

Efficient Model Performance Estimation via Feature Histories

Shengcao Cao

CMU-RI-TR-21-09

April 2021

The Robotics Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213

Thesis Committee:

Kris M. Kitani, Chair

Deva Ramanan

Xiaofang Wang

*Submitted in partial fulfillment of the requirements
for the degree of Master of Science in Robotics.*

Abstract

An important step in the task of neural network design, such as hyper-parameter optimization (HPO) or neural architecture search (NAS), is the evaluation of a candidate model’s performance. Given fixed computational resources, one can either invest more time training each model to obtain more accurate estimates of final performance, or spend more time exploring a greater variety of models in the configuration space. In this work, we aim to optimize this exploration-exploitation trade-off in the context of HPO and NAS for image classification by accurately approximating a model’s maximal performance early in the training process. In contrast to recent accelerated NAS methods customized for certain search spaces, e.g., requiring the search space to be differentiable, our method is flexible and imposes almost no constraints on the search space. Our method uses the evolution history of features of a network during the early stages of training to build a proxy classifier that matches the peak performance of the network under consideration. We show that our method can be combined with multiple search algorithms to find better solutions to a wide range of tasks in HPO and NAS.

Acknowledgments

First and foremost, I would like to thank my thesis committee members: Prof. Kris M. Kitani, Prof. Deva Ramanan, and Xiaofang Wang. I'm grateful for Kris' guidance and support over past years. I always feel fortunate to have an advisor who spends days polishing our papers before every submission deadline, who refines every detail in our presentations, and who never forgets to invite us to BBQ parties at home. I want to thank Xiaofang for all his generous help since I first came to CMU. Our discussions and collaborations have inspired most of my research ideas, and also taught me how a graduate student should think. I also appreciate Deva for joining my thesis committee, and providing me with valuable advice for the thesis.

I would like to thank members of KLab and Smith Hall, who have offered countless interesting presentations and discussions. Although we have highly diverse research interest and topics, it is always beneficial to exchange different ideas about our work. I would like to thank professors of all my courses taken at CMU, whose teaching is of high quality.

I would like to thank my roommates and friends online. Their companion always reminds me of the miraculous connections between me and others. I can hardly imagine the life without them in this year of the COVID pandemic.

Special thanks to my parents, for their unconditional and everlasting supports.

Contents

| | | |
|----------|--|-----------|
| 1 | Introduction | 1 |
| 2 | Related Work | 5 |
| 2.1 | Performance estimation | 5 |
| 2.2 | Hyperparameter optimization | 6 |
| 2.3 | Neural architecture search | 6 |
| 2.4 | Ensemble of history models | 7 |
| 3 | Approach | 9 |
| 3.1 | Performance Estimation via Feature Histories | 9 |
| 3.1.1 | Preliminaries | 9 |
| 3.1.2 | Observation about Feature Histories | 10 |
| 3.1.3 | Classifier Ensemble via Feature Histories | 12 |
| 3.2 | Combination with Search Algorithms | 15 |
| 4 | Experiments | 17 |
| 4.1 | Performance Estimation for CNNs | 17 |
| 4.2 | Neural Architecture Search | 18 |
| 4.2.1 | NAS-Bench-201 | 18 |
| 4.2.2 | DARTS | 21 |
| 4.3 | Hyperparameter Optimization | 22 |
| 4.3.1 | Data Augmentation | 23 |
| 4.3.2 | Scaling Factors for ResNet | 24 |
| 5 | Conclusion | 27 |
| A | Experimental Details | 29 |
| A.1 | Classifier Ensemble | 29 |
| A.2 | Performance Estimation for CNNs | 29 |
| A.3 | NAS-Bench-201 | 30 |
| A.3.1 | Performance Error and Relative Ranking | 30 |
| A.3.2 | Architecture Search | 30 |
| A.4 | DARTS | 31 |
| A.5 | Data Augmentation | 31 |

| | |
|--|-----------|
| A.6 Scaling Factors for ResNet | 32 |
| Bibliography | 33 |

List of Figures

- 1.1 **An illustration of our proposed proxy classifier using feature histories:** Our method aims to approximate the model’s final performance without actually training it until convergence. The major steps are: 1) collecting feature histories of early training stages, 2) finding the optimal linear classifiers for each checkpoint, and 3) ensembling them to build a proxy classifier that can simulate the outputs of the model at convergence. 2

- 3.1 **Performance of the original classifiers, optimized classifiers, and classifier ensembles in training ResNet-18 on CIFAR-100:** In this example, the classifier ensemble already reaches over 77% accuracy at epoch 75, which speeds up the training by about 2.7x compared to full training till convergence. 11

- 3.2 **Classification results of some sampled validation images during training ResNet-18 on CIFAR-100:** A light gray pixel denotes that image is correctly classified by the classifier at the given epoch, and a black pixel denotes a wrong classification. The vertical axis indicates the image indices, sorted by their accuracy for better visualization. 12

- 3.3 **Breakdown of classification results of validation images during training ResNet-18 on CIFAR-100:** Here we compare the classification results with the final model. True positive: Both the current classifier and the final ResNet-18 can correctly classify. True negative: Both the current classifier and the final ResNet-18 cannot correctly classify. False positive: The current classifier is correct but the final ResNet-18 is wrong. False negative: The current classifier is wrong but the final ResNet-18 is correct. 13

- 4.1 **Performance estimation for different architectures including VGG-16, ResNet-18, ResNet-50, and MobileNetV2 on CIFAR-10/CIFAR-100 dataset: Original** denotes the original test accuracy at each epoch. **Ours** denotes the performance of our proposed proxy classifier using the feature histories up to the given epoch. Our method can reach an accuracy closer to the final accuracy very early. 18

- 4.2 **Performance estimation for different training budgets from 25% (50 epochs) to 100% (200 epochs) of training ResNet-18 on CIFAR-10/CIFAR-100 dataset:** The linear learning rate schedule is adjusted accordingly. 19

| | | |
|-----|---|----|
| 4.3 | Absolute difference between the estimated performance and the true performance: Original denotes the original accuracy metrics at each epoch. Ours denotes the predicted performance by our method using the checkpoints up to the given epochs. LCNet denotes the extrapolated final performance given by LCNet. For fair comparison, we slightly shift Ours rightwards because of the computational overhead. | 20 |
| 4.4 | Relative ranking measured by Kendall’s τ coefficient between the estimated performance and the true performance. | 21 |

List of Tables

- 4.1 Search results for NAS-Bench-201 architectures on CIFAR-10. 22
- 4.2 Comparison with state-of-the-art NAS methods that uses the DARTS search space on CIFAR-10. 23
- 4.3 Search results for RandAugment hyperparameter optimization on a subset of CIFAR-10. 24
- 4.4 Search results for scaling factors of ResNet-18 on CIFAR-100. 25
- 4.5 ImageNet performance of scaled ResNet-18 searched on CIFAR-100. 25

Chapter 1

Introduction

Identifying the optimal hyperparameters or the best architecture is important for maximizing the performance of neural networks. Accordingly, algorithms for hyperparameter optimization (HPO) and neural architecture search (NAS) have been proposed to automatically select the optimal hyperparameters and architectures in a data-driven manner. Existing HPO or NAS methods typically require that many candidate configurations of hyperparameters or architectures are evaluated. However, such evaluation is extremely expensive as fully training one model until convergence may take several GPU days given the complexity of modern deep learning models and datasets. This calls for efficient methods that can approximate the model’s final performance early in the training process.

