

NASOA: Towards Faster Task-oriented Online Fine-tuning with a Zoo of Models

Hang Xu
Huawei Noah's Ark Lab

Ning Kang
Huawei Noah's Ark Lab

Gengwei Zhang
Sun Yat-sen University

Chuanlong Xie
Huawei Noah's Ark Lab

Xiaodan Liang*
Sun Yat-sen University

Zhenguo Li
Huawei Noah's Ark Lab

Abstract

Fine-tuning from pre-trained ImageNet models has been a simple, effective, and popular approach for various computer vision tasks. The common practice of fine-tuning is to adopt a default hyperparameter setting with a fixed pre-trained model, while both of them are not optimized for specific tasks and time constraints. Moreover, in cloud computing or GPU clusters where the tasks arrive sequentially in a stream, faster online fine-tuning is a more desired and realistic strategy for saving money, energy consumption, and CO2 emission. In this paper, we propose a joint Neural Architecture Search and Online Adaption framework named NASOA towards a faster task-oriented fine-tuning upon the request of users. Specifically, NASOA first adopts an off-line NAS to identify a group of training-efficient networks to form a pretrained model zoo. We propose a novel joint block and macro level search space to enable a flexible and efficient search. Then, by estimating fine-tuning performance via an adaptive model by accumulating experience from the past tasks, an online schedule generator is proposed to pick up the most suitable model and generate a personalized training regime with respect to each desired task in a one-shot fashion. The resulting model zoo¹ is more training efficient than SOTA models, e.g. 6x faster than RegNetY-16GF, and 1.7x faster than EfficientNetB3. Experiments on multiple datasets also show that NASOA achieves much better fine-tuning results, i.e. improving around 2.1% accuracy than the best performance in RegNet series under various constraints and tasks; 40x faster compared to the BOHB.

1. Introduction

Fine-tuning using pre-trained models becomes the default standard in the field of computer vision because of

its impressive results on various downstream tasks such as fine-grained image classification [36, 50], object detection [19, 23, 54] and segmentation [8, 30]. [25, 19] verified that fine-tuning pre-trained networks outperform training from scratch. It can further help to avoid over-fitting [10] as well as reduce training time significantly [19]. Due to those merits, many cloud computing and AutoML pipelines provide fine-tuning services for an online stream of upcoming users with new data, different tasks and time limits. In order to save the user's time, money, energy consumption, or even CO2 emission, an efficient online automated fine-tuning framework is practically useful and in great demand. Thus, we propose to explore faster online fine-tuning.

The conventional practice of fine-tuning is to adopt a set of predefined hyperparameters for training a predefined model [27]. It has three drawbacks in the current online setting: 1) The design of the backbone model is not optimized for the upcoming fine-tuning task and the selection of the backbone model is not data-specific. 2) A default setting of hyperparameters may not be optimal across tasks and the training settings may not meet the time constraints provided by users. 3) With the incoming tasks, the regular diagram is not suitable for this online setting since it cannot memorize and accumulate experience from the past fine-tuning tasks. Thus, we propose to decouple our faster fine-tuning problem into two parts: finding efficient fine-tuning networks and generating optimal fine-tuning schedules pertinent to specific time constraints in an online learning fashion.

Recently, Neural Architecture Search (NAS) algorithms demonstrate promising results on discovering top-accuracy architectures, which surpass the performance of hand-crafted networks and saves human's efforts [58, 31, 32, 38, 49, 40, 48, 56] as well as studying NAS across tasks and datasets [9, 13]. However, those NAS works usually focus on inference time/FLOPS optimization and their search space is not flexible enough which cannot guarantee the optimality for fast fine-tuning. In contrast, we resort to devel-

* Corresponding Author: Xiaodan Liang xdliang328@gmail.com

¹The efficient training model zoo (ET-NAS) has been released at: <https://github.com/NAS-OA/NASOA>

oping a NAS scheme with a novel flexible search space for fast fine-tuning. On the other hand, hyperparameter optimization (HPO) methods such as grid search [3], Bayesian optimization (BO) [46, 34], and BOHB [14] are used in deep learning and achieve good performance. However, those search-based methods are computationally expensive and require iterative “trial and error”, which violate our goal for faster adaptation time.

In this work, we propose a novel Neural Architecture Search and Online Adaption framework named NASOA. First, we conduct an offline NAS for generating an efficient fine-tuning model zoo. We design a novel block-level and macro-structure search space to allow a flexible choice of the networks. Once the efficient training model zoo is created offline NAS by Pareto optimal models, the online user can enjoy the benefit of those efficient training networks without any marginal cost. We then propose an online learning algorithm with an adaptive predictor to modeling the relation between different hyperparameter, model, dataset meta-info and the final fine-tuning performance. The final training schedule is generated directly from selecting the fine-tuning regime with the best predicted performance. Benefiting from the experience accumulation via online learning, the diversity of the data and the increasing results can further continuously improve our regime generator. Our method behaves in a one-shot fashion and doesn’t involve additional searching cost as HPO, endowing the capability of providing various training regimes under different time constraints. We also theoretically prove the convergence of the optimality of our proposed online model.

Extensive experiments are conducted on multiple widely used fine-tuning datasets. The searched model zoo ET-NAS is more training efficient than SOTA ImageNet models, e.g. 5x training faster than RegNetY-16GF, and 1.7x faster than EfficientNetB3. Moreover, by using the whole NASOA, our online algorithm achieves superior fine-tuning results in terms of both accuracy and fine-tuning speed, i.e. improving around 2.1% accuracy than the best performance in RegNet series under various tasks; saving 40x computational cost comparing to the BOHB method.

Our contributions are summarized as follows:

- We make the first effort to propose a faster fine-tuning pipeline that seamlessly combines the training-efficient NAS and online adaption algorithm. Our NASOA can effectively generate a personalized fine-tuning schedule of each desired task via an adaptive model for accumulating experience from the past tasks.
- The proposed novel joint block/macro level search space enables a flexible and efficient search. The resulting model zoo ET-NAS is more training efficient than very strong ImageNet SOTA models e.g. EfficientNet, RegNet. All the ET-NAS models have been released to help

the community skipping the computation-heavy NAS stage and directly enjoy the benefit of NASOA.

- The whole NASOA pipeline achieves much better fine-tuning results in terms of both accuracy and fine-tuning efficiency than current fine-tuning best practice and HPO method, e.g. , 40x faster compared to the BOHB method.

