a multi-step decision making process, we formalize the adjustment process as an MDP. Formally, the MDP can be defined as a tuple  $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R})$ , where  $\mathcal{S}$  is a finite set of states,  $\mathcal{A}$  is a finite set of actions,  $\mathcal{P} : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$  is the state transition distribution,  $\mathcal{R} : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{R}$  is the reward function. Moreover, a policy  $\pi_\theta$  determines an action given

**Algorithm 2** Training of Network Architecture Adjuster.

**Input:** Datasets  $\{\mathcal{D}_{t-1}, \mathcal{D}_t\}$ ; previous architecture  $\alpha_{t-1}$ , supernet  $\mathcal{N}_t$  and controller  $\pi(\cdot)$  with the parameters  $\theta_t$ ; hyperparameters  $\eta$  and  $M$ .

- 1: Split  $\mathcal{D}_t$  into training and validation sets  $\{\mathcal{D}_{train}, \mathcal{D}_{val}\}$ .
- 2: Fine-tune  $\mathcal{N}_t$  on  $\mathcal{D}_{train}$ .
- 3: Compute the WD ( $d_t$ ) between  $\mathcal{D}_{t-1}$  and  $\mathcal{D}_t$  using Eqn. (1);
- 4: // train the controller model
- 5: **for**  $i=1, \dots, M$  **do**
- 6:   Sample  $\alpha'_t \sim \pi(\alpha_{t-1}, d_t; \theta_t)$ .
- 7:   Sample a batch of data from  $\mathcal{D}_{val}$ .
- 8:   Compute reward  $\mathcal{R}(\alpha'_t)$  based on  $\mathcal{N}_t$  using Eqn. (5).
- 9:   Update  $\theta_t \leftarrow \theta_t + \eta \mathcal{R}(\alpha'_t) \nabla_{\theta_t} \log \pi(\cdot)$ .
- 10: **end for**

**Output:** Supernet  $\mathcal{N}_t$  and controller  $\pi(\cdot; \theta_t)$ .

the current state. In the context of NAA, as illustrated in Figure 2 for time-step  $t$ , we denote state as  $s = [\alpha_{t-1}, d_t] \in \mathcal{S}$ , where  $d_t = \mathcal{W}(\mathcal{D}_t, \mathcal{D}_{t-1})$  is the distribution difference between the current and previous data. Given such a state  $s$ , a policy (controller) takes a series of actions  $a = \pi_\theta(s) \in \mathcal{A}$  to determine each layer’s operation of the adjusted architecture  $\alpha'_t$ . Formally, the action space is defined based on different types of architecture search space. Then, the controller receives a reward  $r = \mathcal{R}(\alpha'_t)$ . More details about the reward design can be found in Sec. 3.4.

**Training of NAA.** The goal of our NAA is to maximize an expectation reward  $\mathbb{E}[\mathcal{R}(\alpha)]$ , represented by solving the following optimization problem:

$$\max_{\theta} \mathbb{E}_{\pi_\theta} [\mathcal{R}(\alpha)]. \quad (2)$$

Following policy gradient methods (Williams, 1992; Schulman et al., 2017), we update  $\theta$  by ascending the gradient:

$$\theta \leftarrow \theta + \eta \mathcal{R}(\alpha) \nabla_{\theta} \log \pi_{\theta}(\alpha). \quad (3)$$

The training details of NAA are summarized in Algorithm 2.

### 3.3. When to Adapt Network Architecture

Given a previous model  $\alpha_{t-1}$  and the current data  $\mathcal{D}_t$ , it should be considered whether the model architecture needs adjustment. For example, given an optimal architecture  $\alpha$  that searched on the MNIST dataset, even if we receive more incoming MNIST images, it is unnecessary to adjust  $\alpha$  since it is already optimal. This indicates the necessity to consider the previous architecture’s feasibility when performing architecture adaptation. Therefore, as shown in Figure 2 (a), for each time new data arrives, we use an adaptation condition to determine the necessity of adjustment, and therefore improve the adaptation efficiency.

Formally, at time-step  $t$ , given the previous model  $\alpha_{t-1}$ , current data  $\mathcal{D}_t$  and previous data  $\mathcal{D}_{t-1}$ , we compute the



following accuracy difference for further decision:

$$H_t = \Phi(\mathcal{D}_{t-1}; \alpha_{t-1}) - \Phi(\mathcal{D}_t; \alpha_{t-1}), \quad (4)$$

where  $\Phi(\mathcal{D}; \alpha)$  is some performance metric of model  $\alpha$  on dataset  $\mathcal{D}$ . For classification models, we choose top-1 accuracy as the metric. Note that for the growing data with new class labels, the previous model will make wrong predictions, since the classifier can not predict the new labels. Based on the above accuracy difference, one can determine whether to adjust the previous architecture  $\alpha_{t-1}$ . Specifically, given  $H_t$  and its threshold  $\epsilon$ , we only adjust the model architecture when  $H_t > \epsilon$ .

