# AUTOHAS: EFFICIENT HYPERPARAMETER AND ARCHITECTURE SEARCH

**Xuanyi Dong**[†‡*] **Mingxing Tan**[†], **Adams Wei Yu**[†], **Daiyi Peng**[†], **Bogdan Gabrys**[‡], **Quoc Le**[†]
[†] Google Research, Brain Team, [‡] AAI, University of Technology Sydney

## ABSTRACT

Efficient hyperparameter or architecture search methods have shown remarkable results, but each of them is only applicable to searching for either hyperparameters (HPs) or architectures. In this work, we propose a unified pipeline, AutoHAS, to efficiently search for both architectures and hyperparameters. AutoHAS learns to alternately update the shared network weights and a reinforcement learning (RL) controller, which learns the probability distribution for the architecture candidates and HP candidates. A temporary weight is introduced to store the updated weight from the selected HPs (by the controller), and a validation accuracy based on this temporary weight serves as a reward to update the controller. In experiments, we show AutoHAS is efficient and generalizable to different search spaces, baselines and datasets. In particular, AutoHAS can improve the accuracy over popular network architectures, such as ResNet and EfficientNet, on CIFAR-10/100, ImageNet, and four more other datasets.

## 1 INTRODUCTION

Deep learning models require intensive efforts in optimizing architectures and hyperparameters. Standard hyperparameter optimization methods, such as grid search, random search or Bayesian optimization, are inefficient because they are *multi-trial*: different configurations are tried in parallel to find the best configuration. As these methods are expensive, there is a trend towards more efficient, single-trial methods for specific hyperparameters. For example, the learning rate can be optimized with the hypergradient method (Baydin et al., 2018). Similarly, many architecture search methods started out multi-trial (Zoph & Le, 2017; Baker et al., 2017; Real et al., 2019), but more recent proposals are single-trial (Pham et al., 2018; Liu et al., 2019). These efficient methods, however, sacrifice generality: each method only works for one aspect or a subset of the hyperparameters or architectures.

In this paper, we generalize those efficient, single-trial methods to include both architectures and a broader range of hyperparameters[1]. One important benefit of the generalization is that we can have a general, efficient method for hyperparameter optimization as a special case. Another benefit is that we can now search for both hyperparameters and architectures in a single model. Practically, this means that our method is an improvement over neural architecture search (NAS) because each model can potentially be coupled with its own best hyperparameters, thus achieving comparable or even better performance than existing NAS with *fixed* hyperparameters.

To this end, we propose AutoHAS, an efficient hyperparameter and architecture search framework. It is, to the best of our knowledge, *the first method that can efficiently handle architecture space, hyperparameter space, or the joint search space.* AutoHAS utilizes the weight sharing technique proposed by Pham et al. (2018). The main idea is to train a super model, where each candidate in the architecture space is its sub-model. Using a super model can avoid training millions of candidates from scratch (Liu et al., 2019; Pham et al., 2018). AutoHAS extends its scope from architecture search to both architecture and hyperparameter search. We not only share the weights of super model with each architecture but also share this super model across hyperparameters. At each search step,

[*] Work done as a research intern at Google.

[1] In this paper, hyperparameters refer all design choices that will affect the training procedure of a model, such as learning rate, weight decay, optimizer, dropout, augmentation policy, etc.

| | learning rate (LR) | weight decay | augmentation | dropout | architecture | **efficient** |
|---|---|---|---|---|---|---|
| Bayesian | √ | √ | √ | √ | √ | × |
| RL or Evolution | √ | √ | √ | √ | √ | × |
| PBT | √ | √ | √ | √ | × | × |
| Gradient Descent on LR | √ | × | × | × | × | √ |
| Hypergradient | × | √ | √ | × | √ | √ |
| NAS (Weight Sharing) | × | × | × | × | √ | √ |
| **AutoHAS** | √ | √ | √ | √ | √ | √ |

Table 1: We compare the different aspects of each algorithm. We consider a search algorithm is efficient if it can complete the search within less than $10\times$ computational costs than training a single model. Bayesian optimization, traditional RL and evolutionary algorithms are applicable to all kinds of hyperparameters and architectures, but they are computationally expensive. Population based training (PBT) methods (Jaderberg et al., 2017; Li et al., 2019) utilized the incomplete observations to accelerate the evolutionary algorithms for HPO, but they are still far from efficient. Recent hypergradient-based HPO methods (Lorraine et al., 2020; Pedregosa, 2016; Shaban et al., 2019; Baydin et al., 2018) and weight sharing-based NAS methods (Liu et al., 2019; Xie et al., 2019; Pham et al., 2018; Dong & Yang, 2019) are efficient but sacrifice the generality. Our AutoHAS takes care of both efficiency and generality.

AutoHAS optimizes the sampled sub-model by a combination of the sampled hyperparameter choices, and the shared weights of super model serves as a good initialization for all hyperparameters at the next step of search (see Fig. 1 and Sec. 2). In order to decouple the shared network weights ($\mathcal{W}$ in Fig. 1) and controller optimization, we also propose to create temporary weights for evaluating the sampled hyperparameters and updating the controller.

A challenge here is that architecture choices (e.g. kernel size) are often categorical values whereas hyperparameter choices (e.g. learning rate) are often continuous values. To address the mixture of categorical and continuous search spaces, we first discretize the continuous hyperparameters into a linear combination of multiple categorical basis. The discretization allows us to unify architecture and hyperparameter choices during search. As explained below, we will use a reinforcement learning (RL) method to search over these discretized choices in Fig. 1. The probability distribution over all candidates is naturally learnt by the RL controller, and it is used as the coefficient in the linear combination to find the best architecture and hyperparameters.

AutoHAS shows generalizability, efficiency and scaling to large datasets in experiments. It consistently improves the baselines' accuracy when searching for architectures, hyperparameters, and both on CIFAR-10, CIFAR-100, ImageNet, Place365, etc. In experiments, AutoHAS shows non-trivial improvements on **eight** datasets, such as 0.8% accuracy gain on highly-optimized EfficientNet and $11\%$ accuracy gain on less-optimized models. AutoHAS reduces the search cost by over $10\times$ than random search and Bayesian search. We also summarize the benefits of our AutoHAS over other search methods in Table 1.

