

Joint Search of Data Augmentation Policies and Network Architectures

Taiga Kashima^{1*}, Yoshihiro Yamada², Shunta Saito²

¹ The University of Tokyo,

² Preferred Networks inc., Japan

Abstract

The common pipeline of training deep neural networks consists of several building blocks such as data augmentation and network architecture selection. AutoML is a research field that aims at automatically designing those parts, but most methods explore each part independently because it is more challenging to simultaneously search all the parts. In this paper, we propose a joint optimization method for data augmentation policies and network architectures to bring more automation to the design of training pipeline. The core idea of our approach is to make the whole part differentiable. The proposed method combines differentiable methods for augmentation policy search and network architecture search to jointly optimize them in the end-to-end manner. The experimental results show our method achieves competitive or superior performance to the independently searched results.

Introduction

Neural network architectures that achieve high accuracy at specific tasks (e.g., image classification (Simonyan and Zisserman 2015; He et al. 2016; Huang et al. 2017)) have been designed through many trials and errors by humans, which requires a high level of expertise and have been a burden to practitioners and researchers in the machine learning community. Thus, AutoML has attracted the attention from many of them, which aims at automatically choosing a better design for each part in the machine learning pipeline, e.g., data augmentation (Cubuk et al. 2019; Lim et al. 2019; Ho et al. 2019; Zhang et al. 2020), network architecture (Liu, Simonyan, and Yang 2019; Zoph, Vasudevan, and Le 2018), loss functions (Li et al. 2019), or learning parameters (Snoek, Larochelle, and Adams 2012; Baydin et al. 2018). In particular, several methods proposed in the field of Neural Architecture Search (NAS) started to achieve comparable accuracy to manually designed networks in a few tasks such as image classification (Zoph, Vasudevan, and Le 2018; Cai, Zhu, and Han 2019; Chen et al. 2019). However, the early methods of NAS typically

*This work was done during internship of Preferred Networks inc., Japan
Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

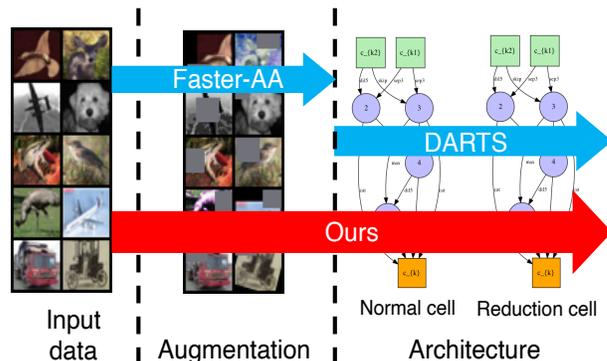


Figure 1: Comparison of search spaces. Our method jointly explores data augmentation policies and network architectures by combining differentiable methods for each part.

have some practical problems mainly due to the large requirements for computing resources such as memories and GPUs (Zoph and Le 2017). Then, recent studies tried to find better architecture with more efficient approaches (Liu, Simonyan, and Yang 2019; Pham et al. 2018; Real et al. 2019; Hu et al. 2020), which lead the optimization goal of NAS to become varied. (Wu et al. 2019) proposed a method to search network architectures that can achieve high accuracy under the limitation of computing resources used for inference, and (He et al. 2019) proposed to consider memory and power efficiency at inference time by model compression. Most of the methods mainly focused on the network architecture as their optimization target because it has large impact on the entire performance, but as the optimization goal becomes diversified, the target of AutoML also becomes varied over different building blocks of the training, e.g., data augmentation and learning parameters.

In terms of search spaces, even a single part in the training pipeline has a large one. For example, an efficient NAS method, ENAS (Pham et al. 2018), still has a large search space over 1.3×10^{11} possible networks. Auto Augment (Cubuk et al. 2019) is a method to automatically choose the data augmentation policies during training, and the search space has roughly 2.9×10^{32} possibilities. It

means that searching over all possibilities of the combination of these two parts will have about 3.8×10^{43} possibilities, which brings significant difficulty to automatic exploration. In addition, searching for network architectures and data augmentation policies can have different objectives. The former explores mainly to minimize a loss function, while the latter also tries to increase the variety of training data. These difficulties have prevented previous research from searching over those two parts jointly, so that most of the studies have focused on automatic exploration of one part at a time. Although there are a few studies that tried to explore learning parameters and network architectures jointly (Dong et al. 2020; Zela et al. 2018), few research have attempted joint optimization of data augmentation policies and network architectures to the best of our knowledge.