In our method, we exploit WD to measure the extent of data difference, and adopts accuracy difference as the adaptation condition of architecture adjustment. The reasons are as follows: (1) For adaptation condition, the accuracy difference is more intuitive for humans. In contrast, since WD is a distribution distance, it is hard to set a suitable threshold for WD to determine the adjustment necessity. (2) To recognize difference extent of data, WD has a stronger discrimination ability than the accuracy difference. Specifically, for varying label space, the accuracy difference is more determined by the number of new data, while WD is computed according to the underlying properties of the data itself.

### 3.4. Reward Design for NAA

The reward function  $\mathcal{R}(\cdot)$  is very important for training our NAA model. In this subsection, we will provide our rewards function. For simplicity, we only illustrate the training of NAA for a single round, and in other rounds, the NAA can be trained in the same way. Given the current and previous datasets  $\mathcal{D}_t$  and  $\mathcal{D}_{t-1}$ , we denote the WD between them as  $d_t = \mathcal{W}(\mathcal{D}_t, \mathcal{D}_{t-1})$ . Then, for the current state  $s$ , the NAA takes a series of actions to obtain an adjusted architecture  $\alpha'_t$ . The final reward is computed as follows:

$$\mathcal{R}(\alpha'_t) = \mathcal{V}(\alpha'_t) - \mathcal{V}(\alpha_{t-1}) - \frac{\lambda}{d_t} (\mathcal{C}(\alpha'_t) - \mathcal{C}(\alpha_{t-1})), \quad (5)$$

where  $\mathcal{V}(\alpha)$  and  $\mathcal{C}(\alpha)$  denote the validation accuracy and computational complexity of model  $\alpha$  respectively, and  $\lambda$  is a trade-off parameter. We adopt MAdds as our metric to measure the computational complexity of  $\alpha$ , and one can also use other metrics, e.g., the inference latency.

To obtain the validation accuracy  $\mathcal{V}(\alpha)$ , one can train it from scratch and then validate it on the validation set. However, this will result in unbearable computational burdens. In this paper, we exploit a weight sharing technique (Pham et al., 2018) to construct a super network, i.e., a large computational graph, where each network architecture shares parameters. In this sense, once the super network is trained, all architectures inherit their weights directly from the super network, and then use these weights for further evaluation.

For the first item  $\mathcal{V}(\alpha'_t) - \mathcal{V}(\alpha_{t-1})$  in Eqn. (5), we hope that the adjusted architecture  $\alpha'_t$  would achieve higher validation accuracy than the original  $\alpha_{t-1}$ . For the second item  $\frac{\lambda}{d_t} \times (\mathcal{C}(\alpha'_t) - \mathcal{C}(\alpha_{t-1}))$  in Eqn. (5), architectures with excessive computational cost are penalized, and the  $d_t$  is used to adaptively regularize the magnitude of the adjusted architecture. If  $d_t$  is small, i.e., the incoming data are very similar to the original, the reward function will devote more attention to constraining the computational complexity of the adjusted architecture.

## 4. Experiments

In this section, we evaluate our AdaXpert with respect to two data growth scenarios, i.e., data volume growth within the same label space (**Scenario I**) and increasing label space (**Scenario II**). Afterward, we conduct ablation studies to verify the effectiveness of each component in our method. Lastly, we compare architectures obtained by our adaptation procedure with those obtained by existing methods. Code is available at <https://github.com/mr-eggplant/adaxpert0>.

### 4.1. Experimental Settings

**Datasets:** We conduct our experiments on ImageNet, a large-scale image classification dataset (Deng et al., 2009). Based on ImageNet, we simulate two data growth scenarios to verify the effectiveness of our proposed method. For convenience, we denote ImageNet-# as a subset of ImageNet, where ‘#’ denotes the number of classes. For instance, ImageNet-100 contains samples of the first 100 classes of the entire ImageNet. We also name our dynamically adjusted architectures in a similar manner, e.g., AdaXpert-20 denotes the architecture obtained on ImageNet-20.

**Search space for architecture adaptation:** Here, we consider the architecture space based on the inverted Mobile Block (Howard et al., 2019). To be specific, the model is divided into 5 units with gradually reduced feature map spatial size and increased number of channels. Each unit consists of 4 layers at most, where only the first layer has stride 2 if the feature map size decreases, and all the other layers have stride 1. In our experiments, we search for the number of layers in each unit (chosen from  $\{2, 3, 4\}$ ), the kernel size in each layer (chosen from  $\{3, 5, 7\}$ ), and the width expansion ratio in each layer (chosen from  $\{3, 4, 6\}$ ).