## 2 AUTOHAS

AutoHAS can handle the general case of NAS and HPO – able to find both architecture $\alpha$ and hyperparameters $h$ that achieve high performance on the validation set $\mathcal{D}_{val}$. This objective can be formulated as a bi-level optimization problem:

$$\min_{\alpha,h} \mathcal{L}(\alpha, \omega^*_{\alpha,h}, \mathcal{D}_{val}) \quad \text{s.t. } \omega^*_{\alpha,h} = f_h(\alpha, \mathcal{D}_{train}), \tag{1}$$

where $\mathcal{L}$ is the objective function (e.g., cross-entropy loss). $\mathcal{D}_{train}$ and $\mathcal{D}_{val}$ denote the training data and the validation data, respectively. $f_h$ represents the algorithm with hyperparameters $h$ to obtain the optimal weights $\omega^*_{\alpha,h}$. For example, $f_h$ could be using SGD to minimize the training loss, where $h$ denotes the hyperparameters of SGD. In this case, $\omega^*_{\alpha,h}$ is the final optimized weights after the SGD converged.

AutoHAS generalizes both NAS and HPO by introducing a broader search space. On the one hand, NAS is a special case of HAS, where $h$ is fixed in Eq. (1). On the other hand, HPO is a special case of HAS, where $\alpha$ is fixed in Eq. (1).

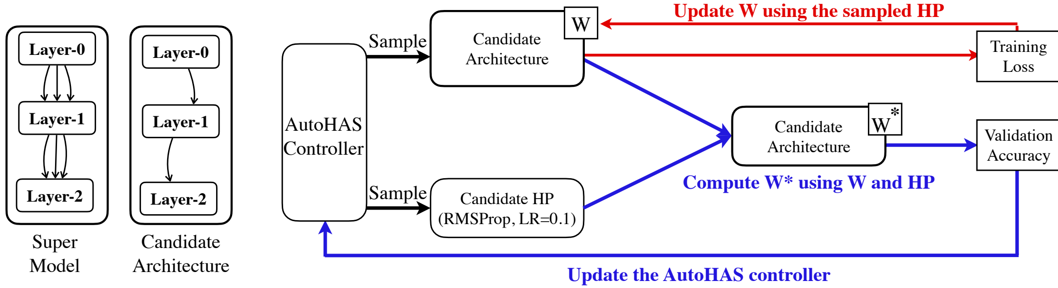


Figure 1: **The overview of AutoHAS.** LEFT: Each candidate architecture's weights are shared with a super model, where each candidate is a sub model within this super model. RIGHT: During the search, AutoHAS alternates between optimizing the shared weights of super model $\mathcal{W}$ and updating the controller. It also creates temporary weights $\mathcal{W}^*$ by optimizing the sampled candidate architecture using the sampled candidate hyperparameter (HP). This $\mathcal{W}^*$ will be used to compute the validation accuracy as a reward so as to update the AutoHAS controller to select better candidates. Finally, $\mathcal{W}^*$ is discarded after updating the controller so as not to affect the original $\mathcal{W}$.

### 2.1 Unified Representation of Hyperparameters and Architectures

The search space in AutoHAS is a Cartesian product of the architecture and hyperparameter candidates. To search over this mixed search space, we need a unified representation of different searchable components, i.e., architectures, learning rates, optimizer, etc.

**Architectures Search Space** We use the simplest case as an example. First of all, let the set of predefined candidate operations (e.g., 3x3 convolution, pooling, etc.) be $\mathcal{O} = \{O_1, O_2, ..., O_n\}$, where the cardinality of $\mathcal{O}$ is $n$ for each layer in the architecture. Suppose an architecture is constructed by stacking multiple layers, each layer takes a tensor $F$ as input and output $\pi(F)$, which serves as the next layer's input. $\pi \in \mathcal{O}$ denotes the operation at a layer and might be different at different layers. Then a candidate architecture $\alpha$ is essentially the sequence for all layers $\{\pi\}$. Further, a layer can be represented as a linear combination of the operations in $\mathcal{O}$ as follows:

$$\pi(F) = \sum\nolimits_{i=1}^{n} C_i^{\alpha}\, O_i(F) \quad \text{s.t.} \; \sum\nolimits_{i=1}^{n} C_i^{\alpha} = 1, C_i^{\alpha} \in \{0, 1\}, \tag{2}$$

where $C_i^{\alpha}$ (the $i$-th element of the vector $C^{\alpha}$) is the coefficient of operation $O_i$ for a layer.

**Hyperparameter Search Space** Now we can define the hyperparameter search space in a similar way. The major difference is that we have to consider both categorical and continuous cases:

$$h = \sum\nolimits_{i=1}^{m} C_i^{h}\, \mathcal{B}_i \quad \text{s.t.} \; \sum\nolimits_{i=1}^{m} C_i^{h} = 1, \text{ and } C_i^{h} \in \begin{cases} [0, 1], & \text{if continuous} \\ \{0, 1\}, & \text{if categorical} \end{cases}, \tag{3}$$

where $\mathcal{B}$ is a predefined set of hyperparameter basis with the cardinality of $m$ and $\mathcal{B}_i$ is the $i$-th basis in $\mathcal{B}$. $C_i^h$ (the $i$-th element of the vector $C^h$) is the coefficient of hyperparameter basis $\mathcal{B}_i$. If we have a continuous hyperparameter, we have to discretize it into a linear combination of basis and unify both categorical and continuous. For example, for weight decay, $\mathcal{B}$ could be {1e-1, 1e-2, 1e-3}, and therefore, all possible weight decay values can be represented as a linear combination over $\mathcal{B}$. For categorical hyperparameters, taking the optimizer as an example, $\mathcal{B}$ could be {Adam, SGD, RMSProp}. In this case, a constraint on $C_i^h$ is applied: $C_i^h \in \{0, 1\}$ as in Eq. (3).

### 2.2 Efficient Hyperparameter and Architecture Search

Given the discretizing strategy in Sec. 2.1, each candidate in the search space can be represented by the value of $\mathbb{C} = \{C^{\alpha}$ for all layers, $C^h$ for all hyperparameters$\}$, which represents the coefficients for all architecture and hyperparameter choices. As a result, AutoHAS converts the searching problem to obtaining the coefficients $\mathbb{C}$. Below we will introduce how to calculate these coefficients.