2. Related Work

Neural Architecture Search (NAS). The goal of NAS is to automatically optimize network architecture and release human effort from this handcraft network architecture engineering. Most previous works [32, 7, 31, 47, 53, 22, 51] aim at searching for CNN architectures with better inference and fewer FLOPS. [2, 5, 57] apply reinforcement learning to train an RNN controller to generate a cell architecture. [32, 53, 7] try to search a cell structure by weight-sharing and differentiable optimization. [48] use a grid search for an efficient network by altering the depth/width of the network with a fixed block structure. On the contrary, our NAS focuses creating an efficient training model zoo for fast fine-tuning. Moreover, the existing search space design cannot meet the purpose of our search.

Generating Hyperparameters for Fine-tuning. HPO methods such as Bayesian optimization (BO) [46, 34], BOHB [14] achieves very promising result but require a lot of computational resources which is contradictory to our original objective of efficient fine-tuning. On the other hand, limited works discuss the model selection and HPO for fine-tuning. [25] finds that ImageNet accuracy and fine-tuning accuracy of different models are highly correlated. [27, 1] suggest that the optimal hyperparameters and model for fine-tuning should be both dataset dependent and domain similarity dependent [10]. HyperStar [35] is a concurrent HPO work demonstrating that a performance predictor can effectively generate good hyper-parameters for a single model. However, those works don’t give an explicit solution about how to perform fine-tuning in a more practical online scenario. In this work, we take the advantage of online learning [21, 43] to build a schedule generator, which allows us to memorize the past training history and provide up-and-coming training regimes for new coming tasks on the fly. Besides, we introduce the NAS model zoo to further push up the speed and performance.

3. The Proposed Approach

The goal of this paper is to develop an online fine-tuning pipeline to facilitate a fast continuous cross-task model adaption. By the preliminary experiments in Section 4.1, we confirm that the model architectures and hyperparameters such as the learning rate and frozen stages will greatly influence the accuracy and speed of the fine-tuning program. Thus, our NASOA includes two parts as shown in the Figure 1: 1) Searching a group of neural architectures

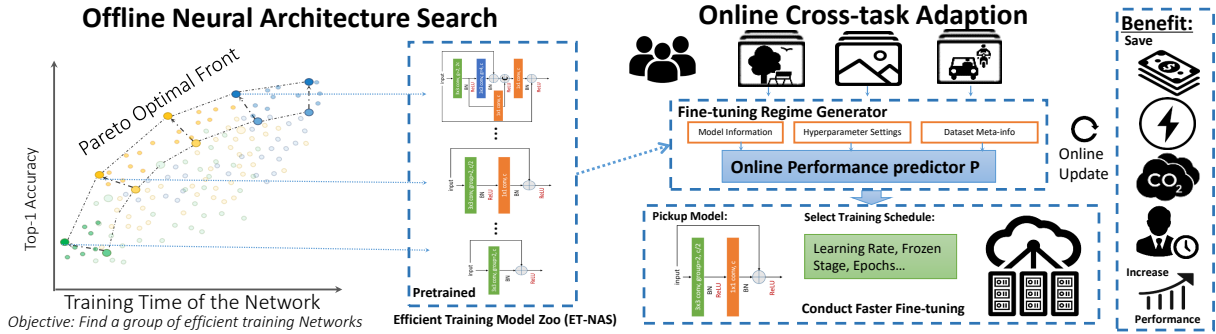


Figure 1. Overview of our NASOA. Our faster task-oriented online fine-tuning system has two parts: a) Offline NAS to generate an efficient training model zoo with good accuracy and training speed; b) An online fine-tuning regime generator to perform a task-specific fine-tuning with a suitable model under user’s time constraint.

with good accuracy and fast training speed to create a pre-trained model zoo; 2) Designing an online task-oriented algorithm to generate an efficient fine-tuning regime with the most suitable model under user’s time constraint.

3.1. Creating an Efficient Training Model Zoo (ET-NAS) by NAS

The commonly used hand-craft backbones for fine-tuning including MobileNet [44], ResNet [20], and ResNeXt [52]. Recently, some state-of-the-art backbone series such as RegNet [39], and EfficientNet [49] are developed by automated algorithms for higher accuracy and faster inference speed. However, the objective of our NAS is to find a group of models with good model generalization ability and training speed. Suggested by [25], the model fine-tuning accuracy (model generalization ability) has a strong correlation between ImageNet accuracy ($r = 0.96$). Meanwhile, the training speed can be measured by the step time of each training iteration. Thus, our NAS can be formulated by a multi-objective optimization problem (MOOP) on the search space S given by:

$$\max_{\mathcal{A} \in S} (\text{acc}(\mathcal{A}), -T_s(\mathcal{A})) \text{ subject to } T_s(\mathcal{A}) \leq T_m \quad (1)$$

where \mathcal{A} is the architecture, $\text{acc}(\cdot)$ is the Top-1 accuracy on ImageNet, $T_s(\cdot)$ is the average step time of one iteration, and T_m is the maximum step time allowed. The step time is defined to be the total time of one iteration, including forward/backward propagation, and parameter update.

Search Space Design is extremely important [39]. As shown in Figure 2, we propose a novel flexible joint block-level and macro-level search space to enable simple to complex block design and fine adjustment of the computation allocation on each stage. Unlike existing topological cell-level search space such as DARTS [32], AmoebaNet[40], and NASBench101[12], ours is more compact and avoids redundant skip-connections which have great memory access cost (MAC). Our block-level search space is more flexible to adjust the width, depth (for each stage), when to down-sample/raise the channels. In contrast, EfficientNet

only scales up/down the total width and depth by a fixed allocation ratio, and RegNet cannot change the number/type of operations in each block.

Block-level Search Space. We consider a search space based on 1-3 successive nodes of 5 different operations. Three skip connections with one fixed residual connection are searched. Element-wise add or channel-wise concat is chosen to combine the features for the skip-connections. For each selected operation, we also search for the ratio of changing channel size: $\times 0.25, \times 0.5, \times 1, \times 2, \times 4$. Note that it can cover many popular block designs such as Bottleneck [20], ResNeXt [52] and MB block [44]. It consists of 5.4×10^6 unique blocks.

Macro-level Search Space. Allocation computation over different stages is crucial for a backbone [29]. Early-stage feature maps in one backbone are larger which captures texture details, while late-stage feature maps are smaller which are more discriminative [28]. Therefore, for macro-level search space, we design a flexible search space to find the optimal channel size (width), depth (total number of blocks), when to down-sample, and when to raise the channels. Our macro-level structure consists of 4 flexible stages. The spatial size of the stages is gradually down-sampled with factor 2. In each stage, we stack a number of block architectures. The positions of the doubling channel block are also flexible. This search space consists of 1.5×10^7 unique architectures.