**Compared methods:** We compare our AdaXpert with three categories of methods. (1) Manually designed networks, including MobileNetV2, MobileNetV2 (1.4 $\times$ ) (Howard et al., 2017), ResNet18, and ResNet50 (He et al., 2016). During the entire data growth process, these models are trained and evaluated with the same fixed architecture. (2) Neural architecture search (NAS) methods. EfficientNet (Tan & Le, 2019) and MnasNet (Tan et al., 2019) are searched on

Table 1. **Scenario I:** Comparison on ImageNet-100 with different sizes of training set. We report Acc. (% ,  $\uparrow$ ) and #MAdds (Million,  $\downarrow$ ).

Methods	10% training set		20% training set		40% training set		80% training set		100% training set	
	Acc.	MAdds	Acc.	MAdds	Acc.	MAdds	Acc.	MAdds	Acc.	MAdds
MobileNetV2	52.72	300	64.02	300	72.08	300	78.36	300	79.60	300
MobileNetV2 (1.4 $\times$ )	54.36	560	64.82	560	73.02	560	78.88	560	80.76	560
ResNet18	52.72	1,814	62.74	1,814	71.54	1,814	77.54	1,814	79.30	1,814
ResNet50	38.86	4,087	53.26	4,087	66.76	4,087	78.62	4,087	80.30	4,087
MnasNet-A1	50.30	323	61.14	323	71.46	323	78.82	323	79.66	323
EfficientNet-B0	50.56	398	62.90	398	72.32	398	<b>79.42</b>	398	80.38	398
Meta-NAS	51.00	559	60.00	559	69.30	559	77.08	559	77.48	559
D-EfficientNets	49.08	<b>145</b>	61.42	178	71.28	203	78.26	229	80.44	278
Progressive NN	54.60	149	61.72	181	71.38	203	78.78	244	80.02	261
DEN	54.60	149	63.50	258	72.20	315	78.85	439	<b>80.84</b>	515
AdaXpert (ours)	<b>54.60</b>	149	<b>64.90</b>	<b>171</b>	<b>73.28</b>	<b>199</b>	79.28	<b>232</b>	80.74	<b>252</b>

inverted Mobile Block search space (as ours) and achieve state-of-the-art performance. We also compare our method with Meta-NAS (Shaw et al., 2019), which first searches a meta architecture on multiple tasks and then adapts it to ImageNet. (3) Dynamic neural networks. Progressive NN (Rusu et al., 2016) and DEN (Yoon et al., 2018) are two methods that dynamically adjust the network architecture from small to large with the growth of data. D-EfficientNets is a width-multiplier method, where the networks are re-scaled with different widths of EfficientNet-B0 (Tan & Le, 2019) to adapt the corresponding data. Please refer to the supplementary for further implementation details.

#### 4.2. I: Growing Data with Same Label Space

In this section, we conduct experiments on the data growth scenario in which the data volume is growing while the label space remains the same.

**Simulation of growing data:** We simulate the data volume growth scenario on ImageNet-100 due to the high computational cost of evaluating all considered architectures on ImageNet-1000. To be specific, the data come with different ratios, i.e.,  $\{10\%, 20\%, 40\%, 80\%, 100\%\}$ , and the number of classes remains unchanged for each data growth scenario. Here, a dataset with a small ratio is a subset of another dataset with a larger ratio.

**Comparison with state-of-the-art methods.** As shown in Table 1, our method achieves the best or comparable accuracy in all cases, suggesting its effectiveness. More critically, the computational cost (i.e., MAdds) of our model is significantly lower than that of other methods, verifying that the proposed reward function enables the consideration of model efficiency. Specifically, our model achieves comparable accuracy with DEN (80.74 vs. 80.84) while the computational cost is much lower (252 M vs. 515 M). Similar phenomena are widely observed in Table 1, indicating that our method achieves a better accuracy/computational efficiency trade-off than other state-of-the-art approaches.

#### 4.3. II: Growing Data with Increasing Label Space

In this section, we conduct experiments on a data growth scenario in which the label space is growing, i.e., the new data have more classes than previous data. This scenario is more challenging since the distribution of new data may be remarkably different (i.e., new classes) from that of previous data. The experimental datasets are constructed as follows.

**Simulation of growing data:** In this experiment, we use the entire ImageNet-1000 to construct our growing datasets. Similar to Scenario I, the dataset grows five times and each subset contains the first  $\{10, 20, 40, 80, 100, 200, 1000\}$  classes of the entire ImageNet, respectively. Similarly, the data of the latter case contain all the data of the former case.

**Comparison with state-of-the-art methods:** As shown in Table 2, our method is able to achieve comparable performance while requiring much lower computational cost. Specifically, for ImageNet-1000, our method outperforms all the baseline methods in terms of accuracy. For ImageNet-100, our method outperforms manual-designed networks such as ResNet50, while requiring  $\times 15.7$  fewer MAdds (257M vs. 4087M). Notably, when the number of classes is small, ResNet18 is significantly better than ResNet50. However, the situation turns around as the number of classes increases, which verifies our motivation that the optimal architecture may vary under different data distributions.