AutoHAS applies reinforcement learning together with weight sharing to search over the discretized space. During search, we learn a controller to sample the candidate architecture and hyperparameters from the discretized space. In AutoHAS, this controller is parameterized by a collection of

independent multinomial variables $\mathcal{P} = \{P^\alpha \text{ for all layers}, P^h \text{ for all hyperparameters}\}$ , which draws the probability distribution of the discretized space. AutoHAS also leverages a super model to share weights $\mathcal{W}$ among all candidate architectures, where each candidate is a sub-model in this super model. Furthermore, AutoHAS extends the scope of weight sharing from architecture to hyperparameters, where $\mathcal{W}$ also serves as the initialization for the algorithm $f_h$.

AutoHAS alternates between learning the shared weights $\mathcal{W}$ and learning the controller using REINFORCE (Williams, 1992). Specifically, at each iteration, the controller samples a candidate — an architecture $\alpha$ and basis hyperparameter $h \in \mathcal{B}$. We estimate its quality $Q(\alpha, h)$ by utilizing the temporary weights $\mathcal{W}^*_\alpha$, which are generated by applying the gradients from training loss to the original weights $\mathcal{W}_\alpha$ of the architecture $\alpha$ with hyperparameters $h$. This estimated quality is used as a reward to update the controller's parameters $\mathcal{P}$ via REINFORCE. Then, we optimize the shared weights $\mathcal{W}$ by minimizing the training loss calculated by the sampled architecture. More details can be found in Alg. 1 in Appendix.

**Discussion** In practice, the training of shared weights in efficient NAS often suffers from the instability problem (Liu et al., 2019; Dong & Yang, 2019; Zela et al., 2020), and that for HAS can be more pronounced. To make AutoHAS more stable, we will sample tens of candidate pairs of $\alpha$ and $h$, and average the gradients from all these pairs for $\mathcal{W}$ as well as that for $\mathcal{P}$ to update $\mathcal{W}$ and $\mathcal{P}$. In this way, the high variance of the gradients from different samples can be significantly reduced. In experiments, this strategy is achieved via training on the distribution system, such that each core can individually sample a different pair.

During the aforementioned strategy, the temporary weights allows us to effectively decouple the shared network weights and controller optimization. If we directly override $\mathcal{W}$ without using temporary weights $\mathcal{W}^*$, it will pollute the shared weights update. This is because that different sampled pairs will use different weights for overriding and quarrel with each other.

### 2.3 Deriving Hyperparameters and Architecture

After AutoHAS learns $\mathcal{P} = \{P^\alpha, P^h\}$, we can derive the coefficient $\mathbb{C}$ as follows:

$$C^\alpha = \text{one_hot}(\arg\max_i \, P^\alpha) \quad C^h = \begin{cases} P^h & \text{if continuous} \\ \text{one_hot}(\arg\max_i \, P^h) & \text{if categorical} \end{cases}, \tag{4}$$

Together with Eq. (2) and Eq. (3), we can derive the final architecture $\alpha$ and hyperparameters $h$. Intuitively speaking, the selected operation in the final architecture has the highest probability over other candidates, and so does the categorical hyperparameter. For the continuous hyperparameter, the final one is the weighted sum of the learnt probability $P^h$ with its basis $\mathcal{B}$.

To evaluate whether the AutoHAS-discovered $\alpha$ and $h$ is good or not, we will use $h$ to re-train $\alpha$ on the whole training set and report its performance on the test sets.

## 3 Experimental Analysis

AutoHAS is a unified, efficient, and general framework for joint architecture and hyperparameter search. We show its applicability and efficiency when integrated with RL-based and differentiable-based searching algorithm in Table 2. It is also easy to switch from RL to a Bayesian-based searching algorithm, which will be investigated in the future.

In Table 2, we use the popular MobileNetV2 (Sandler et al., 2018) as our baseline model. We compare its performance with AutoHAS-discovered architecture and hyperparameters. On the same par of computational resources as training MobileNet-V2, AutoHAS finds a better model. In addition, AutoHAS (REINFORCE) is more efficient than AutoHAS (Differentiable). More empirical analysis can be found in Appendix.

| | #Params (MB) | #FLOPs (M) | ImageNet Accuracy (%) | Search Cost | |
|---|---|---|---|---|---|
| | | | | Memory (GB) | Time (TPU Hour) |
| Baseline model | 1.5 | 35.9 | 50.96 | 1.0 (Training Cost) | 44.8 (Training Cost) |
| AutoHAS (Differentiable) | 1.5 | 36.1 | 52.17 | 6.1 | 92.8 |
| AutoHAS (REINFORCE) | 1.5 | 36.3 | 53.01 | **1.8** | **54.4** |

Table 2: AutoHAS Differentiable Search vs. AutoHAS REINFORCE Search on ImageNet.

**Acknowledgements** We want to thank Gabriel Bender, Hanxiao Liu, Hieu Pham, Ruoming Pang, Barret Zoph and Yanqi Zhou for their help and feedback.

## REFERENCES

Bowen Baker, Otkrist Gupta, Nikhil Naik, and Ramesh Raskar. Designing neural network architectures using reinforcement learning. In *ICLR*, 2017.

Atılım Güneş Baydin, Robert Cornish, David Martínez Rubio, Mark Schmidt, and Frank Wood. Online learning rate adaptation with hypergradient descent. In *ICLR*, 2018.

Gabriel Bender, Hanxiao Liu, Bo Chen, Grace Chu, Shuyang Cheng, Pieter-Jan Kindermans, and Quoc V Le. Can weight sharing outperform random architecture search? an investigation with tunas. In *CVPR*, 2020.

James Bergstra and Yoshua Bengio. Random search for hyper-parameter optimization. *JMLR*, 2012.

Han Cai, Ligeng Zhu, and Song Han. ProxylessNAS: Direct neural architecture search on target task and hardware. In *ICLR*, 2019.

Han Cai, Chuang Gan, and Song Han. Once for all: Train one network and specialize it for efficient deployment. In *ICLR*, 2020.