Multi-objective Searching Algorithm. For MOOP in Eq 1, we define architecture \mathcal{A}_1 dominates \mathcal{A}_2 if (i) \mathcal{A}_1 is no worse than \mathcal{A}_2 in all objectives; (ii) \mathcal{A}_1 is strictly better than \mathcal{A}_2 in at least one objective. \mathcal{A}^* is *Pareto optimal* if there is no other \mathcal{A} that dominate \mathcal{A}^* . The set of all *Pareto optimal* architectures constitutes the *Pareto front*. To solve this MOOP problem, we modify a well-known method named Elitist Non-Dominated sorting genetic algorithm (NSGA-II) [11] to optimize the *Pareto front* \mathcal{P}_f . The main idea of NSGA-II is to rank the sampled architectures by non-dominated sorting and preserve a group of elite architectures. Then a group of new architectures is sampled and




Figure 2. Our joint block/macro-level search space to find efficient training networks. Our block-level search space covers many popular designs such as ResNet, ResNext, MobileNet Block. Our macro-level search space allows small adjustment of the network in each stage thus the resulting models are more flexible and efficient.

trained by mutation of the current elite architectures on the \mathcal{P}_f . The algorithm can be paralleled on multiple computation nodes and lift the \mathcal{P}_f simultaneously. We modify the NSGA-II algorithm to become a NAS algorithm: a) To enable parallel searching on N computational nodes, we modify the non-dominated-sort method to generate exactly N mutated models for each generation, instead of a variable size as the original NSGA-II does. b) We define a group of mutation operations for our block/macro search space for NSGA-II to change the network structure dynamically. c) We add a parent computation node to measure the selected architecture’s training speed and generate the \mathcal{P}_f .

Efficient Training Model Zoo Z_{oo} (ET-NAS). By the proposed NAS method, we then create an efficient-training model zoo Z_{oo} named **ET-NAS** which consists of K Pareto optimal models \mathcal{A}_i^* on \mathcal{P}_f . \mathcal{A}_i^* are pretrained by ImageNet. Details of the search space, its encoding, NSGA-II algorithm and \mathcal{A}_i^* architectures can be found in Appendix.

3.2. Online Task-oriented Fine-tuning Schedule Generation

With the help of efficient training Z_{oo} , the marginal computational cost of each user is minimized while they can enjoy the benefit of NAS. We then need to decide a suitable fine-tuning schedule upon the user’s upcoming tasks. Given user’s dataset D and fine-tuning time constraint T_l , an online regime generator $G(., .)$ is desired:

$$[\text{Regime}_{\text{FT}}, \mathcal{A}_i^*] = G(D, T_l), \quad (2)$$

such that $\text{Acc}(\mathcal{A}_i^{\text{FineTune}}, D_{\text{val}})$ is maximized,

where the $\text{Regime}_{\text{FT}}$ includes all the hyperparameters required, i.e., lr schedule, total training steps, and frozen stages. $G(., .)$ also needs to pick up the most suitable pre-trained model \mathcal{A}_i^* from Z_{oo} . Note that existing search-based HPO methods require huge computational resources and cannot fit in our online one-shot training scenario. Instead, we first propose an online learning predictor Acc_P to model the accuracy on the validation set $\text{Acc}(\mathcal{A}_i^{\text{FT}}, D_{\text{val}})$ by the meta-data information. Then we can use the predictor to construct $G(., .)$ to generate an optimal hyperparameter setting and model.

3.2.1 Online Learning for Modeling $\text{Acc}(\mathcal{A}_i^{\text{FT}}, D_{\text{val}})$

Recently, [27] suggest that the optimal hyperparameters for fine-tuning are highly related to some data statistics such as domain similarity to the ImageNet. Thus, we hypothesis that we can model the final accuracy by a group of predictors, e.g., model information, meta-data description, data statistics $\text{stat}(D)$, domain similarity, and hyperparameters. We list the variables we considered to predict the accuracy result as follows:

Model \mathcal{A}_i^* name (one-hot dummy variable)	
Domain Similarity to ImageNet (EMD) [10]	
Average #. images per class	Std #. images per class
ImageNet Acc. of the \mathcal{A}_i^*	
#.Classes	Number of Iteration
Learning Rate	Frozen Stages

Those variables can be easily calculated ahead of the fine-tuning. One can prepare offline training data by fine-tuning different kinds of dataset and collect the accuracy correspondingly and apply a Multi-layer Perceptron Regression (MLP) offline on it. However, online learning should be a more realistic setting for our problem. In cloud computing service or a GPU cluster, a sequence of fine-tuning requests with different data will arrive from time to time. The predictive model can be further improved by increasing the diversity of the data and the requests over time.

Using a fixed depth of MLP model in the online setting may be problematic. Shallow networks maybe more preferred for small number of instances, while deeper model can achieve better performance when the sample size becomes larger. Inspired by [43], we use an adaptive MLP regression to automatically adapt its model capacity from simple to complex over time. Given the input variables, the prediction of the accuracy is given by:

$$\text{Acc}_P(\mathcal{A}_i^*, \text{Regime}_{\text{FT}}, \text{stat}(D)) = \sum_{l=1}^L \alpha_l f_l, \quad (3)$$

$f_l = h_l W_l$, $h_l = \text{RELU}(\Phi_l h_{l-1})$, $h_0 = [\mathcal{A}_i^*, \text{Regime}_{\text{FT}}, \text{stat}(D)]$. where $l = 0, \dots, L$. The predicted accuracy is a weighted sum of the output f_l of each intermediate fully-connected layer h_l . The W_l and Φ_l are the learnable weights of each fully-connected layer. The α_l is a weight vector assigning the importance to each layer and $\|\alpha\| = 1$. Thus the predictor Acc_P can automatically adapt its model capacity from

simple to complex along with incoming tasks. The learnable weight α_l controls the importance of each intermediate layer and the final predicted accuracy is a weighted sum of f_l of them. The network can be updated by a Hedge Backpropagation [16] in which α_l is updated based on the loss suffered by this layer l as follows:

$$\alpha'_l \leftarrow \alpha_l \beta^{\mathcal{L}(f_l, Acc_{gt})}, W'_l \leftarrow W_l - \eta \alpha_l \nabla_{W_l} \mathcal{L}(f_l, Acc_{gt}) \quad (4)$$

$$\Phi'_l \leftarrow \Phi'_l - \eta \sum_{j=1}^L \alpha_j \nabla_{W_l} \mathcal{L}(f_j, Acc_{gt}), \alpha''_l \leftarrow \frac{\alpha'_l}{\sum \alpha'_l}$$

where $\beta \in (0, 1)$ is the discount rate, the weight α'_l are re-normalized such that $\|\alpha\| = 1$, and η is the learning rate. Thus, during the online update, the model can choose an appropriate depth by α_l based on the performance of each output at that depth. By utilizing the online cumulative results, our generator gains experience that helps future prediction.