#### 4.4. Ablation Studies

##### Effectiveness of the adaptation condition in Eqn. (4).

We conduct experiments to further demonstrate the effectiveness of our adaptation condition for architecture adjustment. Specifically, we first prepare a base dataset  $\mathcal{D}_b$  (20% training set of ImageNet-100 on Scenario I) and a model trained on  $\mathcal{D}_b$ . To demonstrate the necessity of architecture adaptation, we consider two different datasets, i.e.,  $\mathcal{D}_s$  (with a small difference from  $\mathcal{D}_b$ ) and  $\mathcal{D}_l$  (with a large difference from  $\mathcal{D}_b$ ). To construct  $\mathcal{D}_s$ , we apply data augmentation techniques over

Table 2. **Scenario II:** Comparison on ImageNet-1000 with different number of classes. We report Acc. (%, $\uparrow$ ) and #MAdds (Million,  $\downarrow$ ).

Methods	10 classes		20 classes		40 classes		80 classes		100 classes		200 classes		1000 classes	
	Acc.	MAdds	Acc.	MAdds	Acc.	MAdds	Acc.	MAdds	Acc.	MAdds	Acc.	MAdds	Acc.	MAdds
MobileNetV2	81.80	300	85.10	300	81.10	300	76.92	300	79.60	300	80.83	300	72.00	<b>300</b>
MobileNetV2 (1.4x)	81.00	560	85.70	560	81.30	560	<b>77.85</b>	560	80.76	560	81.90	560	74.70	560
ResNet18	<b>82.80</b>	1,814	85.90	1,814	81.90	1,814	75.60	1,814	79.30	1,814	79.89	1,814	72.12	1,814
ResNet50	69.60	4,087	81.80	4,087	78.85	4,087	76.52	4,087	80.30	4,087	<b>82.89</b>	4,087	77.15	4,087
MnasNet-A1	80.60	323	84.60	323	80.80	323	77.03	323	79.66	323	81.95	323	75.20	323
EfficientNet-B0	81.40	398	86.00	398	<b>82.10</b>	398	77.70	398	80.38	398	82.49	398	76.30	398
Meta-NAS	81.20	559	85.50	559	80.75	559	75.03	559	77.48	559	80.53	559	74.30	559
D-EfficientNets	81.00	<b>145</b>	84.20	178	80.40	203	76.63	229	80.44	278	82.03	319	76.30	398
Progressive NN	81.20	149	86.10	181	81.10	203	76.43	244	80.02	261	82.20	329	77.53	427
DEN	81.20	149	86.10	258	81.35	315	77.60	439	<b>80.84</b>	515	82.09	549	72.99	672
AdaXpert (ours)	81.20	149	<b>86.40</b>	<b>176</b>	81.90	<b>195</b>	77.68	<b>242</b>	80.52	<b>257</b>	82.12	<b>293</b>	<b>78.13</b>	395

Table 3. Ablation studies on the adaptation condition. We report the accuracy of No-adjusted and Adjusted models on new current datasets  $\mathcal{D}_b \cup \mathcal{D}_s$  and  $\mathcal{D}_b \cup \mathcal{D}_l$ , respectively.

Dataset	$H_t$ (Eqn. 4)	No-adjusted (Acc. %)	Adjusted (Acc. %)
$\mathcal{D}_b \cup \mathcal{D}_s$	0.65	64.90	65.04 (+0.14)
$\mathcal{D}_b \cup \mathcal{D}_l$	7.76	72.54	73.58 (+1.04)

Table 4. Comparison with nas-for-each (NFE) on ImageNet-100. NFE means “search from scratch for each time of data growth”.

Metric	Method	20% data	40% data	80% data	100% data
Acc. (%)	NFE	64.80	73.46	78.88	80.62
	AdaXpert (ours)	<b>64.90</b>	<b>73.28</b>	<b>79.28</b>	<b>80.74</b>
MAdds (M)	NFE	294	302	313	311
	AdaXpert (ours)	<b>171</b>	<b>199</b>	<b>232</b>	<b>252</b>
Search Cost (GPU days)	NFE	0.8	1.0	1.3	1.5
	AdaXpert (ours)	<b>0.8</b>	<b>0.6</b>	<b>0.6</b>	<b>0.7</b>

the base dataset  $\mathcal{D}_b$ . To construct  $\mathcal{D}_l$ , we use samples from another 20% training set of ImageNet-100.