Xiaoliang Dai, Alvin Wan, Peizhao Zhang, Bichen Wu, Zijian He, Zhen Wei, Kan Chen, Yuandong Tian, Matthew Yu, Peter Vajda, et al. FBNetV3: Joint architecture-recipe search using neural acquisition function. *arXiv preprint arXiv:2006.02049*, 2020.

Erik A Daxberger, Anastasia Makarova, Matteo Turchetta, and Andreas Krause. Mixed-variable bayesian optimization. In *IJCAI*, 2020.

Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *CVPR*, 2009.

Tobias Domhan, Jost Tobias Springenberg, and Frank Hutter. Speeding up automatic hyperparameter optimization of deep neural networks by extrapolation of learning curves. In *IJCAI*, 2015.

Xuanyi Dong and Yi Yang. Searching for a robust neural architecture in four gpu hours. In *CVPR*, 2019.

Daniel Golovin, Benjamin Solnik, Subhodeep Moitra, Greg Kochanski, John Karro, and D Sculley. Google vizier: A service for black-box optimization. In *SIGKDD*, 2017.

Frank Hutter. *Automated configuration of algorithms for solving hard computational problems*. PhD thesis, University of British Columbia, 2009.

Frank Hutter, Holger H Hoos, and Kevin Leyton-Brown. Sequential model-based optimization for general algorithm configuration. In *International Conference on Learning and Intelligent Optimization*, 2011.

Frank Hutter, Lars Kotthoff, and Joaquin Vanschoren. *Automated Machine Learning*. Springer, 2019.

Max Jaderberg, Valentin Dalibard, Simon Osindero, Wojciech M Czarnecki, Jeff Donahue, Ali Razavi, Oriol Vinyals, Tim Green, Iain Dunning, Karen Simonyan, et al. Population based training of neural networks. *arXiv preprint arXiv:1711.09846*, 2017.

Donald R Jones, Matthias Schonlau, and William J Welch. Efficient global optimization of expensive black-box functions. *Journal of Global Optimization*, 1998.

Aaron Klein and Frank Hutter. Tabular benchmarks for joint architecture and hyperparameter optimization. *arXiv preprint arXiv:1905.04970*, 2019.

Ron Kohavi and George H John. Automatic parameter selection by minimizing estimated error. In *Machine Learning Proceedings*, 1995.

Ang Li, Ola Spyra, Sagi Perel, Valentin Dalibard, Max Jaderberg, Chenjie Gu, David Budden, Tim Harley, and Pramod Gupta. A generalized framework for population based training. In *SIGKDD*, 2019.

Liam Li and Ameet Talwalkar. Random search and reproducibility for neural architecture search. In *UAI*, 2020.

Lisha Li, Kevin Jamieson, Giulia DeSalvo, Afshin Rostamizadeh, and Ameet Talwalkar. Hyperband: A novel bandit-based approach to hyperparameter optimization. *JMLR*, 2017.

Hanxiao Liu, Karen Simonyan, and Yiming Yang. Darts: Differentiable architecture search. In *ICLR*, 2019.

Jonathan Lorraine, Paul Vicol, and David Duvenaud. Optimizing millions of hyperparameters by implicit differentiation. In *AISTATS*, 2020.

Dougal Maclaurin, David Duvenaud, and Ryan Adams. Gradient-based hyperparameter optimization through reversible learning. In *ICML*, 2015.

Fabian Pedregosa. Hyperparameter optimization with approximate gradient. In *ICML*, 2016.

Hieu Pham, Melody Y Guan, Barret Zoph, Quoc V Le, and Jeff Dean. Efficient neural architecture search via parameter sharing. In *ICML*, 2018.

Esteban Real, Alok Aggarwal, Yanping Huang, and Quoc V Le. Regularized evolution for image classifier architecture search. In *AAAI*, 2019.

Binxin Ru, Ahsan Alvi, Vu Nguyen, Michael A Osborne, and Stephen Roberts. Bayesian optimisation over multiple continuous and categorical inputs. In *ICML*, 2020.

Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. MobileNetV2: Inverted residuals and linear bottlenecks. In *CVPR*, 2018.

Amirreza Shaban, Ching-An Cheng, Nathan Hatch, and Byron Boots. Truncated back-propagation for bilevel optimization. In *AISTATS*, 2019.

Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P Adams, and Nando De Freitas. Taking the human out of the loop: A review of bayesian optimization. *Proceedings of the IEEE*, 2015.

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 *ICML*, 2015.

Mingxing Tan and Quoc Le. EfficientNet: Rethinking model scaling for convolutional neural networks. In *ICML*, 2019.

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.

Alvin Wan, Xiaoliang Dai, Peizhao Zhang, Zijian He, Yuandong Tian, Saining Xie, Bichen Wu, Matthew Yu, Tao Xu, Kan Chen, et al. FBNetV2: Differentiable neural architecture search for spatial and channel dimensions. In *CVPR*, 2020.

Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 1992.

Bichen Wu, Xiaoliang Dai, Peizhao Zhang, Yanghan Wang, Fei Sun, Yiming Wu, Yuandong Tian, Peter Vajda, Yangqing Jia, and Kurt Keutzer. FBNet: Hardware-aware efficient convnet design via differentiable neural architecture search. In *CVPR*, 2019.

Sirui Xie, Hehui Zheng, Chunxiao Liu, and Liang Lin. SNAS: stochastic neural architecture search. In *ICLR*, 2019.

Arber Zela, Aaron Klein, Stefan Falkner, and Frank Hutter. Towards automated deep learning: Efficient joint neural architecture and hyperparameter search. In *ICML-W*, 2018.

Arber Zela, Thomas Elsken, Tonmoy Saikia, Yassine Marrakchi, Thomas Brox, and Frank Hutter. Understanding and robustifying differentiable architecture search. In *ICLR*, 2020.

Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. In *ICLR*, 2018.

Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. Places: A 10 million image database for scene recognition. *TPAMI*, 2017.

Barret Zoph and Quoc V Le. Neural architecture search with reinforcement learning. In *ICLR*, 2017.

Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V Le. Learning transferable architectures for scalable image recognition. In *CVPR*, 2018.