We propose to leverage *feature histories*, that is, checkpoints of evolving deep features prior to the final fully-connected layer of a neural network during training, to predict what the output of the network might be at convergence. Our proposed approach is motivated by this key observation: Many images which are correctly classified when the network converges, keep changing between being classified correctly and incorrectly at the early stage of training. Only using a particular checkpoint early in the training will lead to an inaccurate approximation for the final accuracy. However, for a particular image, the consensus of classification results by early checkpoints is usually the same as the final model’s decision. Therefore, we propose to build a proxy classifier based on the early feature histories, whose behavior coincides with the network at convergence, and use its evaluated performance as an approximation of the model’s final per-

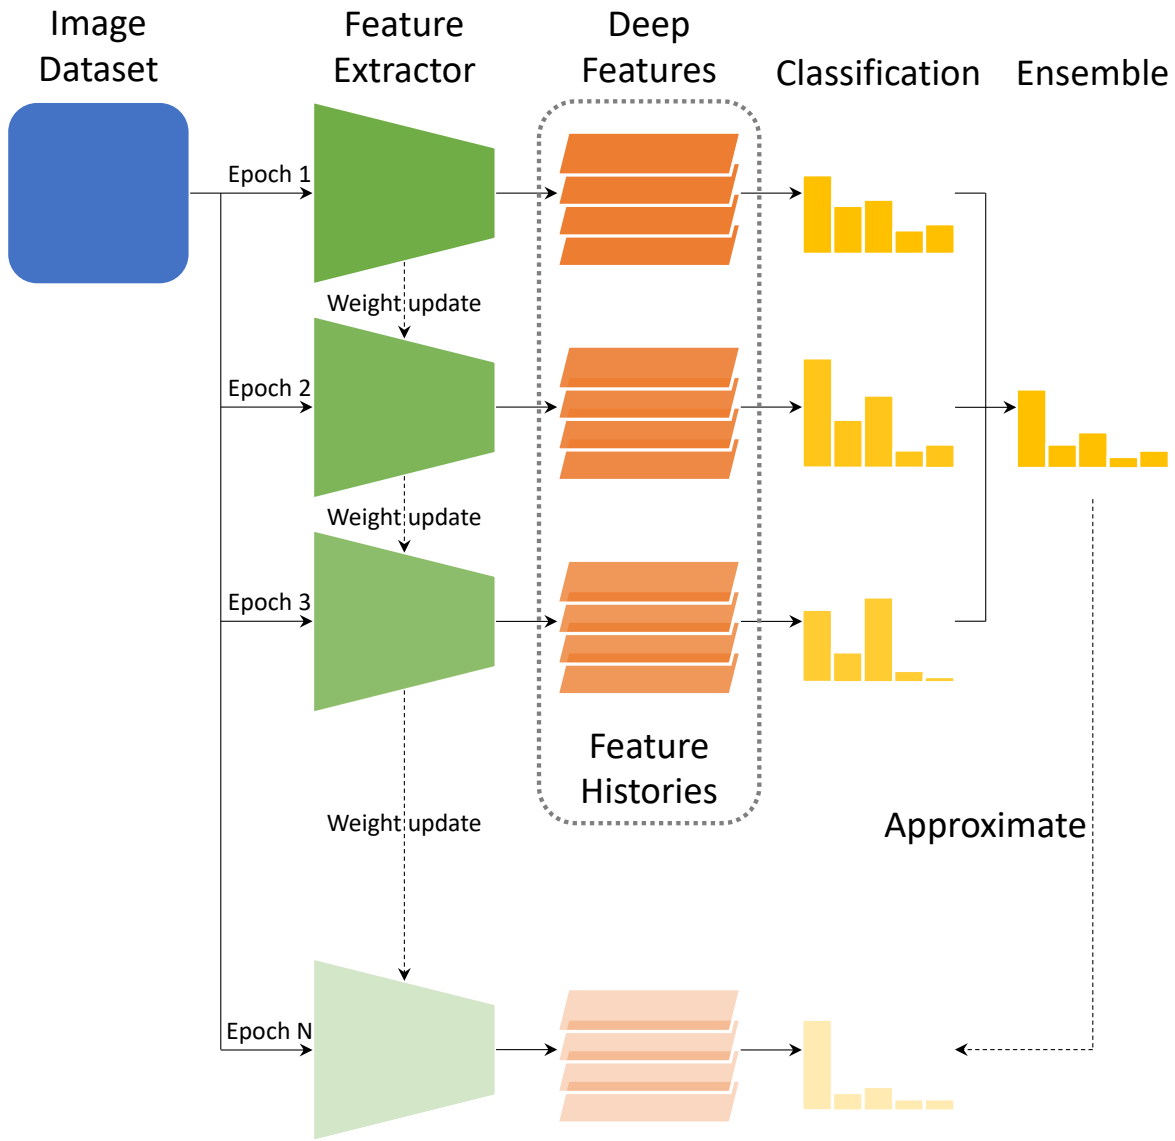


Figure 1.1: **An illustration of our proposed proxy classifier using feature histories:** Our method aims to approximate the model’s final performance without actually training it until convergence. The major steps are: 1) collecting feature histories of early training stages, 2) finding the optimal linear classifiers for each checkpoint, and 3) ensembling them to build a proxy classifier that can simulate the outputs of the model at convergence.

formance. In other words, to approximate a candidate model’s final performance, we no longer need to train it until convergence, but create a proxy based on the feature histories to simulate the converged network. This proxy classifier can bring better assessment of candidate models within limited training budgets, thus benefiting the task of HPO and NAS. A brief illustration of our method is shown in Figure 1.1. We will support this observation with experimental results in Section 3.1.2. Our empirical results show that this method applies to various architecture families, including VGG [30], ResNet [14], MobileNet [28], and also architectures in different NAS tasks, including DARTS [24] and NASBench-201 [11].

Our proposed method has these advantages:

- **Versatility:** Our method does not depend on any pre-training, external data, or specific search space, but only makes use of the feature histories. It can be incorporated into several search algorithms for general purpose, and applied to a wide range of tasks in HPO and NAS.
- **Accuracy:** Experimental results show that the performance of our proposed proxy classifier can accurately approximate the final performance of given candidate model configurations. The relative performance ranking of configurations is also well preserved.
- **Efficiency:** Our method is able to approximate the final performance using only partially trained models. The extra overhead is no more than one epoch of training time.

Chapter 2

Related Work

We propose a novel method for approximating a model’s final performance at an early training stage, which can be applied to hyper-parameter optimization (HPO) and neural architecture search (NAS) tasks. Here we briefly summarize the previous work that our method is based on.

2.1 Performance estimation

To accelerate hyperparameter optimization, [10] proposes to extrapolate the learning curves based on early training stage and terminate bad configurations accordingly. LCNet [19] improves the estimation of the learning curve with a Bayesian neural network. [1] applies this strategy to NAS, using architectural hyperparameters for the learning curve prediction. In contrast to these approaches, our method use feature histories during network optimization, instead of performance metrics, to approximate the final performance. Our method does not require training another prediction network, or manually designed learning curve modeling. It has also been shown that a carefully selected learning rate schedule can produce high-performance models with a limited budget [23]. However, this type of learning rate schedule is impractical in the case where we need to dynamically adjust the budget, which is common in hyperparameter optimization algorithms. Other methods specifically designed for accelerating performance estimation in NAS includes inheriting network weights through network morphisms [4], weight sharing across architectures [2, 25, 29], and optimize budgeted performance estimation [40].

These methods are usually constrained by the customized architecture search space, while our method can be easily extended to various tasks.

2.2 Hyperparameter optimization

Automated hyperparameter optimization can improve the performance of deep learning models while reducing human effort. Model-free methods including random search and grid search, can be considered as the most basic HPO approach. Bayesian optimization, an effective optimization method for computational costly functions, has several variants [3, 31, 32]. More recently, [22] proposes a bandit-based strategy named HyperBand, which dynamically allocates a budget for configurations based on evaluation results. BOHB [12] combines Bayesian optimization and HyperBand to achieve the best of both worlds. Our experimental validation builds on HyperBand and BOHB and shows how our method can further improve hyperparameter optimization.

2.3 Neural architecture search

Neural architecture search Automated architecture design can discover architectures that outperform manually-designed CNNs. [41, 42] first propose to train a controller which generates architectures with reinforcement learning. Later, [26, 36] utilize evolutionary algorithms to search optimal architecture designs. DARTS [24] and its variants [5, 38] takes a different approach to greatly reduce the search cost. They relax the discrete search space of architectures to be continuous, and apply gradient descent to optimize intertwined weights and architecture parameters. More recently, EfficientNet [33] combines neural architecture search and model scaling to produce a series of efficient image backbones for various application scenarios with resource constraints. For better evaluation and comparison of different NAS algorithms, some benchmarks [11, 18, 39] have been proposed.