As shown in Table 3, we report the accuracy differences (based on Eqn. 4), adjusted/no-adjusted model accuracy on new current data  $\mathcal{D}_b \cup \mathcal{D}_s$  and  $\mathcal{D}_b \cup \mathcal{D}_l$ , respectively. From the results, for similar new data, it is unnecessary to adjust the previous model architecture since the improvement is limited (i.e., Acc: 64.90 vs 65.04). In contrast, adjusting the previous architecture for new different data is able to gain a larger performance improvement (i.e., Acc: 72.54 vs 73.58). In this sense, it is important to use an adaptation condition to determine whether adjustments are needed for new data.

**Comparison with search from scratch.** To further verify the superiority of AdaXpert, we also compare it with “search from scratch for each time data growth (namely NAS-for-Each, NFE)”. From Table 4, our AdaXpert achieves better efficiency. At the 100% data snapshot, the architecture performance obtained by AdaXpert is superior to NFE, this mainly benefits from the following two aspects: 1) The AdaXpert

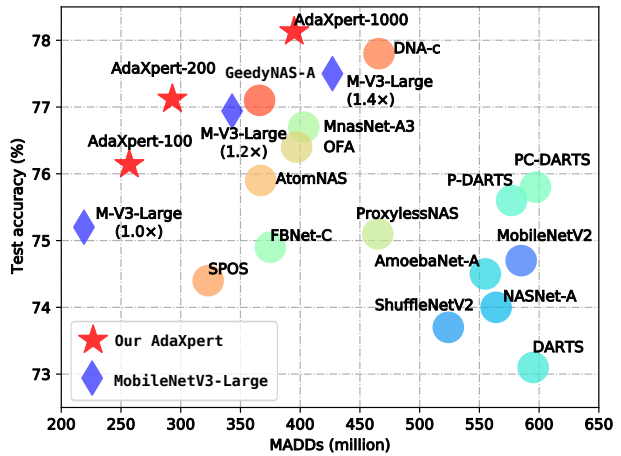


Figure 3. Comparison between AdaXpert and state-of-the-art NAS methods on ImageNet. ‘AdaXpert-#’ denotes our architecture searched on ImageNet-#.

exploits the previously learned knowledge to conduct the current learning. Similar ideas in GAN (e.g., Progressive GAN (Karras et al., 2018)) and NAS (e.g., PNAS (Karras et al., 2018) and CNAS (Guo et al., 2020b)) have proven to be pretty effective. 2) The AdaXpert considers the different extent between current and previous data, and thus to adaptively control the computational cost of adjusted models.

#### 4.5. Comparison on ImageNet-1000

Our proposed method can also be regarded as a neural architecture search (NAS) method, which progressively searches for the optimal architecture on growing datasets. In this section, we compare our intermediate architectures, i.e., AdaXpert-100, AdaXpert-200 and AdaXpert-1000, with existing NAS methods to further verify the effectiveness of our method. Here, AdaXpert-# is searched on ImageNet-#. We re-train each AdaXpert-# model on the whole ImageNet-1000, as well as baseline methods.

Table 5. Comparison of different architectures on ImageNet-1000. Our AdaXpert-# architectures are searched on different subsets of ImageNet and then evaluated on the entire ImageNet dataset. ”-” means unavailable results.

Architecture	Test Accuracy (%)		#MAdds (M)	Search Time (GPU days)	Search Method	Search Space	
	Top-1	Top-5					
ResNet-18 (He et al., 2016)	69.8	89.1	1,814	-			
MobileNetV2 (1.4×) (Sandler et al., 2018)	74.7	-	585	-	manual design	-	
ShuffleNetV2 (2×) (Ma et al., 2018)	73.7	-	524	-			
NASNet-A (Zoph et al., 2018)	74.0	91.6	564	1,800	RL-based evolution	NASNet	
AmoebaNet-A (Real et al., 2019)	74.5	92.0	555	3,150			
DARTS (Liu et al., 2019)	73.1	91.0	595	4	gradient-based	DARTS	
P-DARTS (Chen et al., 2019)	75.6	92.6	577	0.3	gradient-based		
PC-DARTS (Xu et al., 2020)	75.8	92.7	597	3.8	gradient-based		
MobileNetV3-Large (Howard et al., 2019)	75.2	-	219	-	RL-based	Mobile Block	
FBNet-C (Wu et al., 2019)	74.9	-	375	9	gradient-based		
MnasNet-A3 (Tan et al., 2019)	76.7	93.3	403	~3,791	RL-based		
ProxylessNAS (Cai et al., 2019)	75.1	92.3	465	8.3	gradient-based		
SPOS (Guo et al., 2020c)	74.4	91.8	323	12	evolution		
OFA-GPU (Cai et al., 2020)	76.4	-	397	51.7	evolution		
OFA-CPU (Cai et al., 2020)	78.7	-	356	51.7	evolution		
AtomNAS (Mei et al., 2020)	75.9	92.0	367	-	gradient-based		
DNA-c (Li et al., 2020)	77.8	93.7	466	25	greedy search		
GreedyNAS-A (You et al., 2020)	77.1	93.3	366	8	greedy search		
AdaXpert-100 (ours)	76.1	92.7	257	2.5			
AdaXpert-200 (ours)	77.1	93.3	293	3.5	RL-based		Mobile Block
AdaXpert-1000 (ours)	78.1	93.7	395	7			