## A EXPLAIN AUTOHAS ALGORITHM

AutoHAS alternates between optimizing the super model (parameterized by the shared weights $\mathcal{W}$) and the AutoHAS controller (parameterized by the multinomial variables $\mathcal{P}$). Once AutoHAS completes the training of controller, we will follow Sec. 2.3 to derive the final architecture and hyperparameters.

To optimize the AutoHAS controller, it samples a candidate — an architecture $\alpha$ and basis hyperparameter $h \in \mathcal{B}$. Taking the architecture $\alpha$'s weights $\mathcal{W}_\alpha$ as the initialization, we apply hyperparameter $h$ to update $\mathcal{W}_\alpha$ by one or multiple gradient descent steps to obtain the temporary weights $\mathcal{W}_\alpha^*$. We use the temporary weights $\mathcal{W}_\alpha^*$ to calculate the quality $Q(\alpha, h)$ to update the AutoHAS controller. After that, the temporary weights $\mathcal{W}_\alpha^*$ will be discarded and not affect the original weights $\mathcal{W}_\alpha$. When we target better accuracy for HAS, the quality $Q(\alpha, h)$ is the accuracy of a batch of validation data. When we target on FLOP-constrained problem, the quality $Q(\alpha, h)$ becomes a hybrid one combing both accuracy and FLOPs, such as the hard exponential reward function in MNasNet (Tan et al., 2019) or absolute reward function in TuNAS (Bender et al., 2020). In Alg. 1, we show an example of using the simplest RL algorithm – REINFORCE – to optimize the controller. AutoHAS is general and can be integrated with a differentiable searching algorithm to optimize AutoHAS controller[2]. Bayesian optimization can also be easily integrated into AutoHAS, where the acquisition function serves to sample candidates and the quality used to improve the surrogate model in Bayesian optimization. We would investigate these extensions in future.

To optimize the super model's weights, AutoHAS controller samples a candidate — an architecture $\alpha$ and basis hyperparameter $h \in \mathcal{B}$. Then, AutoHAS use the optimization algorithm represented by $h$ to update the weights $\mathcal{W}_\alpha$ in the super model corresponding to the architecture $\alpha$.

In practice, we average the gradients for multiple samples to update the super model or the controller. This can stabilize the training of AutoHAS and can be naturally implemented by using parallel computation of multiple GPU/TPU devices.

Note that, in the main paper, we sometimes omit $\alpha$ in $\mathcal{W}_\alpha^*$ and $\mathcal{W}_\alpha$ for simplicity.

**Scope of AutoHAS**: AutoHAS can search for most kinds of hyperparameters, including learning rate scheduler, momentum coefficient, weight decay, dropout ratio, data augmentation policy, etc. However, it can not be used to search for weight initialization.

**Comparison with weight-sharing NAS**: AutoHAS utilized the *temporary weights* $\mathcal{W}^*$ instead of the raw weights $\mathcal{W}$ in (Pham et al., 2018; Li & Talwalkar, 2020; Dong & Yang, 2019; Zela et al., 2020) to compute the validation accuracy (loss for differentiable NAS). This accuracy over $\mathcal{W}^*$ evaluates "the performance of both architecture and hyperparameters" instead of "the performance of only architecture" in (Pham et al., 2018; Li & Talwalkar, 2020; Dong & Yang, 2019; Zela et al., 2020). In addition, AutoHAS generalizes both ENAS (Pham et al., 2018) and GDAS (Dong & Yang, 2019). ENAS can be viewed as a special case of AutoHAS that uses $\mathcal{W}$ replacing $\mathcal{W}^*$ and searches for architecture only. GDAS can be viewed as a special case of AutoHAS (Differentiable Variant) that uses $\mathcal{W}$ replacing $\mathcal{W}^*$ and searches for architecture only.

[2]See more technical details in our preliminary version: https://arxiv.org/pdf/2006.03656v1.pdf and empirical comparison in Table 5.

**How to create the basis for hyperparameters?** For the categorical hyperparameter, its basis is the set of candidate categorical values. For the continuous hyperparameter, we typically uniform-choose 10 scalars around its default value, where the default value is the commonly used value in previous literature.

**Algorithm 1** The AutoHAS Algorithm.
$\mathcal{W}$ indicates the super model's weights. $\mathcal{W}_\alpha$ indicates the weights of $\alpha$, which is a subset of $\mathcal{W}$.

**Input:** Split the available data into two disjoint sets: $\mathcal{D}_{train}$ and $\mathcal{D}_{val}$
1: Randomly initialize the super model's weights $\mathcal{W}$ and the controller's parameters $\mathcal{P}$
2: **while** not converged **do**
3: Sample $(\alpha, h \in \mathcal{B})$ from the controller
4: Compute the temporary weights $\mathcal{W}_\alpha^*$ by applying $h$ on $\mathcal{W}_\alpha$
5: Calculate the quality $Q(\alpha, h)$ as the reward to update controller by REINFORCE
6: *[Optionally]* Re-sample $(\alpha, h \in \mathcal{B})$ from the controller
7: Optimize $\mathcal{W}_\alpha$ by one step of gradient descent based on $f_h$ and $\mathcal{D}_{train}$
8: **end while**
9: Derive the final architecture $\alpha$ and hyperparameters $h$ by $\mathcal{P}$ (Sec. 2.3)

# B Related Works

We summarize the advantages of AutoHAS over other neural architecture search and hyperparameter search (aka., hyperparameter optimization) methods in Table 1.

**Neural Architecture Search.** Since the seminal works (Baker et al., 2017; Zoph & Le, 2017) show promising improvements over manually designed architectures, more efforts have been devoted to NAS. The accuracy of NAS models has been improved by carefully designed search space (Zoph et al., 2018), better search method (Real et al., 2019), or compound scaling (Tan & Le, 2019). The model size and latency have been reduced by Pareto optimization (Tan et al., 2019; Wu et al., 2019; Cai et al., 2019; 2020) and enlarged search space of neural size (Cai et al., 2020). The efficiency of NAS algorithms has been improved by weight sharing (Pham et al., 2018), differentiable optimization (Liu et al., 2019), or stochastic sampling (Dong & Yang, 2019; Xie et al., 2019). As these NAS methods use fixed hyperparameters during search, we have empirically observed that they often lead to sub-optimal results, because different architectures tend to favor their own hyperparameters. In addition, even if the manual optimization of architecture design is avoided by NAS, they still need to tune the hyperparameters after a good architecture is discovered.