Generating Task-oriented Fine-tuning Schedule. Our schedule generator G then can make use of the performance predictor to find the best training regime $G(D, T_l)$:

$$\arg \max_{\mathcal{A} \in Z_{oo}, Regime_{FT} \in S_{FT}} Acc_P(\mathcal{A}, Regime_{FT}, \text{stat}(D)).$$

Once the time constraint T_l is provided, the max number of iterations for different \mathcal{A}_i^* can be calculated by an offline step-time lookup table for Z_{oo} . The corresponding meta-data variables can be then calculated for the incoming task. The optimal selection of model and hyperparameters is obtained by ranking the predicted accuracy of all possible grid combinations. Details can be found in the Appendix.

Theoretical Analysis. Let the α and \mathcal{L} denote as

$$\alpha = (\alpha_1, \alpha_2, \dots, \alpha_L)^T, \mathcal{L} = (\mathcal{L}_1, \mathcal{L}_2, \dots, \mathcal{L}_L)^T$$

where $\mathcal{L}_l = \mathcal{L}(f_l, Acc_{gt})$ for $l = 1, 2, \dots, L$. At time $0 \leq t \leq T$, we denote α and \mathcal{L} as $\alpha^{(t)}$ and $\mathcal{L}^{(t)}$ respectively.

Theorem 1. *Suppose the number of layers L is a fixed integer, the training time T is sufficiently large and the loss function $\mathcal{L}(f_l, Acc_{gt})$ is bounded in $[0, 1]$. The sequence of the weight vectors: $\{\alpha^{(1)}, \alpha^{(2)}, \dots, \alpha^{(T)}\}$, is learned by the Hedge Backpropagation in (4). The initialized weight vector $\alpha^{(1)}$ is the uniform discrete distribution $\alpha^{(1)} = (\frac{1}{L}, \frac{1}{L}, \dots, \frac{1}{L})$. The discount rate β is fixed during the training procedure and is taken to be $\sqrt{T}/(\sqrt{T} + C)$ given T , where C is a fixed constant. Then the average regret of the online learning algorithm for modelling $Acc(\mathcal{A}_i^{FT}, D_{val})$ satisfies*

$$\frac{1}{T} \sum_{t=1}^T (\alpha^{(t)})^T \mathcal{L}^{(t)} - \min_{\alpha} \frac{1}{T} \sum_{t=1}^T \alpha^T \mathcal{L}^{(t)} \leq O\left(\frac{1}{\sqrt{T}}\right).$$

Proof. The detailed proof is in the Appendix. \square

Remark. This theorem shows that the empirical average regret of the learned sequence $\{\alpha^{(t)}, t = 1, \dots, T\}$ converges

to the optimal achievable average regret as T tends to infinity. For any given α' such that

$$\sum_{t=1}^T (\alpha')^T \mathcal{L}^{(t)} - \min_{\alpha} \sum_{t=1}^T \alpha^T \mathcal{L}^{(t)} > 0$$

the learnt weight vector $\alpha^{(t)}$ finally outperforms α' as long as the training time T is sufficiently large. Obviously, α' can be any one-hot vector. This implies our adaptive learning is **better than the regression with a fixed-depth neural network**. Hence, after accumulating enough experience, **the online learning procedure finds out a solution that approaches the optimality**. The learned weight vector $\alpha^{(t)}$ can capture the optimal model capacity that fully employs the power of depth to learn complex patterns and also guarantees the faster convergence rate of a shallow model.

4. Experimental Results

4.1. Preliminary Experiments

We conduct a complete preliminary experiment to justify our motivation and model settings. Details can be found in the Appendix. According to our experiments, we find that for an efficient fine-tuning, the model matters most. **The suitable model should be selected according to the task and time constraints**. Thus constructing a model zoo with various sizes of training-efficient models and picking up suitable models should be a good solution for faster fine-tuning. We also verify some existing conclusions: Fine-tuning performs better than training from scratch [25] so that our topic is very important for efficient GPU training; Learning rate and frozen stage are crucial for fine-tuning [18], which needs careful adjustment.

4.2. Offline NAS and Model Zoo Results

During the NAS, we directly search on the ImageNet dataset[42]. We first search for a group of efficient block structure, then use those block candidates to conduct the macro-level search. We use a short training setting to evaluate each architecture. It takes about 1 hour on average for evaluating one architecture for the block-level search and 6 hours for the macro-level search. Paralleled on GPUs, it takes about one week on a 64-GPU cluster to conduct the whole search (5K+1K arch). Implementation details and intermediate results can be found in the Appendix.

Faster Fine-tuning Model Zoo (ET-NAS). After identifying the \mathcal{A}_i^* from our search, we fully train those models on ImageNet following common practice. Note that all the models including ET-NAS-L can be easily pretrained on a regular 8-card GPU node since our model is training-efficient. We have released our models for the public to reproduce our results from scratch and let the public to save their energy/CO2/cost. Due to the length of the paper, we put the detailed encoding and architectures of the

DataSets	#.Class	Task	#.Train	#.Test	DataSets	#.Class	Task	#.Train	#.Test
Flowers102 [36]	102	Fine-Grained	6K	2K	Stanford-Car[26]	196	Fine-Grained	8K	8K
CUB-Birds [50]	200	Fine-Grained	10K	2K	MIT67[37]	67	Scene cls.	5K	1K
Caltech101 [15]	101	General	8K	1K	Food101[4]	101	Fine-Grained	75K	25K
Caltech256 [17]	257	General	25K	6K	FGVC Aircrafts[33]	100	Fine-Grained	7K	3K
Stanford-Dog [24]	120	Fine-Grained	12K	8K	Blood-cell[45]	4	Medical Img.	10K	2K

Table 1. Datasets and their statistics used in this paper. Datasets in bold are used to construct the online learning training set. The rest are used to test our NASOA. It is commonly believed that Aircrafts, Flowers102 and Blood-cell deviate from the ImageNet domain.