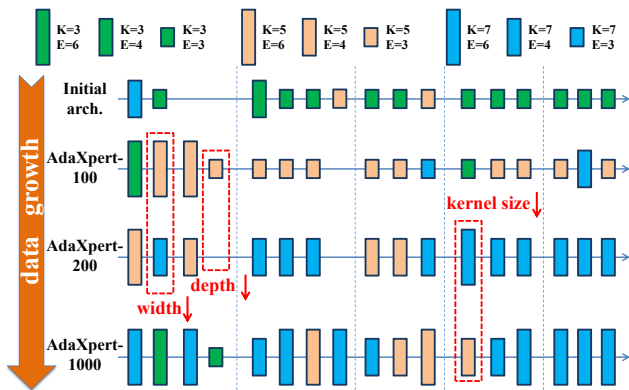


Figure 4. An illustration of the adjusted architectures of our AdaXpert. K and E denote kernel size and expansion ratio, respectively.

As shown in Table 5 and Figure 3, our AdaXpert-1000 achieves 78.1% in terms of top-1 accuracy, which outperforms existing human-designed architectures and mostly considered state-of-the-art NAS models with different search spaces. Surprisingly, our intermediate models, AdaXpert-100 and AdaXpert-200, also achieve comparable performance with most of the baseline methods in top-1 accuracy, but with fewer MAdds. One of the possible reasons is that we exploit previous architectures and measure data distributions’ difference of growing data, while general NAS methods search architectures from scratch. We also provide

the visualization of our adjusted architectures in Figure 4. According to the data difference of each data growth, our AdaXpert adopts different architecture adjustment strategies. Although the model capacity generally increases as the data grow, it is observed that some redundant layers may be removed and some other layers’ kernel size and expansion ratio may be reduced.

## 5. Conclusion

In this paper, we have proposed a new neural architecture adaptation method to efficiently adapt suitable neural architectures for growing data. Unlike existing methods that neglect the knowledge from previous architectures, our method exploits the previous architecture and the data difference extent between current and previous data for achieving effective adaptation. Moreover, we have devised an adaptation condition to avoid unnecessary adjustments, thus further improving the network adaptation efficiency. Experimental results show that our method achieves state-of-the-art performance while enjoying less computational cost in two data growth scenarios (increasing data volume or number of classes). More critically, compared with the architectures searched on the entire ImageNet dataset by existing NAS methods, our architectures are able to achieve comparable accuracy/computational cost with fewer training data (i.e., the subset of ImageNet). In future work, it would be interesting to extend our method to adapt neural architectures for the growing data from diverse data domains.



**Acknowledgements.** This work was partially supported by the National Key R&D Program of China (No. 2020AAA0106900), National Natural Science Foundation of China (NSFC) 62072190, Key-Area Research and Development Program of Guangdong Province (2018B010107001), Program for Guangdong Introducing Innovative and Entrepreneurial Teams 2017ZT07X183, Fundamental Research Funds for the Central Universities D2191240, Tencent AI Lab Rhino-Bird Focused Research Program (No. JR201902).