**Hyperparameter Optimization (HPO).** Black-box and multi-fidelity HPO methods have a long standing history (Bergstra & Bengio, 2012; Hutter, 2009; Hutter et al., 2011; 2019; Kohavi & John, 1995). Black-box methods, e.g., grid search and random search (Bergstra & Bengio, 2012), regard the evaluation function as a black-box. They sample some hyperparameters and evaluate them one by one to find the best. Bayesian methods can make the sampling procedure in random search more efficient (Jones et al., 1998; Shahriari et al., 2015; Snoek et al., 2015). They employ a surrogate model and an acquisition function to decide which candidate to evaluate next. Multi-fidelity optimization methods accelerate the above methods by evaluating on a proxy task, e.g., using less training epochs or a subset of data (Domhan et al., 2015; Jaderberg et al., 2017; Kohavi & John, 1995; Li et al., 2017). These HPO methods are computationally expensive to search for deep learning models.

Recently, gradient-based HPO methods have shown better efficiency (Baydin et al., 2018; Lorraine et al., 2020), by computing the gradient with respect to the hyperparameters. For example, Maclaurin et al. (2015) calculate the extract gradients w.r.t. hyperparameters. Pedregosa (2016) leverages the implicit function theorem to calculate approximate hypergradient. Following that, different approximation methods have been proposed (Lorraine et al., 2020; Pedregosa, 2016; Shaban et al., 2019). Despite of their efficiency, they can only be applied to differentiable hyperparameters such as weight decay, but not non-differentiable hyperparameters, such as learning rate (Lorraine et al., 2020) or optimizer (Shaban et al., 2019). Our AutoHAS is not only as efficient as gradient-based HPO methods but also applicable to both differentiable and non-differentiable hyperparameters. Moreover, we show significant improvements on state-of-the-art models with large-scale datasets, which supplements the lack of strong empirical evidence in previous HPO methods.

| Method | Search Space | CIFAR-10 | CIFAR-100 | Stanford Cars | Oxford Flower | SUN-397 |
|---|---|---|---|---|---|---|
| MobileNetV2 (baseline) | | 94.1 | 76.3 | 83.8 | 74.0 | 46.3 |
| AutoHAS | Weight Decay | 95.0 | 77.8 | 89.0 | 84.4 | 49.1 |
| AutoHAS | MixUp | 94.1 | 77.0 | 85.2 | 79.6 | 47.4 |
| AutoHAS | Arch | 94.5 | 76.8 | 84.1 | 76.4 | 46.3 |
| AutoHAS | MixUp + Arch | 94.4 | 77.4 | 84.8 | 78.2 | 47.3 |
| AutoHAS | Weight Decay + MixUp | **95.0** (+0.9) | **78.4** (+2.1) | **89.9** (+6.1) | **84.4** (+10.4) | **50.5** (+4.2) |

Table 3: Classification top-1 accuracy (%) of AutoHAS for different search space on five datasets. Weight decay and MixUp Zhang et al. (2018) are for hyperparameters, and Arch is for architectures. Each experiment is repeated three times and the average accuracy is reported (standard deviation is about 0.2%).

Our method is also related to population based training (PBT) (Jaderberg et al., 2017; Li et al., 2019). They are asynchronous optimization algorithms that jointly optimize a population of models and their hyperparameters. Our AutoHAS is more efficient and general than PBT methods due to two factors. First, PBT needs to maintain a large number of model population in parallel, whereas AutoHAS only needs to handle a single model and its temporary weights at each iteration. Second, PBT can not be directly used for architecture search.

Recent HPO methods also pay attention to the mixture search space (Ru et al., 2020; Daxberger et al., 2020). However, due to their multi-trial nature, it is impractical to apply them for large-scale models and datasets.

**Hyperparameter and Architecture Search.** Few approaches have been developed for the joint searching of hyperparameter and architecture (Klein & Hutter, 2019; Zela et al., 2018). However, they focus on small datasets and small search spaces. These methods are more computationally expensive than AutoHAS. Concurrent to our AutoHAS, FBNet-V3 (Dai et al., 2020) learns an acquisition function to predict the performance for the pair of hyperparameter and architecture. They require to evaluate thousands of pairs to optimize this function and thus costs much more computational resources than ours.

# C Experiments

We evaluate AutoHAS on eight datasets, including two large-scale datasets, ImageNet (Deng et al., 2009) and Places365 (Zhou et al., 2017). We will briefly introduce the experimental settings in Sec. C.1. We demonstrate the generalizability of AutoHAS in Sec. C.2 and show its efficiency in Sec. C.3. Furthermore, we verify AutoHAS's scalability in Sec. C.4. Lastly, we ablatively study AutoHAS in Sec. C.5.

## C.1 Experimental Settings

**Searching settings.** We call the hyperparameters that control the behavior of AutoHAS as meta hyperparameters – the optimizer and learning rate for RL controller, the momentum ratio for RL baseline, and the warm-up ratio. Warm-upping the REINFORCE algorithm indicates that we do not update the parameters of the controller at the beginning. In addition, when the search space includes architecture choices, we also uses the warm-up technique described in Bender et al. (2020). For these meta hyperparameters, we use Adam, momentum as 0.95, warm-up ratio as 0.3. The meta learning rate is selected from $\{0.01, 0.02, 0.05, 0.1\}$ according to the validation performance. When the architecture choices are in the search space, we will use the absolute reward function (Bender et al., 2020) to constrain the FLOPs of the searched model to be the same as the baseline model. For experiments on ImageNet and Places365, we use the batch size of 4096, search for 100 epochs, and use $4\times4$ Cloud TPU V3 chips. For experiments on other datasets, we use the batch size of 512, search for 15K steps, and use Cloud TPUv3-8 chips.

**Training settings**. Once we complete the searching procedure, we re-train the model using the AutoHAS-discovered hyperparameter and architecture. For the components that are not searched for,

we keep it the same as the baseline models. For each experiment, we run three times and report the mean of the accuracy.