In this paper, we propose a joint optimization method for both data augmentation policies and network architectures. As stated above, these two parts can have large search spaces in total and different objectives, which implies that combining existing methods for each part straightforwardly would be intractable. Additionally, during the progress of architecture search, networks in different phases of training can desire different data augmentation policies for better generalization ability. Motivated by these intuitions, we propose an end-to-end differentiable approach to optimize both parts in the training pipeline simultaneously. Fig. 1 shows the difference of our search space from previous successful methods for each part. Specifically, we jointly optimize the differentiable approaches for augmentation policy search (Hataya et al. 2020) and architecture search (Liu, Simonyan, and Yang 2019). We firstly apply differentiable operations to the input data, and then use the transformation outputs as the inputs for the differentiable NAS method. It enables to optimize the augmentation policies with the gradients come from the upstream NAS method because the entire pipeline is fully differentiable, so that we can train both parts simultaneously in the end-to-end manner. We consider a combination of existing methods each of which is performed independently on either of augmentation policy search or architecture search as the baseline, and compare the performance with ours. The experimental results show that our method achieves competitive or superior performance in common benchmarks of image classification.

Preliminaries

As the differentiable methods for augmentation policy search and architecture search, we adopt Faster Auto Augmentation (Hataya et al. 2020) and DARTS (Liu, Simonyan, and Yang 2019). We briefly summarize them in this section.

Differentiable Data Augmentation

Data augmentation is a series of transformation applied on the input data. Typically, we have to choose which operations should be applied with what magnitudes. Several methods for automatic search of probability distribution on the selection of operators and their magnitudes have been proposed (Cubuk et al. 2019; Lim et al. 2019;

Hataya et al. 2020). The Faster-AA considers that a policy consists of L sub-policies each of which has K consecutive operations. Each operation out of $\#\mathcal{O}$ operations has a probability $p_O \in [0, 1]$ which represents how likely the operation is adopted in a sub-policy and a magnitude $\mu_O \in [0, 1]$ which controls the transformation. Therefore, the search space is $(\#\mathcal{O} \times [0, 1] \times [0, 1])^{KL}$ in total, where \mathcal{O} is a set of possible operations. Applying an operation O to an input data X is formulated as:

$$X \rightarrow \begin{cases} O(X; \mu_O) & \text{(with the probability } p_O) \\ X & \text{(with the probability of } 1 - p_O), \end{cases} \quad (1)$$

and note that the Gumbel trick (Jang, Gu, and Poole 2016) is used to make the probability differentiable. A sub-policy is a series of K operations, and during training, the output of k -th operation X' from an input X is calculated as a weighted sum over all possible $\#\mathcal{O}$ operations as follows:

$$X' = \sum_{n=1}^{\#\mathcal{O}} [\sigma_\eta(\mathbf{z}_k)]_n O_k^{(n)}(X; \mu_k^{(n)}, p_k^{(n)}), \quad (2)$$

$$\text{s.t. } \sum_{n=1}^{\#\mathcal{O}} [\sigma_\eta(\mathbf{z}_k)]_n = 1, \quad (3)$$

where σ_η is a softmax function with a temperature parameter $\eta > 0$, and $\mathbf{z}_k \in \mathbb{R}^{\#\mathcal{O}}$ denotes the learnable parameter for the distribution of operator selection. During inference, the k -th operation is sampled from the categorical distribution $\text{Cat}(\sigma_\eta(\mathbf{z}_k))$, so that we obtain transformed data X' by Eq. (1).