2.4 Ensemble of history models

[35] proposes to combine the outputs of networks from earlier epochs to improve final predictions. [16] uses cyclic learning rate to efficiently obtain multiple well-trained models and their ensemble. [13] proposes a method to discover high-accuracy paths between trained models and their ensemble. Our method can be considered as an ensemble of history models as well. In contrast to previous methods, we focus on ensembling simple linear classifiers based on feature histories, which greatly improves the efficiency of model performance evaluation.

Chapter 3

Approach

We first describe the observation that motivates this work and the details of our proposed method for efficient performance estimation in Section 3.1, and then explain how to combine our method with other search algorithms for neural architecture search (NAS) and hyperparameter optimization (HPO) applications in Section 3.2.

3.1 Performance Estimation via Feature Histories

3.1.1 Preliminaries

A modern CNN designed for image classification [14, 20, 30, 37] can be divided into two parts: the feature extractor and the linear classifier. The feature extractor part typically consists of multiple convolutional layers, normalization layers, and pooling layers. The learned feature mapping can often be transferred to other datasets or tasks including object detection [27] and segmentation [15]. In contrast, the linear classifier is one single fully connected layer, which outputs logits for each image class and learns the optimal linear boundaries with respect to the loss function. Compared with the linear classifier, the feature extractor has more parameters, is more computationally expensive and harder to optimize. Formally, we can consider a CNN as the composition of two functions:

$$f(\mathbf{x}; \mathbf{w}^{fea}, \mathbf{w}^{cls}) = h(g(\mathbf{x}; \mathbf{w}^{fea}); \mathbf{w}^{cls})$$

where $f : \mathbb{X} \rightarrow \mathbb{R}^c$ is a CNN which maps the input image space \mathbb{X} to c class logits, $g : \mathbb{X} \rightarrow \mathbb{R}^d$ is the feature extractor with parameters w^{fea} which maps images to a d -dim deep feature space, and $h : \mathbb{R}^d \rightarrow \mathbb{R}^c$ is the linear classifier with parameters w^{cls} which maps deep features to c class logits. The task of the feature extractor is to learn a mapping from images to deep features such that the features corresponding to different image classes are linearly separable. Given a fixed feature extractor mapping, it is easy to find the optimal linear boundaries in the feature space.

3.1.2 Observation about Feature Histories

During the process of CNN optimization and evaluation, we observe these phenomena:

- **Classification oscillation:** During training, the features of most images fluctuates between being classified correctly and incorrectly by the optimal linear classifier. As training continues, the classification results will become stabilized and more features will lie on the correct side of the decision boundary.
- **Consistency between historical models:** If the feature of one validation image lies on the correct side of the optimal linear boundaries most of the time, with high probability it will be correctly classified at the end of optimization. Similarly, if the feature is mostly on the wrong side of the decision boundary in history, it will be more likely to be mis-classified by the model at convergence.

Next we will use an example of training ResNet-18 on CIFAR-100 to demonstrate this observation qualitatively and quantitatively. In fact, this observation generalizes well to various CNN architectures, as we will show in Chapter 4. We train ResNet-18 for 200 epochs and save the deep features of training and validation images collected in every epoch, i.e., feature histories throughout training. We compare three classifiers:

- **Original:** The saved linear classifier at the given epoch. Evaluating this classifier with corresponding saved deep features is identical to evaluating the original ResNet-18 checkpoints.
- **Optimized:** We can also optimize the linear classifiers based on the saved training image features at given epochs and ground truth labels. This classifier represents the optimal

linear boundaries of feature histories. Note that separately optimizing the linear classifiers with saved features requires much less computational resources compared with training the whole ResNet-18.

- **Ensemble:** Based on the optimal linear classifiers, we can further create an ensemble of the most recent $K = 10$ epochs, by averaging the softmax outputs.

The performance of the three types of classifiers are shown in Figure 3.1. The classifier ensemble is able to reach the maximal performance of ResNet-18 at an earlier training stage and the performance remains stable. This suggests that in order to reach the final performance of the given model, we can build a classifier ensemble as a proxy instead of train the model until convergence, which greatly reduces the computational cost.

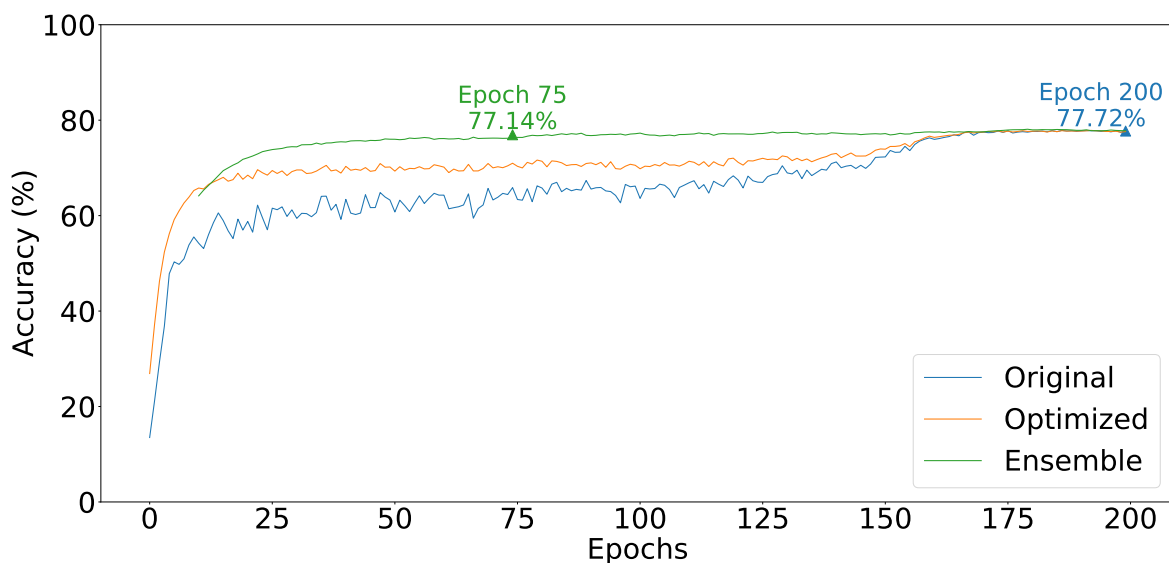


Figure 3.1: **Performance of the original classifiers, optimized classifiers, and classifier ensembles in training ResNet-18 on CIFAR-100:** In this example, the classifier ensemble already reaches over 77% accuracy at epoch 75, which speeds up the training by about 2.7x compared to full training till convergence.

We can also take a closer look at the individual validation images. Figure 3.2 shows the qualitative classification results of three types of classifiers by sampling 200 validation images. Original classification results are rather noisy and unstable, which leads to under-estimated per-

formance in early training stages. Even with the optimized linear classifiers, most validation image features keep moving back and forth between the correct and wrong sides of the optimal linear boundaries. Nevertheless, with classifier ensembles, we are able to obtain stable classification results, whose decisions are more consistent with the fully trained model.

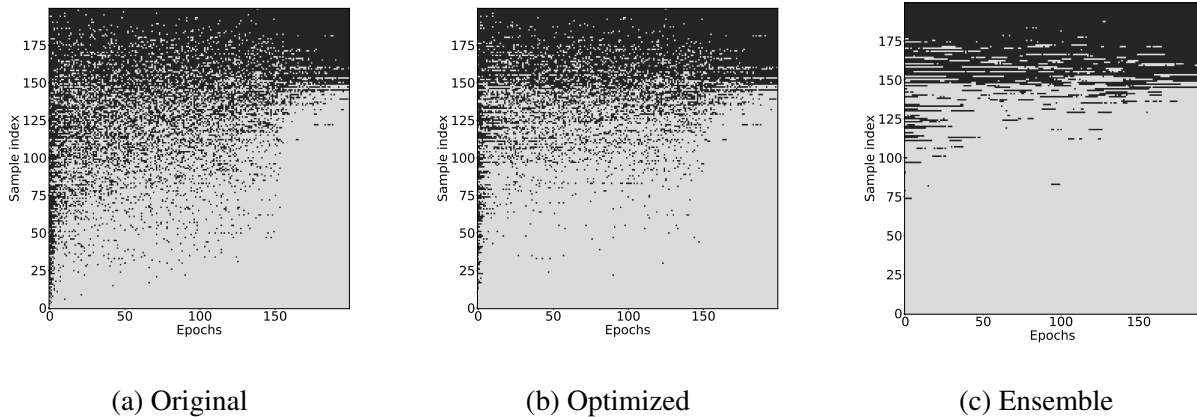


Figure 3.2: **Classification results of some sampled validation images during training ResNet-18 on CIFAR-100:** A light gray pixel denotes that image is correctly classified by the classifier at the given epoch, and a black pixel denotes a wrong classification. The vertical axis indicates the image indices, sorted by their accuracy for better visualization.