Model Name	Top-1 Acc.	Inf Time (ms)	Training Step Time (ms)	Training GPU Usage (MB)
RegNetY-200MF[39]	70.40	14.25	62.30	2842
ET-NAS-C	71.29	8.94	26.28	2572
RegNetY-400MF[39]	74.10	20.57	90.61	4222
ET-NAS-D	74.46	14.54	36.30	3184
RegNetY-600MF[39]	75.50	22.15	90.11	4498
MobileNet-V3-Large[22]	75.20	16.88	71.65	12318
OFANet[6]	76.10	17.81	73.10	-
MNASNet-A1[47]	75.2	28.65	125.1	5642
ET-NAS-E	76.87	25.34	61.95	4922
EfficientNet-B0[48]	77.70	24.30	120.29	7778
RegNetY-1.6GF[39]	78.00	45.59	170.96	6338
ET-NAS-F	78.80	33.83	93.04	5800
EfficientNet-B2[48]	80.40	58.78	277.60	14258
RegNetY-16GF[39]	80.40	192.78	677.68	19258
ET-NAS-G	80.41	53.08	133.97	8120
ET-NAS-H	80.92	76.80	193.40	9140
EfficientNet-B3[48]	81.50	97.33	455.86	22368
ET-NAS-I	81.38	94.60	265.13	10732
ET-NAS-J	82.08	131.92	370.28	13774
ET-NAS-L	82.65	191.89	542.52	20556

Table 2. Comparison of our ET-NAS models and SOTA ImageNet models. Inference time, training step time and training GPU memory consumption are measured on single Nvidia V100, with $bs = 64$. Our models show a great advantage in terms of training speed and GPU memory usage.

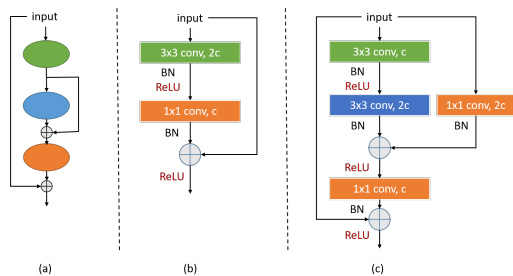


Figure 3. Two examples of our searched block in our ET-NAS. We found that smaller models should use simpler structures of blocks while bigger models prefer complex blocks for fast training.

final searched models in the Appendix. Surprisingly, we found that smaller models should use simpler structures of blocks while bigger models prefer complex blocks as shown in Figure 3. Comparing our searched backbone to the conventional ResNet/ResNeXt, we find that early stages in our

models are very short which is more efficient since feature maps in an early stage are very large and the computational cost is comparably large. This also verified our findings in preliminary experiments.

Comparison with the state-of-the-art ImageNet models. We compare the training/inference efficiency of our searched ET-NAS with the SOTA ImageNet models such as MobileNetV3 [22], RegNet series [39], and EfficientNet series [48] as shown in Table 2 and Figure 4. Overall, our searched models outperform other SOTA ImageNet models in terms of training accuracy and training speed from Figure 4 (left). Specifically, ET-NAS-G is about 6x training faster than RegNetY-16GF, and ET-NAS-I is about 1.5x training faster than EfficientNetB3. Our models are also better than mobile setting models such as MobileNetV2/V3[22] and RegNetY-200MF. Although our model is optimized for fast training, we also compare the inference speed in Figure 4(right). Our models still have a very strong performance in terms of inference speed, outperforming RegNet series and achieving comparable performance with EfficientNet.

Comparison with the NAS models. We also compare our method with state-of-the-art NAS methods: AmoebaNet[41], OFA[6], Darts[32], PCDarts[55], EfficientNet[48], RegNet[39] and so on. In Figure 5, it can be found that our searched models are more training efficient than other NAS results, e.g., some evolution-based NAS methods such as AmoebaNet, OFANet, and some weight sharing methods such as Darts and AmoebaNet. This is because of our flexible and effective search space, which considers both macro and micro-level structure.

What makes our network training efficient? To answer this, we define an *efficiency score* and conduct a statistical analysis of different factors for efficient-training (The analysis can be found in Appendix). We have the following conclusions: a) By observing optimal \mathcal{A}_i^* , smaller models should use simpler blocks while bigger models prefer complex blocks. Using the same block structure for all sizes of models [48, 39] may not be optimal. b) Adding redundant skip-connections which have great memory access cost will decrease the training efficiency of the model thus existing topological cell-level search space such as DARTS [32], AmoebaNet [40], and NASBench101 [12] is not efficient. c) The computation allocation on different stages is crucially important. Simply increasing depth/width to ex-

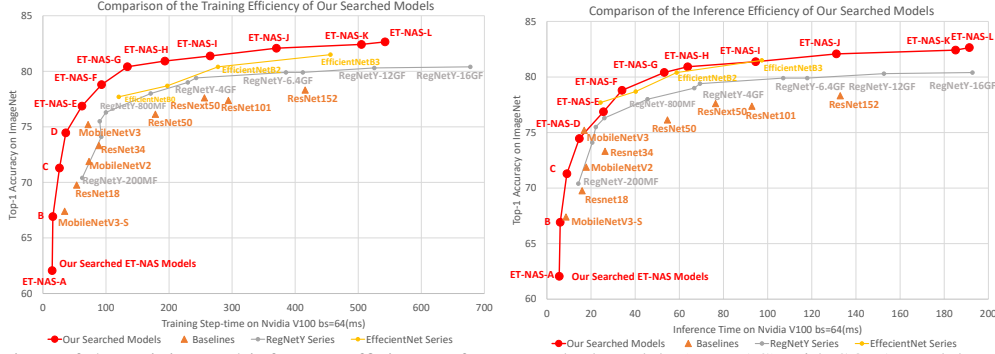


Figure 4. Comparison of the training and inference efficiency of our searched models (ET-NAS) with SOTA models on ImageNet. Our searched models are faster, e.g., ET-NAS-G is 6x training faster than RegNetY-16GF, and ET-NAS-I is 1.5x training faster than EfficientNetB3. Although our models are optimized for fast training, the inference speed is comparable to EfficientNet and better than RegNet.




Figure 5. Comparison of the training efficiency of our searched models (ET-NAS) with 8 other NAS results on ImageNet. It can be found that our method is more training efficient than some recent evolution-based NAS methods such as AmoebaNet, OFANet because of our effective search space.

pand the model as [48] may not be optimal will downgrade the performance. To conclude our novel joint search space contributes most to the training efficiency.