Differentiable NAS

DARTS (Liu, Simonyan, and Yang 2019) is a differentiable neural architecture search method which focuses on searching the inside structures of normal cells and reduction cells that are finally stacked up to build a deep network. Each cell is represented as a directed acyclic graph (DAG) consisting of N nodes which represent intermediate features (e.g., feature maps). A cell takes two input nodes and one output node, and an edge $f^{(i,j)}$ between two nodes i, j represents an operation such as convolution or pooling. A node j has the connections with all the previous nodes $i < j$ in topological ordering of the DAG, so that the search space of DARTS is roughly

$$|\mathcal{F}|^2 \prod_{k=1}^{N-3} \frac{k(k+1)}{2}, \quad (4)$$

where \mathcal{F} is a set of candidate operations. To make the search space continuous, DARTS relaxes the categorical choice of an operation to a softmax over possible operations:

$$\bar{f}^{(i,j)} = \sum_{f \in \mathcal{F}} \frac{\exp(\alpha_f^{(i,j)})}{\sum_{f'} \exp(\alpha_{f'}^{(i,j)})} f(x). \quad (5)$$

After the training, a discrete architecture is determined by replacing each mixed operation $\bar{f}^{(i,j)}$ with the most likely operation $f^{(i,j)} = \arg \max_{f \in \mathcal{F}} \alpha_f^{(i,j)}$.

DARTS performs a bilevel optimization on the architecture parameter α and the model weights w using two sets of training and validation data.

$$\min_{\alpha} \mathcal{L}_{val}(w^*(\alpha), \alpha), \quad (6)$$

$$\text{s.t. } w^*(\alpha) = \arg \min_w \mathcal{L}_{train}(w, \alpha), \quad (7)$$

where \mathcal{L}_{val} and \mathcal{L}_{train} are loss functions calculated with validation data and train data, respectively. $w^*(\alpha)$ denotes the optimal model weights for an architecture α . Eq. (6) has an inner optimization for $w^*(\alpha)$, so that evaluating the gradient of \mathcal{L}_{val} w.r.t. α can be prohibitive. Therefore, DARTS approximates $w^*(\alpha)$ by adapting w only after a single training step on the training data as follows:

$$w^*(\alpha) = w - \xi \nabla_w \mathcal{L}_{train}(w, \alpha). \quad (8)$$

Then, during a single step in the iterative optimization procedure of DARTS, it solves Eq. (6) and Eq. (7) alternately.

Method

We sequentially combine Faster-AA and DARTS, then optimize both in the end-to-end manner. We solve another bilevel optimization problem for our entire search space. Specifically, augmentation policies and network architectures are both optimized by minimizing a loss function on the validation dataset, while the network weights are optimized using the training dataset. This bilevel optimization is formulated as:

$$\min_{\alpha, \mathbf{z}_k, p_O, \mu_O} \mathcal{L}_{val}(w^*(\alpha), \alpha, \mathbf{z}_k, p_O, \mu_O), \quad (9)$$

$$\text{s.t. } w^*(\alpha) = \arg \min_w \mathcal{L}_{train}(w, \alpha, \mathbf{z}_k, p_O, \mu_O), \quad (10)$$

and we adopt the first-order approximation for the architecture gradient for Eq. (6) same as in DARTS to speed-up the optimization, i.e., we set ξ to 0 in Eq. (8). Then, we iteratively solve Eq. (9) and Eq. (10). As a loss function, we use the cross-entropy loss for both \mathcal{L}_{train} and \mathcal{L}_{val} .

To solve Eq. (9), we first apply a series of differentiable data augmentation operations $O^{(n)}$ ($n = 1, \dots, \#\mathcal{O}$), then give the transformed data to the network to solve Eq. (10). We outline this algorithm in Alg. 1.

The total computational resources we require for the joint optimization additionally to what DARTS requires are relatively small. Specifically, the additional space complexity is only $KL(\#\mathcal{O} \times 3)$, while the search space of Faster-AA is $(\#\mathcal{O} \times [0, 1] \times [0, 1])^{KL}$ which is large. In addition to it, our entire system is end-to-end differentiable, and augmentation policy search and network architecture search are jointly performed to minimize the same loss function, so that the gradients for updating policy and architecture parameters are obtained via a single backpropagation. This advantage of our end-to-end differentiable approach enables to conduct joint optimization for policy search and architecture search with few additional space and time complexity compared to the case if we apply Faster-AA and DARTS independently and combine the results.