Figure 3.3 shows the breakdown of classification results. This further demonstrates that the classifier ensemble can not only reach the maximal performance earlier, but also make similar decisions for most images with the final model. In conclusion, the classifier ensemble can act as an excellent proxy for simulating the behavior of the converged model, thus providing more accurate performance estimation at an early training stage.

3.1.3 Classifier Ensemble via Feature Histories

Based on this key observation of classification oscillation and consistency between historical models, we propose to use the classifier ensemble based on the history information of deep features to approximate the final performance of a CNN at an early training stage. Given saved feature histories, we can obtain the optimal linear classifiers for each checkpoint. Then we can test the features of validation images with these classifiers, and decide if the features are correctly

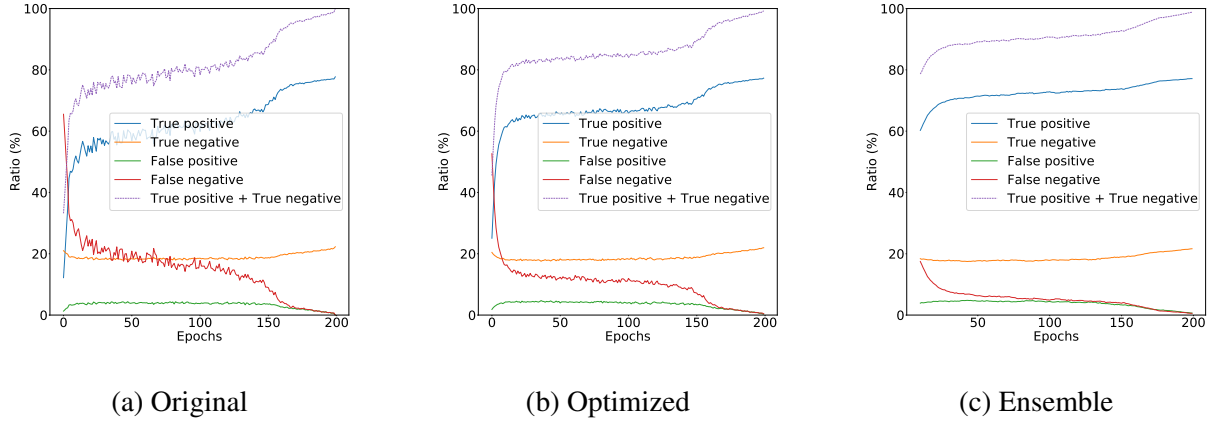


Figure 3.3: **Breakdown of classification results of validation images during training ResNet-18 on CIFAR-100:** Here we compare the classification results with the final model. True positive: Both the current classifier and the final ResNet-18 can correctly classify. True negative: Both the current classifier and the final ResNet-18 cannot correctly classify. False positive: The current classifier is correct but the final ResNet-18 is wrong. False negative: The current classifier is wrong but the final ResNet-18 is correct.

positioned most of the time during training. Thus, we can approximate the set of images which will be correctly classified at the end of training, and give a more accurate estimation of the final performance.

The implementation is straight-forward and simple, as summarized in Algorithm 1. In addition to a typical CNN training loop, we only need to explicitly save intermediate features (line 4, 7), optimize the linear classifiers based on saved features (line 9, 10, 11), and ensemble them (line 12, 13) to acquire a performance estimation closer to the final performance that the network can achieve. More details are discussed in Appendix A.1 .

One concern may be the extra computational overhead of this performance estimation strategy. In fact, compared with a typical training loop, our method usually only requires less than one half epoch time of training the original network, which is almost negligible since we often need to train the network for multiple epochs. The computational overhead of our method has two main parts: 1) Saving the intermediate features. Note the computation of deep features is a part of the original training process, we only need to slightly modify the network implemen-

Algorithm 1: Training loop with our performance estimation

Input: Initialized network $f(\mathbf{x}; \mathbf{w}_0^{fea}, \mathbf{w}_0^{cls})$,
training dataset $(\mathbf{X}^{train}, \mathbf{Y}^{train})$,
validation dataset $(\mathbf{X}^{val}, \mathbf{Y}^{val})$,
epoch budget N , window size K

Output: Improved performance estimation \bar{E}_N

- 1 **for** $i = 1$ to N **do**
- 2 Set $\mathbf{w}_i^{fea}, \mathbf{w}_i^{cls} = \mathbf{w}_{i-1}^{fea}, \mathbf{w}_{i-1}^{cls}$
- 3 **for** *Sampled batch* $\mathbf{X}_{i,j}^{train}, \mathbf{Y}_{i,j}^{train}$ in $\mathbf{X}^{train}, \mathbf{Y}^{train}$ **do**
- 4 Compute and save intermediate features $\mathbf{H}_{i,j}^{train} = g(\mathbf{X}_{i,j}^{train}; \mathbf{w}_i^{fea})$
- 5 Compute outputs $\mathbf{Z}_{i,j}^{train} = h(\mathbf{H}_{i,j}^{train}; \mathbf{w}_i^{cls})$
- 6 Optimize $\mathbf{w}_i^{fea}, \mathbf{w}_i^{cls}$ with one gradient step
- 7 Compute and save intermediate features $\mathbf{H}_i^{val} = g(\mathbf{X}^{val}; \mathbf{w}_i^{fea})$
- 8 Save checkpoint $\mathbf{w}_i^{fea}, \mathbf{w}_i^{cls}$
- 9 **for** $k = N - K + 1$ to N **do**
- 10 Initialize $\mathbf{v}_k^{cls} = \mathbf{w}_k^{cls}$
- 11 Optimize \mathbf{v}_k^{cls} with $\text{Loss}(h(\mathbf{H}_k^{train}; \mathbf{v}_k^{cls}), \mathbf{Y}^{train})$
- 12 Compute outputs $\mathbf{Z}_k^{val} = h(\mathbf{H}_k^{val}; \mathbf{v}_k^{cls})$
- 13 Average predicted distributions $\bar{\mathbf{Z}}_N^{val} = \text{Mean}(\mathbf{Z}_{N-K+1}^{val}, \dots, \mathbf{Z}_N^{val})$
- 14 Evaluate performance $\bar{E}_N = \text{Eval}(\bar{\mathbf{Z}}_N^{val}, \mathbf{Y}^{val})$
- 15 **return** \bar{E}_N

tation, let it return the intermediate results, and save them in the storage after each epoch. 2) Optimization and ensemble of the linear classifiers. As discussed above, the linear classifiers are computationally cheaper and easier to optimize than the feature extraction part. We also have good weight initialization from the saved checkpoints, so a small number of iterations are adequate for convergence. We also observe the memory usage increase is moderate. The image features that we need to save have reduced dimension compared with earlier layers. Therefore, the overhead only takes up a small fraction of the whole training process.

3.2 Combination with Search Algorithms

Our method is able to give accurate performance estimation at an early training stage, thus better assess the quality of a configuration with limited computational resources. In order to apply our method in a search task like NAS and HPO, we need to integrate our performance estimation into existing search algorithms for general purpose. In the experiments, we mainly focus on three algorithms: random search, HyperBand [22], and BOHB [12].