## References

- Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. Language models are few-shot learners. In *Advances in Neural Information Processing Systems*, 2020.
- Cai, H., Zhu, L., and Han, S. Proxylessnas: Direct neural architecture search on target task and hardware. In *International Conference on Learning Representations*, 2019.
- Cai, H., Gan, C., Wang, T., Zhang, Z., and Han, S. Once for all: Train one network and specialize it for efficient deployment. In *International Conference on Learning Representations*, 2020.
- Cao, J., Mo, L., Zhang, Y., et al. Multi-marginal wasserstein gan. In *Advances in Neural Information Processing Systems*, pp. 1774–1784, 2019.
- Chaudhry, A., Ranzato, M., Rohrbach, M., and Elhoseiny, M. Efficient lifelong learning with a-gem. In *International Conference on Learning Representations*, 2019a.
- Chaudhry, A., Rohrbach, M., Elhoseiny, M., Ajanthan, T., Dokania, P., Torr, P., and Ranzato, M. Continual learning with tiny episodic memories. *ArXiv*, abs/1902.10486, 2019b.
- Chen, X., Xie, L., Wu, J., and Tian, Q. Progressive differentiable architecture search: Bridging the depth gap between search and evaluation. In *IEEE International Conference on Computer Vision*, pp. 1294–1303, 2019.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. Imagenet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 248–255, 2009.
- Devlin, J., Chang, M., Lee, K., and Toutanova, K. BERT: pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics*, pp. 4171–4186, 2019.
- Elsken, T., Staffler, B., Metzen, J. H., and Hutter, F. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 12362–12372, 2020.
- Fuglede, B. and Topsoe, F. Jensen-shannon divergence and hilbert space embedding. In *Proceedings of the IEEE International Symposium on Information Theory*, 2004.
- Grantz, K., Meredith, H. R., Cummings, D., Metcalf, C., Grenfell, B., Giles, J., Mehta, S., Solomon, S., Labrique, A., Kishore, N., Buckee, C., and Wesolowski, A. The use of mobile phone data to inform analysis of covid-19 pandemic epidemiology. *Nature Communications*, 11, 2020.
- Guo, Y., Chen, J., Wang, J., Chen, Q., Cao, J., Deng, Z., Xu, Y., and Tan, M. Closed-loop matters: Dual regression networks for single image super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5407–5416, 2020a.
- Guo, Y., Chen, Y., Zheng, Y., Zhao, P., Chen, J., Huang, J., and Tan, M. Breaking the curse of space explosion: Towards efficient nas with curriculum search. In *Proceedings of the International Conference on Machine Learning*, 2020b.
- Guo, Z., Zhang, X., Mu, H., Heng, W., Liu, Z., Wei, Y., and Sun, J. Single path one-shot neural architecture search with uniform sampling. In *International Conference on Learning Representations*, 2020c.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770–778, 2016.
- Hoi, S., Sahoo, D., Lu, J., and Zhao, P. Online learning: A comprehensive survey. *ArXiv*, abs/1802.02871, 2018.
- Howard, A., Sandler, M., Chu, G., Chen, L.-C., Chen, B., Tan, M., Wang, W., Zhu, Y., Pang, R., Vasudevan, V., et al. Searching for mobilenetv3. In *IEEE International Conference on Computer Vision*, pp. 1314–1324, 2019.
- Howard, A. G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., and Adam, H. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017.
- Hu, J., Shen, L., and Sun, G. Squeeze-and-excitation networks. In *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 7132–7141, 2018.
- Karras, T., Aila, T., Laine, S., and Lehtinen, J. Progressive growing of gans for improved quality, stability, and variation. In *International Conference on Learning Representations*, 2018.

- Kirkpatrick, J., Pascanu, R., Rabinowitz, N. C., Veness, J., Desjardins, G., Rusu, A. A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., Hassabis, D., Clopath, C., Kumaran, D., and Hadsell, R. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114:3521–3526, 2017.
- Lange, M. D., Jia, X., Parisot, S., Leonardis, A., Slabaugh, G., and Tuytelaars, T. Unsupervised model personalization while preserving privacy and scalability: An open problem. In *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 14451–14460, 2020.
- Langley, P. Crafting papers on machine learning. In *Proceedings of the International Conference on Machine Learning*, pp. 1207–1216, 2000.
- Li, C., Peng, J., Yuan, L., Wang, G., Liang, X., Lin, L., and Chang, X. Block-wisely supervised neural architecture search with knowledge distillation. In *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1989–1998, 2020.
- Li, X., Zhou, Y., Wu, T., Socher, R., and Xiong, C. Learn to grow: A continual structure learning framework for overcoming catastrophic forgetting. In *Proceedings of the International Conference on Machine Learning*, pp. 3925–3934, 2019.
- Liang, H., Tsui, B., Ni, H., et al. Evaluation and accurate diagnoses of pediatric diseases using artificial intelligence. *Nature Medicine*, 25:433–438, 2019.
- Liu, H., Simonyan, K., and Yang, Y. Darts: Differentiable architecture search. In *International Conference on Learning Representations*, 2019.
- Liu, X., Masana, M., Herranz, L., van de Weijer, J., López, A. M., and Bagdanov, A. D. Rotate your networks: Better weight consolidation and less catastrophic forgetting. In *International Conference on Pattern Recognition*, pp. 2262–2268, 2018.
- Lu, Z., Sreekumar, G., Goodman, E., Banzhaf, W., Deb, K., and Boddeti, V. Neural architecture transfer. *IEEE transactions on pattern analysis and machine intelligence*, PP, 2021.
- Ma, N., Zhang, X., Zheng, H.-T., and Sun, J. ShuffleNet V2: Practical guidelines for efficient CNN architecture design. In *European Conference on Computer Vision*, pp. 116–131, 2018.
- Mei, J., Li, Y., Lian, X., Jin, X., Yang, L., Yuille, A., and Yang, J. Atomnas: Fine-grained end-to-end neural architecture search. In *International Conference on Learning Representations*, 2020.
- Nguyen, X., Wainwright, M. J., and Jordan, M. I. Nonparametric estimation of the likelihood ratio and divergence functionals. In *IEEE International Symposium on Information Theory*, pp. 2016–2020. IEEE, 2007.
- Pham, H., Guan, M. Y., Zoph, B., Le, Q. V., and Dean, J. Efficient neural architecture search via parameter sharing. In *Proceedings of the International Conference on Machine Learning*, pp. 4092–4101, 2018.
- Piergiovanni, A. J., Angelova, A., and Ryoo, M. S. Tiny video networks. *arXiv preprint arXiv:1910.06961*, 2019.
- Real, E., Aggarwal, A., Huang, Y., and Le, Q. V. Regularized evolution for image classifier architecture search. In *AAAI Conference on Artificial Intelligence*, 2019.
- Rebuffi, S.-A., Kolesnikov, A., Sperl, G., and Lampert, C. H. icarl: Incremental classifier and representation learning. In *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5533–5542, 2017.
- Rolnick, D., Ahuja, A., Schwarz, J., Lillicrap, T., and Wayne, G. Experience replay for continual learning. In *Advances in Neural Information Processing Systems*, 2019.
- Rosenfeld, A. and Tsotsos, J. K. Incremental learning through deep adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42:651–663, 2020.
- Rusu, A. A., Rabinowitz, N. C., Desjardins, G., Soyer, H., Kirkpatrick, J., Kavukcuoglu, K., Pascanu, R., and Hadsell, R. Progressive neural networks. *ArXiv*, abs/1606.04671, 2016.
- Sandler, M., Howard, A. G., Zhu, M., Zhmoginov, A., and Chen, L. MobileNetV2: Inverted residuals and linear bottlenecks. In *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4510–4520, 2018.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Shaw, A., Wei, W., Liu, W., Song, L., and Dai, B. Meta architecture search. In *Advances in Neural Information Processing Systems*, pp. 11225–11235, 2019.
- Sriperumbudur, B. K., Fukumizu, K., Gretton, A., Schölkopf, B., and Lanckriet, G. R. Non-parametric estimation of integral probability metrics. In *IEEE International Symposium on Information Theory*, pp. 1428–1432. IEEE, 2010.
- Takatsu, A. et al. Wasserstein geometry of gaussian measures. *Osaka Journal of Mathematics*, 48(4):1005–1026, 2011.