4.3. Results for Online Adaptive Predictor Acc_P

Experimental Settings. We evaluate our online algorithm based on ten widely used image classification datasets, that cover various fine-tuning tasks as shown in Table 1. Five of them (in bold) are chosen to be the online learning training set (meta-training). 30K samples are collected by continually sampling a subset of each dataset and fine-tuning with a randomized hyperparameters on it. Each subset varies from #. classes and #. images. The variables in Section 3.2.1 are calculated accordingly. The fine-tuning accuracy is evaluated on the test set. 30K sample is split into 24K meta-training and 6K meta-validation.

Then an adaptive MLP regression in Eq 3 are used to fit the data and predict the $Acc(A_i^{FT}, D_{val})$. We use $L = 10$ with 64 units in each hidden layer. We use a learning rate

Models	All Cumulative Err.		Segment 20-40%		Segment 80-100%	
	MAE	MSE	MAE	MSE	MAE	MSE
Fixed MLP (L=3)	10.07%	1.94%	8.99%	1.56%	7.99%	1.21%
Fixed MLP (L=6)	9.12%	1.71%	9.03%	1.62%	7.16%	1.04%
Fixed MLP (L=10)	8.45%	1.59%	8.46%	1.53%	6.68%	0.96%
Fixed MLP (L=14)	11.24%	2.91%	8.34%	1.54%	4.62%	0.46%
Ours Adaptive MLP	7.51%	1.36%	7.55%	1.11%	3.73%	0.28%

Table 3. Online error rate of our method and fixed MLP. Our adaptive MLP with hedge backpropagation is better in the online setting of predicting the fine-tuning accuracy.

of 0.01 and $\beta = 0.99$. As baselines, we also compare the results of using fixed MLP with different layers ($L = 3, 6, 10, 14$). MAE (mean absolute error) and MSE (mean square error) are performance metrics to measure the cumulative error with different segments of the tasks stream.

Comparison of online learning method. The cumulative error obtained by all the baselines and the proposed method to predict the fine-tuning accuracy is shown in Table 4. Our adaptive MLP with hedge backpropagation is better than fixed MLP in terms of the cumulative error of the predicted accuracy. Our method enjoys the benefit from the adaptive depth which allows faster convergence in the initial stage and strong predictive power in the later stage.

4.4. Final NASOA Results

To evaluate the performance of our whole NASOA, we select four time constraints on the testing datasets and use $Acc_P(\cdot)$ to test the fine-tuning accuracy. The testing datasets are MIT67, Food101, Aircrafts, Blood-cell, and Stanford-Car. The shortest/longest time constraint are the time of fine-tuning 10/50 epochs for ResNet18/ResNet101. The rest are equally divided into the log-space. For our NASOA, we generate the fine-tuning schedules by maximizing the predicted accuracy in Eq 2. We also conduct fine-tuning on various candidates of baselines such as ResNet (R18 to R152), RegNet (200MF to 16GF), and EfficientNet (B0 to B3) with the default hyperparameter setting in [27].

Comparison of the final fine-tuning results with the SOTA networks. We plot the time versus accuracy com-




Figure 6. Comparison of the final fine-tuning results under four time constraints for the testing dataset. Red square lines are the results of our NASOA in one-shot. The dots on the other solid line are the best performance of all the models in that series can perform. The model and training regime generated by our NASOA can outperform the upper bound of other methods in most cases. Our methods can improve around 2.1%~7.4% accuracy than the upper bound of RegNet/EfficientNet series on average.

Methods	Search Cost	Aircrafts	MIT67	Sf-Car	Sf-Dog
Random Search (HPO only)	x40	63.07%	75.60%	67.47%	86.25%
BOHB (HPO only)	x40	72.70%	77.61%	70.94%	87.41%
Random Search	x40	81.07%	79.93%	88.99%	89.06%
BOHB	x40	82.34%	79.85%	89.01%	89.49%
Our Z_{oo} with Random Search	x40	83.71%	80.97%	87.84%	92.75%
Our Z_{oo} with BOHB	x40	84.67%	82.34%	89.03%	93.74%
Our OA only	x1	81.22%	79.33%	84.56%	89.70%
Our Z_{oo} with Fixed MLP (Offline)	x1	81.31%	75.97%	88.81%	88.58%
Our final NASOA	x1	82.54%	80.30%	88.20%	92.30%

Table 4. Comparison of the final NASOA results with other HPO methods. “HPO only” means only optimizing the hyperparameters with RegNetY-16GF. Other HPO methods optimize both selecting hyperparameters and model from RegNet series models. “OA only” is our online schedule generator with RegNet series models. “Our Z_{oo} ” means using our models zoo to find suitable model. “Fixed MLP Predictor” is the offline baseline with fixed MLP predictor. “Our NASOA” is our whole pipeline with both the model zoo and online adaptive scheduler. Without additional search cost (x40), NASOA can reach similar performance of BOHB.

parison in Figure 6. As can be seen, the model and training regime generated by our NASOA can outperform the upper bound of other methods in most cases. On average, our methods can improve around 2.1%/7.4% accuracy than the best model of RegNet/EfficientNet series under various time constraints and tasks. It is noteworthy that our NASOA performs better especially in the case of short time constraint, which demonstrates that our schedule generator is capable provide both efficient and effective regimes.

Comparison of the final fine-tuning results with the HPO methods. In Table 4, we compare our method with the HPO methods which optimizing the hyperparameters and picking up models in ResNet, RegNetY, and EfficientNet series. “HPO only” means the method only optimizes the hyperparameters with a fixed model RegNetY-16GF. “OA only” is our online schedule generator with RegNet series models. “Our Z_{oo} ” means using our ET models zoo to find suitable model. “Fixed MLP Predictor” is the offline baseline with fixed MLP predictor (L=10) with our model zoo. “Our NASOA” is the our whole pipeline with

Methods	Fixed Model	Existing Models	NAS Models	Adaptive Scheduler	Comp Cost	Avg. Finetune Accuracy
BOHB[14]	✓				x40	77.2%
+ Our Z_{oo}			✓		x40 ^{-0x}	87.5% ^{+10.3%}
Our OA		✓		✓	x1 ^{-40x}	83.7% ^{+6.5%}
NASOA			✓	✓	x1 ^{-40x}	85.8% ^{+8.6%}

Table 5. Ablative study. We calculates the average fine-tuning accuracy over the test datasets.

both training efficient model zoo and online adaptive scheduler. Comparing to the offline baseline with our NASOA, our online adaption module can boost the average performance by 2.17%. It can also be found that our method can save up to 40x computational cost compared to HPO methods while reaching similar performance. With more computational budget, our model zoo with BOHB search can reach even higher accuracy (+avg. 10.28%).