In random search, we generate a pool of configurations by random sampling from the search space, evaluate each configuration after training for some fixed budget, and pick the best one based on the latest performance. HyperBand is a bandit strategy which adaptively allocates training resources for configurations in the sample pool based on their current performance. BOHB can be considered as a variant of HyperBand, which includes a Bayesian optimization component to better generate the initial configuration pool. These search algorithms do not have assumptions about the search space, making them applicable to most NAS and HPO tasks. When combined with our method, these search algorithm make decisions based on the improved performance estimation instead of the latest performance acquired. In HyperBand and BOHB, this performance metric also determines the configurations which will be allocated with more budgets.

Chapter 4

Experiments

We first show that our method can efficiently estimate the performance of various CNN architecture families. Then we combine our method with search algorithms and demonstrate its effectiveness on neural architecture search (NAS) and hyperparameter optimization (HPO).

4.1 Performance Estimation for CNNs

We consider the following architecture families: VGG [30], ResNet [14], and MobileNetV2 [28] and use the CIFAR-10 and CIFAR-100 dataset [21]. We train each architecture for 200 epochs (see experimental details in Appendix A.2). We report the test accuracy every 10 epochs, which is the original performance evaluated at that epoch. We also apply our method for checkpoints during training every 10 epochs, which gives an improved estimation of the final performance. The results for comparison are shown in Figure 4.1. By utilizing the feature histories during network optimization, our method produces stable and accurate estimation of the final performance at an early training stage for all these CNN architectures.

In a resource-limited search setting, one may need to shrink the training budget and adjust the learning rate schedule accordingly, at the cost of reducing the final performance. We also investigate the example of ResNet-18 with fewer training epochs, as shown in Figure 4.2. Our method can still accurately approximate the final performance early in this setting.

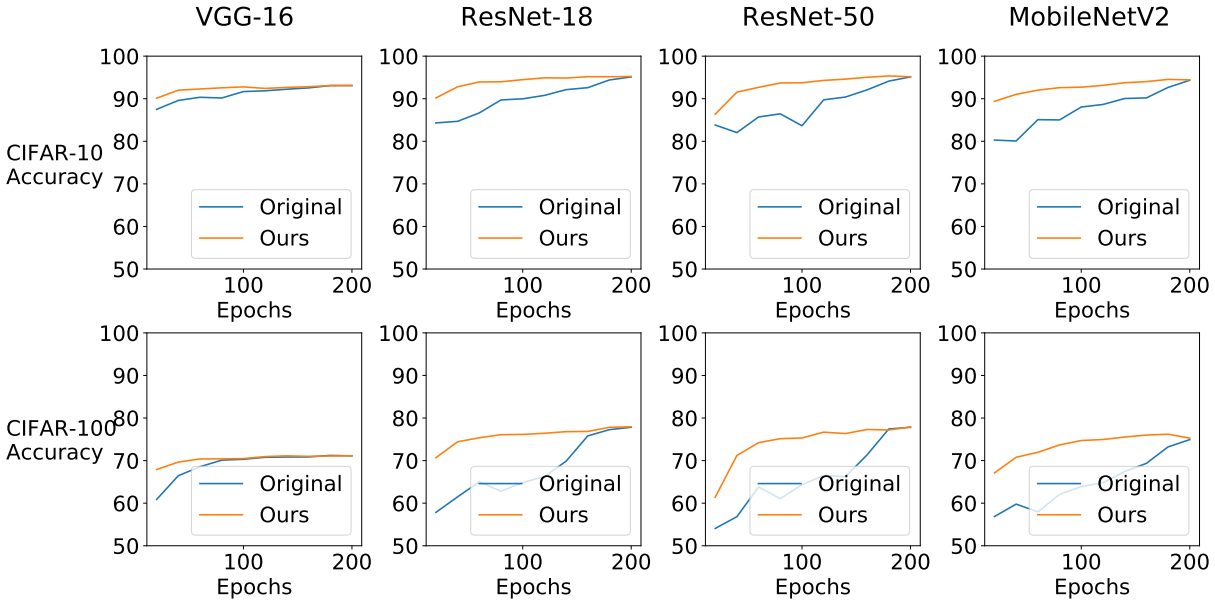


Figure 4.1: **Performance estimation for different architectures including VGG-16, ResNet-18, ResNet-50, and MobileNetV2 on CIFAR-10/CIFAR-100 dataset: Original** denotes the original test accuracy at each epoch. **Ours** denotes the performance of our proposed proxy classifier using the feature histories up to the given epoch. Our method can reach an accuracy closer to the final accuracy very early.

4.2 Neural Architecture Search

For NAS applications, we conduct experiments in two search spaces: NAS-Bench-201 [11] and DARTS [24].

4.2.1 NAS-Bench-201

NAS-Bench-201 [11] is a public benchmark for testing NAS algorithms. It defines a search space consisting of 15,625 architectures, and includes full training log of all the architectures on CIFAR-10, CIFAR-100 [21], and downsampled ImageNet [8]. In the following experiments, we mainly use the information about the “true performance” on CIFAR-10, which is defined as the average top-1 test accuracy of three independent runs of training the given architecture for 200 epochs.

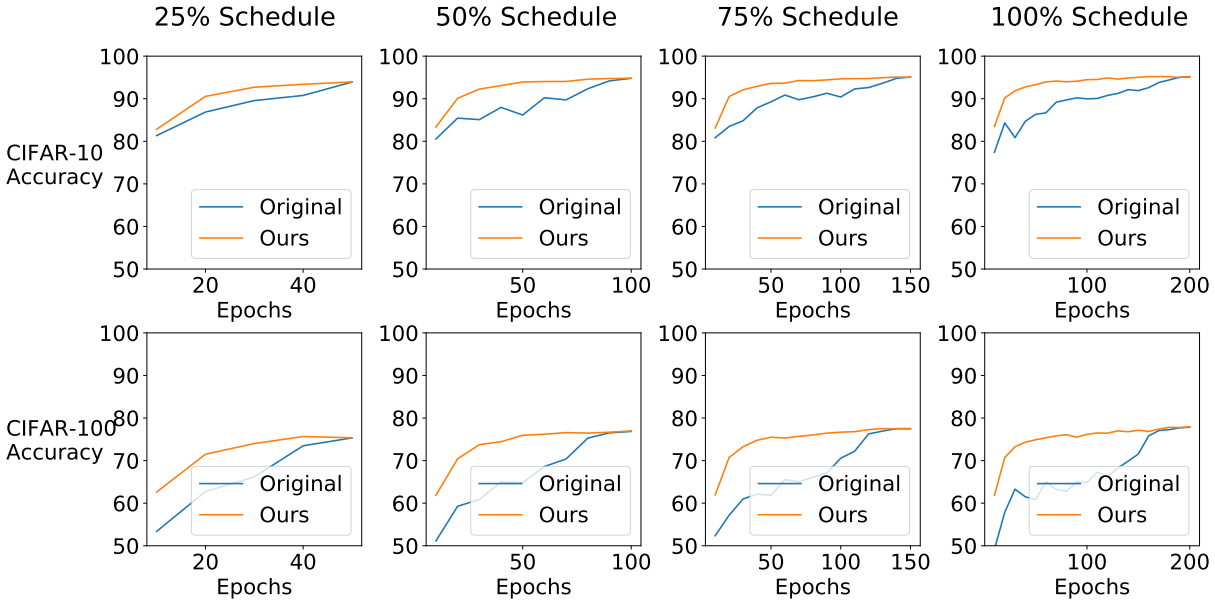


Figure 4.2: **Performance estimation for different training budgets from 25% (50 epochs) to 100% (200 epochs) of training ResNet-18 on CIFAR-10/CIFAR-100 dataset:** The linear learning rate schedule is adjusted accordingly.

First, we verify that our method is helpful for comparing different architecture configurations. To pick the best architecture among candidates, usually we need to train the architectures for some epochs and evaluate on the validation set to acquire a performance estimate, and the longer we train, the estimate is closer to the true performance. Here we randomly sample 100 architectures from the search space, train each for 200 epochs, evaluate both the original models and our proposed proxy classifier every 10 epochs, and compare their performance with the true performance (see experimental details in Appendix A.3.1). We also include another performance prediction baseline LCNet [19] for better comparison. LCNet trains a Bayesian neural network which extrapolates the learning curve using partially observed performance data. It can approach candidate models’ final performance by extrapolating the observed learning curve, but its representation power is limited for accurate predictions. Figure 4.3 shows that our method can acquire significantly more accurate estimates for true performance of NAS-Bench-201 architectures, compared to both baselines.