### C.2 Generalizability of AutoHAS

To evaluate the generalization ability of AutoHAS, we have evaluated AutoHAS in different hyperparameter and architecture spaces for five datasets. For simplicity, we choose the standard MobileNetV2 (Sandler et al., 2018) as our baseline model[3]. Table 3 shows the results. We observe that AutoHAS achieves up to 10% accuracy gain on the Flower dataset, suggesting that AutoHAS could be more useful for less optimized or new model/dataset scenarios.

### C.3 Efficiency of AutoHAS

We evaluate the efficiency of AutoHAS on ImageNet in Fig. 2a. We still choose MobileNetV2 (Sandler et al., 2018) as the baseline model. We search for the mixup ratio from [0, 0.2] and drop-path ratio from [0, 0.5] for each MBConv layer. We use the training schedule in (Bender et al., 2020). Results compared with four representative HPO methods are shown in Fig. 2a. Multi-trial search methods – Random Search (Bergstra & Bengio, 2012) or Bayesian optimization (Golovin et al., 2017) – must train and evaluate many candidates from *scratch*, and thus are inefficient. Even using 10× more time, they still cannot match the accuracy of AutoHAS. This is because AutoHAS uses weight sharing and only requires negligible cost for every new sampled candidate. AutoHAS can traverse hundreds of more samples, within the same amount of time, than the traditional Random Search and Bayesian optimization.

Comparing AutoHAS to differentiable HPO methods, HGD (Baydin et al., 2018) can only search for the learning rate and the searched learning rate is much worse than the baseline. IFT (Lorraine et al., 2020) is an efficient gradient-based HPO method. With the same search space, AutoHAS gets higher accuracy than IFT.

| Model | Method | #Params (M) | #FLOPs (M) | Top-1 Accuracy (%) |
|---|---|---|---|---|
| ResNet-50 | Human | 25.6 | 4110 | 77.20 |
| | AutoHAS | 25.6 | 4110 | **77.83** (+0.63) |
| EfficientNet-B0 | NAS | 5.3 | 398 | 77.15 |
| | AutoHAS | 5.2 | 418 | **77.92** (+0.77) |

Table 4: AutoHAS improves popular models on the ImageNet dataset.

As another proof of AutoHAS's efficiency, we follow MNasNet (Tan et al., 2019) and ProxylessNAS Cai et al. (2019) to design an architecture search space (i.e., kernel size {3x3, 5x5} and expansion ratio {3, 6} on top of MobileNetV2), and a joint search space with additional hyperparameter search options (i.e., mixup and dropout ratio). We then compare AutoHAS performance on these two search spaces. With architecture-only search, AutoHAS achieves comparable results (e.g., 73.9% accuracy @ 300M flops) as MNasNet/ProxylessNAS. With the joint search, AutoHAS can further improve accuracy by 0.2% with negligible additional cost. Notably, AutoHAS has the the same level of computational costs as efficient NAS methods, but NAS methods are infeasible to optimize the hyperparameters.

### C.4 AutoHAS Can Scale to Large Datasets

To investigate the effect of AutoHAS over the large-scale datasets, we apply AutoHAS to two popular ImageNet models. Firstly, we choose ResNet-50. The baseline strategy is to train it by 200 epochs, start the learning rate at 1.6 and decay it by 0.1 for every $\frac{1}{3}$ of the whole training procedure, use EMA with the decay rate of 0.9999, and apply SGD with the momentum of 0.9. This can provide higher accuracy than the original paper. For reference, the reported top-1 accuracy is 76.15% for ResNet-50 in TorchVision, whereas our baseline is 77.2% accuracy. Since previous methods usually do not tune the architecture of ResNet-50, we only use AutoHAS to search for its hyperparameters including the learning rate and the mixup ratio. From Table 4, AutoHAS improves this ResNet baseline by 0.63%.

[3] Baseline: the weight decay as 5e-6 and the mixup ratio as 0.

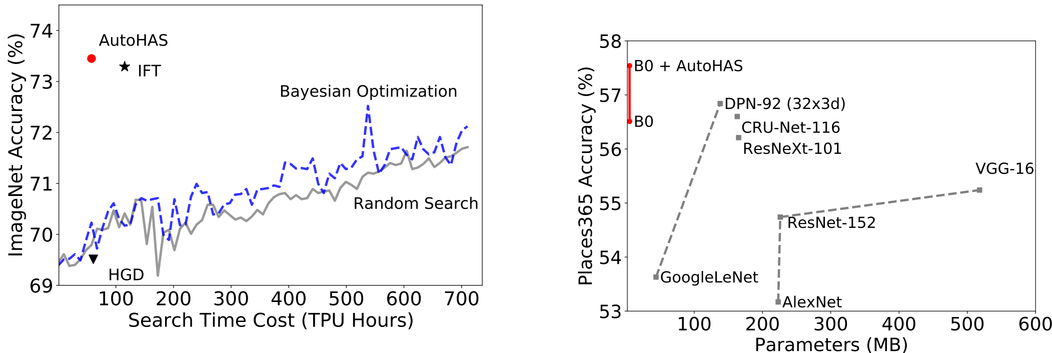


(a) Comparison between AutoHAS and previous HPO methods on ImageNet. AutoHAS achieves higher accuracy than HGD and IFT, and uses much less search time cost than others.

(b) AutoHAS improves accuracy by 1% for EfficientNet-B0 on Places365.

Figure 2: LEFT: Comparison between AutoHAS and other HPO methods. RIGHT: Comparison between AutoHAS and other popular models.

| | #Params (MB) | #FLOPs (M) | Accuracy (%) | Search Cost Memory (GB) | Search Cost Time (TPU Hour) |
|---|---|---|---|---|---|
| Baseline model | 1.5 | 35.9 | 50.96 | 1.0 | 44.8 |
| AutoHAS (Differentiable Variant) | 1.5 | 36.1 | 52.17 | 6.1 | 92.8 |
| AutoHAS (REINFORCE) | 1.5 | 36.3 | 53.01 | **1.8** | **54.4** |

Table 5: AutoHAS Differentiable Search vs. AutoHAS REINFORCE Search – Both are applied to the same baseline model with the same hyperparameter and architecture search space. Baseline model has no search cost, but we list its standalone training cost as a reference. Compared to the differentiable search, AutoHAS achieves slightly better accuracy with much less search memory cost.