The original Faster-AA uses a critic network to consider a classification loss and the WGAN-GP loss (Gulrajani et

Algorithm 1 Joint optimization

```

1: while not converged do
2:   // solve equation (9)
3:   sample  $X \sim D_{val}$ 
4:   apply equation (2) to  $X$ 
5:   calculate  $L_{val}(w, \alpha, \mathbf{z}_k, p_O, \mu_O)$  with  $X'$ 
6:   calculate gradients of  $L_{val}$  w.r.t.  $\mathbf{z}_k, p_O, \mu_O, \alpha$ 
7:   update Faster-AA parameters ( $\mathbf{z}_k, p_O, \mu_O$ )
8:   update DARTS parameters  $\alpha$  by gradient descent
9:   // solve equation (10)
10:  sample  $X \sim D_{train}$ 
11:  apply equation (2) to  $X$ 
12:  calculate  $L_{train}(w, \alpha, \mathbf{z}_k, p_O, \mu_O)$  with  $X'$ 
13:  calculate gradients of  $L_{train}$  w.r.t.  $w$ 
14:  update the network weights  $w$  by gradient descent
15: end while
16: Derive the final policy and architecture

```

al. 2017) that encourage the distribution of transformed data to be as close to the original data distribution as possible. For the same purpose, we can also exploit the network under searching with DARTS as a critic network to encourage transformed data to remain in the same classes before transformations. However, DARTS only uses a cross-entropy as the loss for architecture search, while the critic network in Faster-AA also considers WGAN-GP. In this paper, we adopt a single unified loss function for the both of policy search and architecture search and did not introduce any critic network for simplicity of the entire framework and computational efficiency.

Experiments

We compare our joint optimization model with the original DARTS and the baseline which combines the results of Faster-AA and DARTS that are optimized independently from each other. In the baseline, the learned policy by Faster-AA is transferred to be used for the training of DARTS. We conducted the comparison on three datasets, CIFAR-10, CIFAR-100, and SVHN.

We first carefully re-implemented Faster-AA¹ and DARTS² based on their authors' implementation. Then, we confirmed our implementation successfully reproduced the reported scores in those papers. It should be noted that we adopt the same search space as the original paper for DARTS, while we exclude the cutout operation from the search space for the Faster-AA part in our framework as the authors do in their implementation. According to a comment which has been left in their code, the cutout operation makes the optimization unstable. Therefore, the target operations in the augmentation policy search are *shear X*, *shear Y*, *translate X*, *translate Y*, *rotate*, *auto contrast*, *horizontal flip*, *invert*, *equalize*, *solarize*, *posterize*, *contrast*, *color brightness*, *sharpness*, and *sample pairing*. We set the num-

¹source code from: https://github.com/moskomule/dda/tree/fasteraa/faster_autoaugment

²source code from: <https://github.com/quark0/darts>

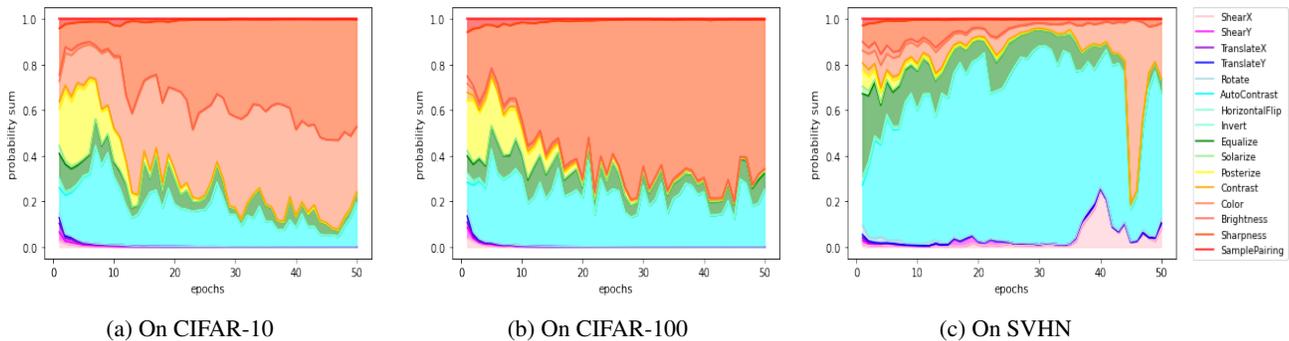


Figure 2: Probability distribution of the augmentation policy selection over time.

ber of sub-policies $L = 10$ and the number of operations in a sub-policy $K = 2$, which are the same in the original settings. Additionally, as preprocessings in the baseline, random cropping with zero-padding and random horizontal flipping are always applied. After those preprocessings, we apply transformations under the policy search with Faster-AA. The cutout operation is always applied after the Faster-AA part in the baseline, although it is not included in the search space.