- Tan, M. and Le, Q. Efficientnet: Rethinking model scaling for convolutional neural networks. In *Proceedings of the International Conference on Machine Learning*, pp. 6105–6114, 2019.
- Tan, M., Chen, B., Pang, R., Vasudevan, V., Sandler, M., Howard, A., and Le, Q. V. Mnasnet: Platform-aware neural architecture search for mobile. In *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2820–2828, 2019.
- Wang, J., Wu, J., Bai, H., and Cheng, J. M-nas: Meta neural architecture search. In *AAAI Conference on Artificial Intelligence*, 2020.
- Williams, R. J. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256, 1992.
- Wu, B., Dai, X., Zhang, P., Wang, Y., Sun, F., Wu, Y., Tian, Y., Vajda, P., Jia, Y., and Keutzer, K. FBNet: Hardware-aware efficient convnet design via differentiable neural architecture search. In *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 10734–10742, 2019.
- Xu, Y., Xie, L., Zhang, X., Chen, X., Qi, G.-J., Tian, Q., and Xiong, H. PC-DARTS: Partial channel connections for memory-efficient architecture search. In *International Conference on Learning Representations*, 2020.
- Yoon, J., Yang, E., Lee, J., and Hwang, S. J. Lifelong learning with dynamically expandable networks. In *International Conference on Learning Representations*, 2018.
- You, S., Huang, T., Yang, M., Wang, F., Qian, C., and Zhang, C. Greedynas: Towards fast one-shot nas with greedy supernet. In *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1999–2008, 2020.
- Zeng, R., Xu, H., Huang, W., Chen, P., Tan, M., and Gan, C. Dense regression network for video grounding. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- Zhang, M., Li, H., Pan, S., Chang, X., and Su, S. Overcoming multi-model forgetting in one-shot nas with diversity maximization. In *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 7806–7815, 2020a.
- Zhang, Y., Zhao, P., Niu, S., Wu, Q., et al. Online adaptive asymmetric active learning with limited budgets. *IEEE Transactions on Knowledge and Data Engineering*, 2019.
- Zhang, Y., Wei, Y., Wu, Q., Zhao, P., Niu, S., Huang, J., and Tan, M. Collaborative unsupervised domain adaptation for medical image diagnosis. *IEEE Transactions on Image Processing*, 29:7834–7844, 2020b.
- Zoph, B. and Le, Q. V. Neural architecture search with reinforcement learning. In *International Conference on Learning Representations*, 2017.
- Zoph, B., Vasudevan, V., Shlens, J., and Le, Q. V. Learning transferable architectures for scalable image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 8697–8710, 2018.