Secondly, we choose a NAS-searched model, EfficientNet-B0 (Tan & Le, 2019). The baseline strategy is to train it by 600 epochs and use the same learning rate schedule as in the original paper. As EfficientNet-B0 already tunes the kernel size and expansion ratio, we choose a different architecture space. Specifically, in each MBConv layer, we search for the number of groups for all the 1-by-1 convolution layer, the number of depth-wise convolution layer, whether to use a residual branch or not. In terms of the hyperparameter space, we search for the per-layer drop-connect ratio, mixup ratio, and the learning rate. We use AutoHAS to first search for the architecture and then for the hyperparameters. From Table 4, we improves the strong EfficientNet-B0 baseline by 0.77% ImageNet top-1 accuracy.

**AutoHAS improves SoTA Places365 models.** Beside ImageNet, we have also evaluated AutoHAS on another popular dataset: Places365. Similarly, we apply AutoHAS to EfficientNet-B0 to search for better architectures and hyperparameters on this dataset. Fig. 2b shows the results: Although EfficientNet-B0 is a strong baseline with significantly better parameter-accuracy trade-offs than other models, AutoHAS can still further improve its accuracy 1% and obtain a new state-of-the-art accuracy on Places365. Note that `B0` and `B0 + AutoHAS` only uses single crop evaluation, while other models use 10 crops.

### C.5 Ablation Studies and Visualization

**Why choose RL instead of a differentiable strategy?** Differentiable search methods have been extensively studied for its simplicity in many previous literature (Liu et al., 2019; Dong & Yang, 2019; Wan et al., 2020; Xie et al., 2019), but these methods usually require much higher memory cost in order to train the entire super model. In our AutoHAS framework, we employ a simple reinforcement learning algorithm – REINFORCE (Williams, 1992) – to optimize the controller: instead of training the whole super model, we only train a subset of the super model and therefore significantly reduce the training memory cost. Notably, the REINFORCE could also be replaced by a differentiable-based algorithm with the supervision of validation loss. We investigate the difference between differentiable

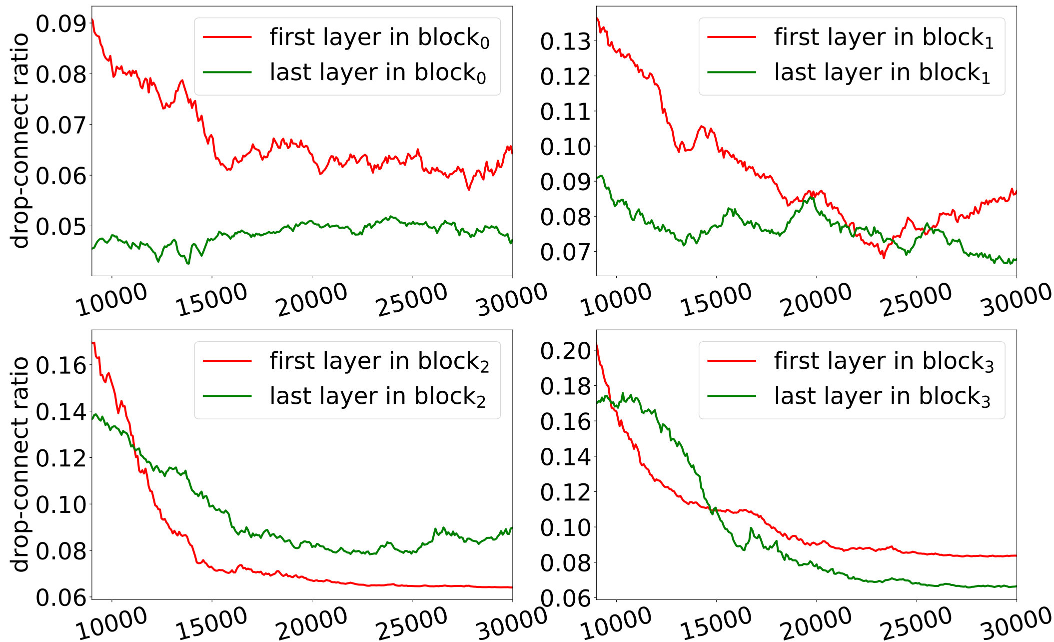


Figure 3: We visualize the AutoHAS-discovered drop-connect ratio for some layers in EfficientNet-B0. The x-axis and y-axis indicate the search step and drop-connect ratio, respectively.

and REINFORCE search in Table 5. Not surprisingly, differentiable search requires much higher memory cost (6.1x more than baseline) as it needs to maintain the feature or gradient tensors for all the super model, whereas our REINFORMCE-based AutoHAS is much more memory efficient: reducing the memory cost by 70% than the differentiable approach. Empirically, we observe they achieve similar accuracy gains in this case, but AutoHAS enables us to search for much larger models such as EfficientNet-B0 and ResNet-50 as shown in Table 4.

**Visualization**. We show the intermediate search results of drop-connect ratios in Fig. 3. Human experts have prior knowledge that drop-connect is crucial to the performance. However, it is prohibitive to manually tune a proper ratio for each of those 10+ drop-connect layers in EfficientNet-B0. Our AutoHAS is suitable for such a scenario, it can discover suitable ratios for many drop-connect layers in a single trial.

## D Conclusion & Future Work

In this paper, we proposed an automated and unified framework AutoHAS, which can efficiently search for hyperparameters and architectures. AutoHAS provides a novel perspective of AutoML algorithms by generalizing the weight sharing technique from architectures to hyperparameters. AutoHAS integrates several techniques to be memory friendly, efficient, generalized, and scalable. Experimentally, AutoHAS improves the baseline models on *eight* datasets over different kinds of search spaces. For the highly-optimized ResNet/EfficientNet, it improves ImageNet top-1 accuracy by 0.8%; for other less-optimized scenarios (e.g., Oxford Flower), it improves the accuracy by 11.4%.

In the future, we would evaluate AutoHAS on NAS or HAS benchmarks. In addition, we would comprehensively investigate AutoHAS's performance by integrating different searching algorithms in AutoHAS.