For the baseline, we first obtain a Faster-AA policy independently from DARTS by following the same experimental settings used in the original paper, i.e., the baseline uses WideResnet40-2 (Zagoruyko and Komodakis 2016) as the architecture to search policies for 20 epochs. Next, the network architecture for the baseline is explored by DARTS in the same manner with the original paper. Those policy and architecture obtained by the existing methods are combined together, then we train the model weights for 600 epochs. As for the searching epochs of architectures, we conduct two different total epochs (25 and 50) both for the baseline and proposed framework. Because we found that training of DARTS is unstable on CIFAR-100 and SVHN as mentioned in (Bi et al. 2019). *We repeated experiments three times to see average scores with standard deviations.*

Discussion

Table 1: Comparison in classification accuracy. (1) “DARTS” only searches architecture. (2) “Baseline” separately searches policies and architectures. (3) “Ours” is the proposed joint search method. Our method achieves competitive or superior results compared with the baseline.

method	CIFAR-10	CIFAR-100	SVHN
searching epoch is 50			
DARTS	97.33±0.12	76.21±3.76	97.85±0.08
Baseline	97.55±0.30	77.96±3.12	98.02±0.03
Ours	97.40±0.03	79.02±2.14	97.92±0.12
searching epoch is 25			
DARTS	97.14±0.04	82.91±0.30	97.94±0.06
Baseline	97.29±0.03	84.17±0.29	98.03±0.05
Ours	97.46±0.09	83.81±0.49	97.82±0.08

We present the experimental results in Table 1. Our proposed method achieves competitive or superior results compared to the baseline in both searching epochs. As stated above, the baseline uses WideResnet40-2 during policy search, and it can be different from the architecture found by DARTS for the final stage to learn the model weights for 600 epochs. Although the policies found with WideResnet40-2 show high accuracy at several results, which might be derived from the suitability of WideResnet40-2 for augmentation policy search. If that is the case, the model selection largely affects on the policy search, so that the human expertise is still required. On the other hand, our proposed method does not require humans to select the network architecture for policy search and achieves competitive or even superior results to Faster-AA. We believe that the joint optimization approach has more potential to obtain high performance and reduces human expertise.

Fig. 2 shows how the categorical distribution to choose augmentation policies changes over the time. We found that the augmentation policies obtained with our method choose color enhancement operations such as *color* or *auto contrast* more often than geometric operations such as *rotate* from the transition. This trend of the resulting policies is also reported in other papers of automatic augmentation policy search (Cubuk et al. 2019). In Figure 2c, the policy dramatically changes after 45 epochs, and we found that the network architecture under searching also largely changed at the same timing, which may imply that the optimal policy is different depending on the network architecture.

Conclusion

In this paper, we proposed a method to jointly optimize data augmentation policy and network architecture. The proposed method combines differentiable methods for policy search and architecture search to jointly optimize them in the end-to-end manner. The experimental results showed that our method achieves competitive or superior performance to independently searched results in common benchmarks of image classification. Our joint optimization approach may be able to include the other parts such as learning rates. Hence, we will attempt to bring more automation to the design of training pipeline with this end-to-end differentiable approach for the future work.