Ablative interpretation of performance superiority. Table 5 calculates the average fine-tuning accuracy over tasks. Our NAS model zoo can greatly increase the fine-tuning average accuracy from 77.17% to 87.45%, which is the main contribution of the performance superiority. Using our online adaptive scheduler instead of BOHB can significantly reduce the computational cost (-40x).

5. Conclusion

We propose the first efficient task-oriented fine-tuning framework aiming at saving the resources for GPU clusters and cloud computing. The joint NAS and online adaption strategy achieves much better fine-tuning results in terms of both accuracy and speed. The searched architectures are more training-efficient than very strong baselines. We also theoretically prove our online model is better than fixed-depth model. Our experiments on multiple datasets show our NASOA achieves 40x speed-up comparing to BOHB. Generalization to more tasks such as detection and segmentation can be considered for future work.

References

- [1] Alessandro Achille, Michael Lam, Rahul Tewari, Avinash Ravichandran, Subhansu Maji, Charles C Fowlkes, Stefano Soatto, and Pietro Perona. Task2vec: Task embedding for meta-learning. In *ICCV*, pages 6430–6439, 2019. 2
- [2] Bowen Baker, Otkrist Gupta, Nikhil Naik, and Ramesh Raskar. Designing neural network architectures using reinforcement learning. In *ICLR*, 2017. 2
- [3] James Bergstra and Yoshua Bengio. Random search for hyper-parameter optimization. *JMLR*, 2012. 1
- [4] Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101—mining discriminative components with random forests. In *ECCV*, 2014. 4.1
- [5] Han Cai, Tianyao Chen, Weinan Zhang, Yong Yu, and Jun Wang. Efficient architecture search by network transformation. In *AAAI*, 2018. 2
- [6] Han Cai, Chuang Gan, Tianzhe Wang, Zhekai Zhang, and Song Han. Once-for-all: Train one network and specialize it for efficient deployment. *arXiv preprint arXiv:1908.09791*, 2019. 4.2, 4.2
- [7] Han Cai, Ligeng Zhu, and Song Han. Proxylessnas: Direct neural architecture search on target task and hardware. In *ICLR*, 2019. 2
- [8] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *TPAMI*, 2017. 1
- [9] Xin Chen, Yawen Duan, Zewei Chen, Hang Xu, Zihao Chen, Xiaodan Liang, Tong Zhang, and Zhenguo Li. Catch: Context-based meta reinforcement learning for transferrable architecture search. In *ECCV*, 2020. 1
- [10] Yin Cui, Yang Song, Chen Sun, Andrew Howard, and Serge Belongie. Large scale fine-grained categorization and domain-specific transfer learning. In *CVPR*, 2018. 1, 2, 3.2.1
- [11] Kalyanmoy Deb, Samir Agrawal, Amrit Pratap, and Tanaka Meyarivan. A fast elitist non-dominated sorting genetic algorithm for multi-objective optimization: Nsga-ii. In *International conference on parallel problem solving from nature*, pages 849–858. Springer, 2000. 3.1
- [12] Xuanyi Dong and Yi Yang. Nas-bench-201: Extending the scope of reproducible neural architecture search. In *ICLR*, 2019. 3.1, 4.2
- [13] Yawen Duan, Xin Chen, Hang Xu, Zewei Chen, Xiaodan Liang, Tong Zhang, and Zhenguo Li. Transnas-bench-101: Improving transferability and generalizability of cross-task neural architecture search. In *CVPR*, 2021. 1
- [14] Stefan Falkner, Aaron Klein, and Frank Hutter. BOHB: robust and efficient hyperparameter optimization at scale. In *ICML*, 2018. 1, 2, 4.4
- [15] Li Fei-Fei, Rob Fergus, and Pietro Perona. One-shot learning of object categories. *TPAMI*, 2006. 4.1
- [16] Yoav Freund and Robert E Schapire. Adaptive game playing using multiplicative weights. *Games and Economic Behavior*, 29(1-2):79–103, 1999. 3.2.1
- [17] Gregory Griffin, Alex Holub, and Pietro Perona. Caltech-256 object category dataset. 2007. 4.1
- [18] Yunhui Guo, Honghui Shi, Abhishek Kumar, Kristen Grauman, Tajana Rosing, and Rogerio Feris. Spottune: transfer learning through adaptive fine-tuning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4805–4814, 2019. 4.1
- [19] Kaiming He, Ross B. Girshick, and Piotr Dollár. Rethinking imagenet pre-training. In *ICCV*, 2019. 1
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 3.1, 3.1
- [21] Steven CH Hoi, Doyen Sahoo, Jing Lu, and Peilin Zhao. Online learning: A comprehensive survey. *arXiv preprint arXiv:1802.02871*, 2018. 2
- [22] Andrew Howard, Mark Sandler, Grace Chu, Liang-Chieh Chen, Bo Chen, Mingxing Tan, Weijun Wang, Yukun Zhu, Ruoming Pang, Vijay Vasudevan, et al. Searching for mobilenetv3. In *CVPR*, 2019. 2, 4.2, 4.2
- [23] Chenhan Jiang, Hang Xu, Xiangdan Liang, and Liang Lin. Hybrid knowledge routed modules for large-scale object detection. In *NeurIPS*, 2018. 1
- [24] Aditya Khosla, Nityananda Jayadevaprakash, Bangpeng Yao, and Li Fei-Fei. Novel dataset for fine-grained image categorization. In *Workshop on Fine-Grained Visual Categorization, CVPR*, 2011. 4.1
- [25] Simon Kornblith, Jonathon Shlens, and Quoc V Le. Do better imagenet models transfer better? In *CVPR*, 2019. 1, 2, 3.1, 4.1
- [26] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *4th International IEEE Workshop on 3D Representation and Recognition (3dRR-13)*, Sydney, Australia, 2013. 4.1
- [27] Hao Li, Pratik Chaudhari, Hao Yang, Michael Lam, Avinash Ravichandran, Rahul Bhotika, and Stefano Soatto. Rethinking the hyperparameters for fine-tuning. In *ICLR*, 2020. 1, 2, 3.2.1, 4.4
- [28] Zeming Li, Chao Peng, Gang Yu, Xiangyu Zhang, Yangdong Deng, and Jian Sun. Detnet: A backbone network for object detection. In *ECCV*, 2018. 3.1
- [29] Feng Liang, Chen Lin, Ronghao Guo, Ming Sun, Wei Wu, Junjie Yan, and Wanli Ouyang. Computation reallocation for object detection. In *ICLR*. OpenReview.net, 2020. 3.1
- [30] Chenxi Liu, Liang-Chieh Chen, Florian Schroff, Hartwig Adam, Wei Hua, Alan L. Yuille, and Fei-Fei Li. Auto-deeplab: Hierarchical neural architecture search for semantic image segmentation. In *CVPR*, 2019. 1
- [31] Chenxi Liu, Barret Zoph, Maxim Neumann, Jonathon Shlens, Wei Hua, Li-Jia Li, Li Fei-Fei, Alan Yuille, Jonathan Huang, and Kevin Murphy. Progressive neural architecture search. In *ECCV*, 2018. 1, 2
- [32] Hanxiao Liu, Karen Simonyan, and Yiming Yang. Darts: Differentiable architecture search. In *ICLR*, 2018. 1, 2, 3.1, 4.2, 4.2
- [33] Subhansu Maji, Esa Rahtu, Juho Kannala, Matthew Blaschko, and Andrea Vedaldi. Fine-grained visual classification of aircraft. *arXiv preprint arXiv:1306.5151*, 2013. 4.1