Besides the absolute numbers, it is important for a NAS algorithm to obtain the relative rank-

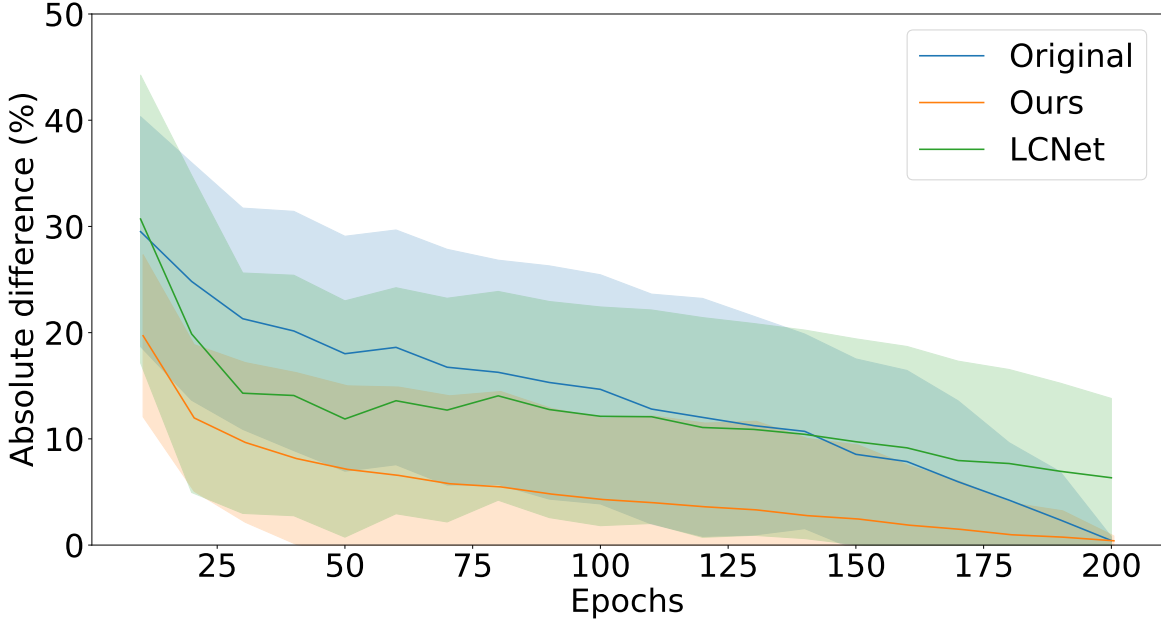


Figure 4.3: **Absolute difference between the estimated performance and the true performance:** **Original** denotes the original accuracy metrics at each epoch. **Ours** denotes the predicted performance by our method using the checkpoints up to the given epochs. **LCNet** denotes the extrapolated final performance given by LCNet. For fair comparison, we slightly shift **Ours** rightwards because of the computational overhead.

ing of architectures that is consistent with the true performance, so that the algorithm can return an architecture among the best ones in the search space. Given the true performance and estimated performance, we can measure the consistency of their relative ranking with the Kendall’s τ rank correlation coefficient [17]. The τ coefficient ranges in $[-1, 1]$, and when the two observations have perfectly matched relative rankings, τ reaches its maximum value 1. Figure 4.4 shows that the performance produced by our method is not only closer to the true performance for each individual architecture, but also keeps better relative ranking for distinguishing different architectures. Using this method, we can find better architectures with limited training budgets.

Next, we compare several search algorithms with and without our performance estimation, as described in Section 3.2 (see experimental details in Appendix A.3.2). The search budget is around 2 hours per run, and we quantify the performance of each search algorithm by the




Figure 4.4: **Relative ranking measured by Kendall’s τ coefficient between the estimated performance and the true performance.**

test accuracy of the searched architecture. The results are shown in Table 4.1. Our method can improve all three search algorithms in terms of the final searched architectures. Especially, the combination of HyperBand and our method can robustly find an architecture which has $< 0.7\%$ test accuracy gap from the global optimum. The global optimal architecture in NAS-Bench-201 search space has accuracy 94.37%.

4.2.2 DARTS

Differentiable ARchiTecture Search (DARTS) [24] is a recent NAS algorithm which relaxes the search space to be continuous, achieving search speed orders of magnitude faster than previous methods. Here we directly transfer the method that we use on NAS-Bench-201 to the DARTS search space, to demonstrate our performance estimation strategy is also helpful in another NAS setting.

We use the combination of HyperBand and our method, which performs the best on NAS-

Table 4.1: Search results for NAS-Bench-201 architectures on CIFAR-10.

| Method | Test Accuracy (%) |
|----------------------|------------------------------------|
| Random Search | 93.21 \pm 0.78 |
| Random Search + Ours | 93.45 \pm 0.27 |
| HyperBand | 93.36 \pm 0.30 |
| HyperBand + Ours | 93.77 \pm 0.09 |
| BOHB | 93.35 \pm 0.55 |
| BOHB + Ours | 93.45 \pm 0.37 |

Bench-201, and follow the practice of architecture evaluation in DARTS (see experimental details in Appendix A.4). The results are summarized in Table 4.2. We also list some more recent DARTS variants which also use the same search space and evaluation configuration, including P-DARTS [5], PC-DARTS [38], and UNAS [34]. Our method is able to find better architectures than DARTS with similar search cost. Our result is also close to more recent state-of-the-art NAS methods. It is notable that our method is the only one that does not make use of the gradient-based method, so we can easily parallelize the search process. In fact, we use 4 parallel workers in the experiment, reducing the wall-clock search time to 0.2 day per run. We expect to further reduce search time with even more parallel workers.

4.3 Hyperparameter Optimization

DARTS and its variants are NAS algorithms customized for their search space design. We are also interested in other tasks where the search space is non-differentiable and thus these NAS algorithms may not apply. Here we investigate two examples of hyperparameter optimization, to show our method is robust and versatile to a wider range of tasks.

Table 4.2: **Comparison with state-of-the-art NAS methods that uses the DARTS search space on CIFAR-10.**

| Method | Test Error (%) | Params (M) | Search Cost (GPU days / run) | Search Method |
|----------------------|-----------------|------------|------------------------------|---------------|
| DARTS (first order) | 3.00 ± 0.14 | 3.3 | 0.4 | Gradient |
| DARTS (second order) | 2.76 ± 0.09 | 3.3 | 1 | Gradient |
| P-DARTS | 2.50 | 3.4 | 0.3 | Gradient |
| PC-DARTS | 2.57 ± 0.07 | 3.6 | 0.1 | Gradient |
| UNAS | 2.53 | 3.3 | 4.3 | Gradient&RL |
| Ours | 2.64 ± 0.04 | 3.4 | $0.2 \times 4^*$ | HyperBand |

* We use 4 parallel workers. The wall-clock time is 0.2 day.

4.3.1 Data Augmentation

In this task, we consider the problem of optimizing hyperparameters in data augmentation for CNN training. We use the search space defined in RandAugment [7], which has only two discrete hyperparameters: N is the number of augmentation transformations to apply, and M is the magnitude for the transformations. The search space is significantly smaller than previous work such as AutoAugment [6], but still computationally expensive for the grid search that [7] does for each CNN architecture and dataset.

To compare HyperBand and BOHB with and without our performance estimation, we search the optimal RandAugment configuration for VGG-16, ResNet-18, and MobileNetV2 on a subset of CIFAR-10 (see experimental details in Appendix A.5). The results are summarized in Table 4.3. In most cases, our method consistently improves the baseline search algorithms. This task further demonstrates the effectiveness and versatility of our method.