References

- Baydin, A. G.; Cornish, R.; Rubio, D. M.; Schmidt, M.; and Wood, F. 2018. Online learning rate adaptation with hyper-gradient descent. In *ICLR*.
- Bi, K.; Hu, C.; Xie, L.; Chen, X.; Wei, L.; and Tian, Q. 2019. Stabilizing darts with amended gradient estimation on architectural parameters. In *Arxiv: 1910.11831*.
- Cai, H.; Zhu, L.; and Han, S. 2019. Proxylessnas: Direct neural architecture search on target task and hardware. In *ICLR*.
- Chen, X.; Xie, L.; Wu, J.; and Tian, Q. 2019. Progressive differentiable architecture search: Bridging the depth gap between search and evaluation. In *ICCV*.
- Cubuk, E. D.; Zoph, B.; Manè, D.; Vasudevan, V.; and Le, Q. V. 2019. Autoaugment: Learning augmentation strategies from data. In *CVPR*.
- Dong, X.; Tan, M.; Yu, A. W.; Peng, D.; Gabrys, B.; and Le, Q. V. 2020. Autohas: Differentiable hyper-parameter and architecture search. In *Arxiv: 2006.03656*.
- Gulrajani, I.; Ahmad, F.; Arjovsky, M.; Dumoulin, V.; and Courville, A. 2017. Improved training of wasserstein gans. In *NeurIPS*.
- Hataya, R.; Zdenek, J.; Yoshizoe, K.; and Nakayama, H. 2020. Faster autoaugment: Learning augmentation strategies using backpropagation. In *ECCV*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *CVPR*.
- He, Y.; Li, J.; Liu, Z.; Wang, H.; Li, L.; and Han, S. 2019. Amc: Automl for model compression and acceleration on mobile devices. In *ECCV*.
- Ho, D.; Liang, E.; Stoica, I.; Abbeel, P.; and Chen, X. 2019. Population based augmentation: Efficient learning of augmentation policy schedules. In *ICML*.
- Hu, S.; Xie, S.; Zheng, H.; Liu, C.; Shi, J.; Liu, X.; and Lin, D. 2020. Dsnas: Direct neural architecture search without parameter retraining. In *CVPR*.
- Huang, G.; Liu, Z.; van der Maaten, L.; and Weinberger, K. Q. 2017. Densely connected convolutional networks. In *CVPR*.
- Jang, E.; Gu, S.; and Poole, B. 2016. Categorical reparametrization with gumbel-softmax. In *ICLR*.
- Li, C.; Yuan, X.; Lin, C.; Guo, M.; and W. Wu, J. Yan, W. O. 2019. Am-lfs: Automl for loss function search. In *ICCV*.
- Lim, S.; Kim, I.; Kim, T.; Kim, C.; and Kim, S. 2019. Fast autoaugment. In *NeurIPS*.
- Liu, H.; Simonyan, K.; and Yang, Y. 2019. Darts: Differentiable architecture search. In *ICLR*.
- Pham, H.; Guan, M. Y.; Zoph, B.; Le, Q. V.; and Dean, J. 2018. Efficient neural architecture search via parameter sharing. In *ICML*.
- Real, E.; Aggarwal, A.; Huang, Y.; and Le, Q. V. 2019. Regularized evolution for image classifier architecture search. In *AAAI*.
- Simonyan, K., and Zisserman, A. 2015. Very deep convolutional networks for large-scale image recognition. In *ICLR*.
- Snoek, J.; Larochelle, H.; and Adams, R. P. 2012. Practical bayesian optimization of machine learning algorithms. In *NeurIPS*.
- Wu, B.; Dai, X.; Zhang, P.; Wang, Y.; Sun, F.; Wu, Y.; Tian, Y.; Vajda, P.; Jia, Y.; and Keutzer, K. 2019. Fbnet: Hardware-aware efficient convnet design via differentiable neural architecture search. In *CVPR*.
- Zagoruyko, S., and Komodakis, N. 2016. Wide residual networks. In *BMVC*.
- Zela, A.; Klein, A.; Falkner, S.; and Hutter, F. 2018. Towards automated deep learning: Efficient joint neural architecture and hyperparameter search. In *W.S. in ICLR*.
- Zhang, X.; Wang, Q.; Zhang, J.; and Zhong, Z. 2020. Adversarial autoaugment. In *ICLR*.
- Zoph, B., and Le, Q. V. 2017. Neural architecture search with reinforcement learning. In *ICLR*.
- Zoph, B.; Vasudevan, V.; and Le, J. S. Q. V. 2018. Learning transferable architectures for scalable image recognition. In *CVPR*.