- [34] Hector Mendoza, Aaron Klein, Matthias Feurer, Jost Tobias Springenberg, and Frank Hutter. Towards automatically-tuned neural networks. In *Workshop on Automatic Machine Learning*, pages 58–65, 2016. 1, 2
- [35] Gaurav Mittal, Chang Liu, Nikolaos Karianakis, Victor Fragoso, Mei Chen, and Yun Fu. Hyperstar: Task-aware hyperparameters for deep networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8736–8745, 2020. 2
- [36] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *2008 Sixth Indian Conference on Computer Vision, Graphics & Image Processing*, pages 722–729. IEEE, 2008. 1, 4.1
- [37] Ariadna Quattoni and Antonio Torralba. Recognizing indoor scenes. In *CVPR*, 2009. 4.1
- [38] Ilija Radosavovic, Justin Johnson, Saining Xie, Wan-Yen Lo, and Piotr Dollár. On network design spaces for visual recognition. In *ICCV*, pages 1882–1890, 2019. 1
- [39] Ilija Radosavovic, Raj Prateek Kosaraju, Ross Girshick, Kaiming He, and Piotr Dollár. Designing network design spaces. *arXiv preprint arXiv:2003.13678*, 2020. 3.1, 3.1, 4.2, 4.2, 4.2
- [40] Esteban Real, Alok Aggarwal, Yanping Huang, and Quoc V Le. Regularized evolution for image classifier architecture search. In *AAAI*, 2019. 1, 3.1, 4.2
- [41] Esteban Real, Alok Aggarwal, Yanping Huang, and Quoc V Le. Regularized evolution for image classifier architecture search. In *Proceedings of the aaai conference on artificial intelligence*, volume 33, pages 4780–4789, 2019. 4.2
- [42] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *IJCV*, 115(3):211–252, 2015. 4.2
- [43] Doyen Sahoo, Quang Pham, Jing Lu, and Steven CH Hoi. Online deep learning: Learning deep neural networks on the fly. In *IJCAI*, 2017. 2, 3.2.1
- [44] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *CVPR*, pages 4510–4520, 2018. 3.1, 3.1
- [45] Ishpreet Singh, Narinder Pal Singh, Harnoor Singh, Saharsh Bawankar, and Alioune Ngom. Blood cell types classification using cnn. In *International Work-Conference on Bioinformatics and Biomedical Engineering*, pages 727–738. Springer, 2020. 4.1
- [46] Emma Strubell, Ananya Ganesh, and Andrew McCallum. Energy and policy considerations for deep learning in NLP. In *ACL*, 2019. 1, 2
- [47] Mingxing Tan, Bo Chen, Ruoming Pang, Vijay Vasudevan, Mark Sandler, Andrew Howard, and Quoc V. Le. Mnasnet: Platform-aware neural architecture search for mobile. In *CVPR*, 2019. 2, 4.2
- [48] Mingxing Tan and Quoc V. Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *ICML*, volume 97 of *Proceedings of Machine Learning Research*, pages 6105–6114. PMLR, 2019. 1, 2, 4.2, 4.2, 4.2
- [49] Mingxing Tan, Ruoming Pang, and Quoc V Le. Efficientdet: Scalable and efficient object detection. *arXiv preprint arXiv:1911.09070*, 2019. 1, 3.1
- [50] P. Welinder, S. Branson, T. Mita, C. Wah, F. Schroff, S. Belongie, and P. Perona. Caltech-UCSD Birds 200. Technical Report CNS-TR-2010-001, California Institute of Technology, 2010. 1, 4.1
- [51] Yangxin Wu, Gengwei Zhang, Hang Xu, Xiaodan Liang, and Liang Lin. Auto-panoptic: Cooperative multi-component architecture search for panoptic segmentation. In *NeurIPS*, 2020. 2
- [52] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *CVPR*, 2017. 3.1, 3.1
- [53] Sirui Xie, Hehui Zheng, Chunxiao Liu, and Liang Lin. Snas: stochastic neural architecture search. In *ICLR*, 2019. 2
- [54] Hang Xu, Lewei Yao, Wei Zhang, Xiaodan Liang, and Zhen-guo Li. Auto-fpn: Automatic network architecture adaptation for object detection beyond classification. In *ICCV*, 2019. 1
- [55] Yuhui Xu, Lingxi Xie, Xiaopeng Zhang, Xin Chen, Guo-Jun Qi, Qi Tian, and Hongkai Xiong. Pc-darts: Partial channel connections for memory-efficient architecture search. *arXiv preprint arXiv:1907.05737*, 2019. 4.2
- [56] Lewei Yao, Hang Xu, Wei Zhang, Xiaodan Liang, and Zhen-guo Li. SM-NAS: Structural-to-modular neural architecture search for object detection. 2020. 1
- [57] Zhao Zhong, Junjie Yan, Wei Wu, Jing Shao, and Cheng-Lin Liu. Practical block-wise neural network architecture generation. In *CVPR*, 2018. 2
- [58] Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V Le. Learning transferable architectures for scalable image recognition. In *CVPR*, 2018. 1