Table 4.3: **Search results for RandAugment hyperparameter optimization on a subset of CIFAR-10.**

| Method | Test Accuracy (%) | | |
|------------------|------------------------------------|------------------------------------|------------------------------------|
| | VGG-16 | ResNet-18 | MobileNetV2 |
| HyperBand | 85.38 ± 0.24 | 85.83 ± 0.43 | 86.53 ± 0.23 |
| HyperBand + Ours | 85.60 ± 0.28 | 86.06 ± 0.64 | 86.62 ± 0.48 |
| BOHB | 85.35 ± 0.44 | 86.04 ± 1.16 | 86.92 ± 0.46 |
| BOHB + Ours | 85.61 ± 0.38 | 86.53 ± 0.64 | 86.80 ± 0.71 |

4.3.2 Scaling Factors for ResNet

EfficientNet [33] is a series of CNN architectures with state-of-the-art image classification performance, which is created by compound scaling up the depth, width, and resolution of a small base network. Inspired by EfficientNet, we can consider scaling up the ResNet-18 architecture to build stronger networks by multiplying the number of building blocks and channels at each stage with some scaling factors, under the constraint of FLOPS. We test HyperBand with and without our method, to search the optimal scaling factors for ResNet-18 on CIFAR-100 (see experimental details in Appendix A.6). Table 4.4 shows that our method can help HyperBand to find better scaling factors in this task.

We have also transferred the best-performing solutions searched on CIFAR-100 to ImageNet, and the result in Table 4.5 further demonstrates the improvement brought by our method.

Table 4.4: Search results for scaling factors of ResNet-18 on CIFAR-100.

| Architecture | Test Accuracy (%) | Params (M) | FLOPS (G) |
|----------------------------------|---------------------|--------------|-------------|
| ResNet-18 | 77.83 | 11.22 | 0.56 |
| ResNet-34 | 78.49 | 21.33 | 1.16 |
| 2× ResNet-18 by HyperBand | 78.40 ± 0.21 | 18.35 ± 0.42 | 1.04 ± 0.05 |
| 2× ResNet-18 by HyperBand + Ours | 78.75 ± 0.03 | 15.83 ± 1.75 | 1.02 ± 0.08 |
| 4× ResNet-18 by HyperBand | 79.00 ± 0.12 | 22.10 ± 3.36 | 1.57 ± 0.26 |
| 4× ResNet-18 by HyperBand + Ours | 79.40 ± 0.34 | 28.31 ± 6.23 | 1.93 ± 0.24 |

Table 4.5: ImageNet performance of scaled ResNet-18 searched on CIFAR-100.

| Architecture | Top-1 Accuracy (%) |
|-------------------------|--------------------|
| 2× ResNet-18 by HB | 71.46 |
| 2× ResNet-18 by HB+Ours | 71.77 |
| 4× ResNet-18 by HB | 73.01 |
| 4× ResNet-18 by HB+Ours | 73.14 |

Chapter 5

Conclusion

We propose a novel efficient model performance estimation method, which effectively use the saved feature histories during optimization to produce accurate approximation of the final performance. Our method is simple to implement and applicable to various tasks in NAS and HPO, and leads to improvement for general search algorithms, as demonstrated by our extensive experiments.

Our work reveals the informativeness of the evolution in deep features. By analyzing the dynamics of feature histories, one may discover signals that indicate the future directions of a neural network being trained. Therefore, for future work, we believe it would be possible to use our proposed performance estimation strategy to guide and further accelerate CNN training.

Appendix A

Experimental Details

A.1 Classifier Ensemble

The training procedure of the classifier ensemble has the same following configuration across all experiments. We use the the saved features from the most recent $K = 10$ epochs. We first obtain K optimal classifiers by training them on the saved features of images from the training set and their labels. The weights are initialized from the saved network weights of the corresponding epochs. We use the cross entropy loss and SGD optimizer with momentum 0.9 and weight decay 5×10^{-4} . The learning rate starts from 0.05 and decays following a linear schedule [23]. The batch size is 1024, and the number of epochs is 5. To build the ensemble, we collect the classification results of the optimized linear classifiers for the validation image features, and output the mean of the softmax-ed probability distribution, which is then used for evaluation and estimation. The whole process introduces little computational cost compared with the original training loop. For example, in NAS-Bench-201, we observe the overhead of our method is typically 0.5 epoch time for each architecture.

A.2 Performance Estimation for CNNs

In the first part, we train each CNN on CIFAR-10 and CIFAR-100 for 200 epochs with batch size 256. The initial learning rate is set to 0.1 for ResNet and MobileNetV2, 0.01 for VGG, and we

use a linearly decaying learning rate schedule as suggested by [23]. We use SGD optimizer with momentum 0.9 and weight decay 5×10^{-4} .

In the second part, we train ResNet-18 on CIFAR-100 for 50, 100, 150, 200 epochs, corresponding to 25%, 50%, 75%, 100% budgets. The learning rate still follows the linear schedule, which drops to zero at the end of each budget. Other settings are the same as the first part.

A.3 NAS-Bench-201

A.3.1 Performance Error and Relative Ranking

In each run of this experiment, we first randomly sample 100 architectures from NAS-Bench-201 [11], and train each for 200 epochs with batch size 256. The initial learning rate is set to 0.1 and we use the linear learning rate schedule. We use SGD optimizer with momentum 0.9 and weight decay 5×10^{-4} . We split the original CIFAR-10 training set into two subsets: 40000 images for training and 10000 images for validation.

For the performance error, we compute absolute difference between the current performance and the true performance, and for the relative ranking, we compute the Kendall’s τ coefficient [17] between the current performance and the true performance. The experiment is repeated 3 times.

A.3.2 Architecture Search

For random search, we randomly sample 64 architectures, and train each for 128 epochs with batch size 256. For HyperBand and BOHB, the training budget for each architecture ranges in $[1, 128]$ epochs, and the factor for increasing training budget and shrinking sample pool is set as $\eta = 2$. We use 4 parallel workers, each using an NVIDIA GeForce RTX 2080 Ti GPU. The total search time is about the same for the three algorithms, around 2 hours. During search, the initial learning rate is set to 0.1 and we use the linear learning rate schedule with the maximal budget set to 128. We use SGD optimizer with momentum 0.9 and weight decay 5×10^{-4} . we use a subset from the original CIFAR-10 training set with 8000 images for training, and 2000 images

for validation. After search, we calculate the regret as the difference in the true performance between the architecture given by each algorithm and the global optimal architecture, provided by the benchmark. Each experiment is repeated 5 times with different random seeds.

A.4 DARTS

During search, the search space relaxation is not applicable in our method. We still randomly sample from the discrete search space defined in DARTS [24], and select the best architectures using HyperBand and our performance estimation. We decompose the process of sampling architecture into sampling the operations and connections for each node in the cells from the uniform distribution of valid choices. At search time, we set the batch size to 96, initial number of channels to 32, number of cells to 12. We also use Cutout [9], path dropout, and auxiliary towers. The initial learning rate is 0.025 and we use the linear learning rate schedule with the maximal budget set to 128. We use SGD optimizer with momentum 0.9 and weight decay 3×10^{-4} . Other search settings are the same as Appendix A.3.2.

For architecture evaluation, we closely follow the setup of DARTS: We repeat the search process for 4 times and pick the best architecture for fully train and evaluation. An enlarged architecture with 36 initial channels is trained for 600 epochs with batch size 96, using Cutout, path dropout, and auxiliary towers. The only difference in architecture evaluation is we increase the number of cells from 20 to 22. The reason is that our searched cells have more parameter-free operations like skip connections and pooling layers compared with DARTS, so we increase the number of cells to match the number of parameters in DARTS.

A.5 Data Augmentation

To better show the influence of data augmentation, we only use a small subset of the original CIFAR-10 training set with 8000 images for training and 2000 images for validation. Besides RandAugment [7], default data augmentation for training images also includes random flips, pad-and-crop and Cutout, following RandAugment. Other search settings are the same as Ap-

pendix A.3.2. Each experiment is repeated 3 times with different random seeds.

A.6 Scaling Factors for ResNet

We use ResNet-18 [14] as our base model. The $\{C_2, C_3, C_4, C_5\}$ stages of ResNet-18 have $\{2, 2, 2, 2\}$ basic building blocks and $\{64, 128, 256, 512\}$ channels respectively. We search for 8 hyper-parameters $\{d_2, d_3, d_4, d_5, w_2, w_3, w_4, w_5\}$ for scaling up the depth and width of ResNet-18. The scaled network has $\{d_2 \times 2, d_3 \times 2, d_4 \times 2, d_5 \times 2\}$ blocks and $\{w_2 \times 64, w_3 \times 128, w_4 \times 256, w_5 \times 512\}$ channels for each stage. Meanwhile, we constrain the size of the scaled network by its FLOPS. For CIFAR-100, ResNet-18 has 0.56 GFLOPS, so $2 \times$ ResNet-18 should have no larger than 1.12 GFLOPS, and $4 \times$ ResNet-18 should have no larger than 2.24 GFLOPS. Other search settings are the same as Appendix A.3.2. Each experiment is repeated 3 times with different random seeds.

Bibliography

- [1] Bowen Baker, Otkrist Gupta, Ramesh Raskar, and Nikhil Naik. Accelerating neural architecture search using performance prediction. *arXiv preprint arXiv:1705.10823*, 2017. 2.1
- [2] Gabriel M. Bender, Pieter jan Kindermans, Barret Zoph, Vijay Vasudevan, and Quoc Le. Understanding and simplifying one-shot architecture search. In *ICML*, 2018. 2.1
- [3] James S Bergstra, Rémi Bardenet, Yoshua Bengio, and Balázs Kégl. Algorithms for hyperparameter optimization. In *Advances in neural information processing systems*, pages 2546–2554, 2011. 2.2
- [4] Tianqi Chen, Ian Goodfellow, and Jonathon Shlens. Net2net: Accelerating learning via knowledge transfer. *arXiv preprint arXiv:1511.05641*, 2015. 2.1
- [5] Xin Chen, Lingxi Xie, Jun Wu, and Qi Tian. Progressive differentiable architecture search: Bridging the depth gap between search and evaluation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1294–1303, 2019. 2.3, 4.2.2
- [6] Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. Autoaugment: Learning augmentation policies from data. *arXiv preprint arXiv:1805.09501*, 2018. 4.3.1
- [7] Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 702–703, 2020. 4.3.1, A.5
- [8] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-

- scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 4.2.1
- [9] Terrance DeVries and Graham W Taylor. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*, 2017. A.4
- [10] Tobias Domhan, Jost Tobias Springenberg, and Frank Hutter. Speeding up automatic hyperparameter optimization of deep neural networks by extrapolation of learning curves. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*, 2015. 2.1
- [11] Xuanyi Dong and Yi Yang. Nas-bench-102: Extending the scope of reproducible neural architecture search. *arXiv preprint arXiv:2001.00326*, 2020. 1, 2.3, 4.2, 4.2.1, A.3.1
- [12] Stefan Falkner, Aaron Klein, and Frank Hutter. Bohb: Robust and efficient hyperparameter optimization at scale. *arXiv preprint arXiv:1807.01774*, 2018. 2.2, 3.2
- [13] Timur Garipov, Pavel Izmailov, Dmitrii Podoprikin, Dmitry P Vetrov, and Andrew G Wilson. Loss surfaces, mode connectivity, and fast ensembling of dnns. In *Advances in Neural Information Processing Systems*, pages 8789–8798, 2018. 2.4
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 1, 3.1.1, 4.1, A.6
- [15] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. 3.1.1
- [16] Gao Huang, Yixuan Li, Geoff Pleiss, Zhuang Liu, John E Hopcroft, and Kilian Q Weinberger. Snapshot ensembles: Train 1, get m for free. *arXiv preprint arXiv:1704.00109*, 2017. 2.4
- [17] Maurice G Kendall. A new measure of rank correlation. *Biometrika*, 30(1/2):81–93, 1938. 4.2.1, A.3.1
- [18] Aaron Klein and Frank Hutter. Tabular benchmarks for joint architecture and hyperparameter optimization. *arXiv preprint arXiv:1905.04970*, 2019. 2.3

- [19] Aaron Klein, Stefan Falkner, Jost Tobias Springenberg, and Frank Hutter. Learning curve prediction with bayesian neural networks. 2016. 2.1, 4.2.1
- [20] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012. 3.1.1
- [21] Alex Krizhevsky et al. Learning multiple layers of features from tiny images. 2009. 4.1, 4.2.1
- [22] Lisha Li, Kevin Jamieson, Giulia DeSalvo, Afshin Rostamizadeh, and Ameet Talwalkar. Hyperband: A novel bandit-based approach to hyperparameter optimization. *The Journal of Machine Learning Research*, 18(1):6765–6816, 2017. 2.2, 3.2
- [23] Mengtian Li, Ersin Yumer, and Deva Ramanan. Budgeted training: Rethinking deep neural network training under resource constraints. *arXiv preprint arXiv:1905.04753*, 2019. 2.1, A.1, A.2
- [24] Hanxiao Liu, Karen Simonyan, and Yiming Yang. Darts: Differentiable architecture search. *arXiv preprint arXiv:1806.09055*, 2018. 1, 2.3, 4.2, 4.2.2, A.4
- [25] Hieu Pham, Melody Guan, Barret Zoph, Quoc Le, and Jeff Dean. Efficient neural architecture search via parameters sharing. In *ICML*, 2018. 2.1
- [26] Esteban Real, Sherry Moore, Andrew Selle, Saurabh Saxena, Yutaka Leon Suematsu, Jie Tan, Quoc Le, and Alex Kurakin. Large-scale evolution of image classifiers. *arXiv preprint arXiv:1703.01041*, 2017. 2.3
- [27] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015. 3.1.1
- [28] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4510–4520, 2018. 1, 4.1
- [29] Han Shi, Renjie Pi, Hang Xu, Zhenguo Li, James Kwok, and Tong Zhang. Bridging the

- gap between sample-based and one-shot neural architecture search with bonas. *Advances in Neural Information Processing Systems*, 33, 2020. 2.1
- [30] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 1, 3.1.1, 4.1
- [31] Jasper Snoek, Hugo Larochelle, and Ryan P Adams. Practical bayesian optimization of machine learning algorithms. In *Advances in neural information processing systems*, pages 2951–2959, 2012. 2.2
- [32] Jasper Snoek, Oren Rippel, Kevin Swersky, Ryan Kiros, Nadathur Satish, Narayanan Sundaram, Mostofa Patwary, Mr Prabhat, and Ryan Adams. Scalable bayesian optimization using deep neural networks. In *International conference on machine learning*, pages 2171–2180, 2015. 2.2
- [33] Mingxing Tan and Quoc V Le. Efficientnet: Rethinking model scaling for convolutional neural networks. *arXiv preprint arXiv:1905.11946*, 2019. 2.3, 4.3.2
- [34] Arash Vahdat, Arun Mallya, Ming-Yu Liu, and Jan Kautz. Unas: Differentiable architecture search meets reinforcement learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11266–11275, 2020. 4.2.2
- [35] Jingjing Xie, Bing Xu, and Zhang Chuang. Horizontal and vertical ensemble with deep representation for classification. *arXiv preprint arXiv:1306.2759*, 2013. 2.4
- [36] Lingxi Xie and Alan Yuille. Genetic cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 1379–1388, 2017. 2.3
- [37] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1492–1500, 2017. 3.1.1
- [38] Yuhui Xu, Lingxi Xie, Xiaopeng Zhang, Xin Chen, Guo-Jun Qi, Qi Tian, and Hongkai Xiong. Pc-darts: Partial channel connections for memory-efficient differentiable architecture search. *arXiv preprint arXiv:1907.05737*, 2019. 2.3, 4.2.2
- [39] Chris Ying, Aaron Klein, Eric Christiansen, Esteban Real, Kevin Murphy, and Frank Hut-

- ter. Nas-bench-101: Towards reproducible neural architecture search. In *International Conference on Machine Learning*, pages 7105–7114, 2019. 2.3
- [40] Xiawu Zheng, Rongrong Ji, Qiang Wang, Qixiang Ye, Zhenguo Li, Yonghong Tian, and Qi Tian. Rethinking performance estimation in neural architecture search. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11356–11365, 2020. 2.1
- [41] Barret Zoph and Quoc V Le. Neural architecture search with reinforcement learning. *arXiv preprint arXiv:1611.01578*, 2016. 2.3
- [42] Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V Le. Learning transferable architectures for scalable image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8697–8710, 2018. 2.3