HELP: Hardware-Adaptive Efficient Latency Predictor for NAS via Meta-Learning

Anonymous Author(s)
Affiliation
Address
email

Abstract

For deployment, neural architecture search should be hardware-aware, in order to satisfy the device-specific constraints (e.g., memory usage, latency and energy consumption) and enhance the model efficiency. Existing methods on hardware-aware NAS collect a large number of samples (e.g., accuracy and latency) from a target device, either builds a lookup table or a latency estimator. However, such approach is impractical in real-world scenarios as there exist numerous devices with different hardware specifications, and collecting samples from such a large number of devices will require prohibitive computational and monetary cost. To overcome such limitations, we propose Hardware-adaptive Efficient Latency Predictor (HELP), which formulates the device-specific latency estimation problem as a meta-learning problem, such that we can estimate the latency of a model's performance for a given task on an unseen device with a few samples. To this end, we introduce novel hardware embeddings to embed any devices considering them as black-box functions that output latencies, and meta-learn the hardware-adaptive latency predictor in a device-dependent manner, using the hardware embeddings. We validate the proposed HELP for its latency estimation performance on unseen platforms, on which it achieves high estimation performance with as few as 10 measurement samples, outperforming all relevant baselines. We also validate end-to-end NAS frameworks using HELP against ones without it, and show that it largely reduces the total time cost of the base NAS method, in latency-constrained settings.

1 Introduction

Neural Architecture Search (NAS) [37, 1, 38, 24, 21, 20, 34, 5], which aims to search for the optimal architecture for a given task, has achieved a huge practical success by overcoming the sub-optimality of manual architecture designs. However, for NAS to be truly practical in real-world scenarios, it should be hardware-aware. That is, we need to search for the neural architectures that satisfy diverse device constraints (e.g., memory footprint, inference latency, and energy consumption). Due to the practical importance of the problem, many existing works propose to take into account the hardware efficiency constraints (mostly latency) in the search process [28, 3, 33, 31, 6, 2, 35, 4, 32].

However, one of the main challenges with such hardware-aware NAS is that collecting training samples (e.g., architecture-latency pairs on each target device) to build reliable prediction models for the efficiency metrics, is computationally costly or requires the knowledge of the hardware devices. Existing hardware-aware NAS methods usually require a large number of samples (e.g., 5k) and train metric predictors for each device from scratch. Additionally, most works [28, 3, 33, 31, 6] are task-specific, and thus such collection of performance samples should be done from scratch, for a new task. Thus, the sample collection process is prohibitively costly when we consider real-world deployment of a model to a large number of hardware devices, for any tasks. OFA [4] alleviate the

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

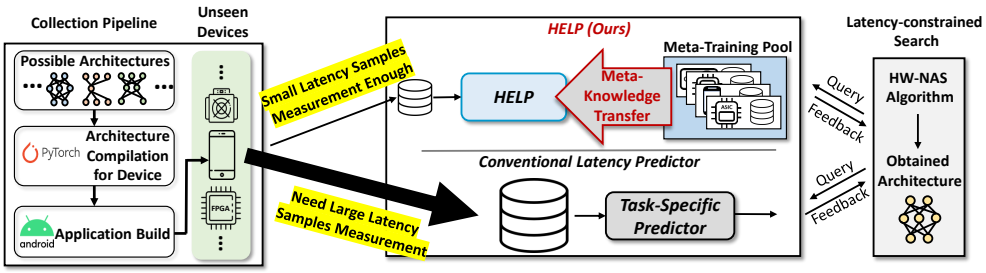

Figure 1: **Concept.** Conventional latency estimation methods require a large number of architecture-latency pairs to build a prediction model separately for each device, which is inefficient. Contrarily, the proposed HELP uses a single meta-latency predictor that can fast adapt to any unknown device by collecting only a few latency measurements from it, by utilizing the meta-knowledge of the source device pool.

search cost by utilizing a high-performance network trained on ImageNet as the supernet, which does not require training the models from scratch. Yet, they are still sub-optimal since building device-specific predictors for metric still requires a large number of samples to be collected, to build a layer-wise latency predictor for each device. This could take multiple hours depending on the task and the device, and becomes a bottleneck for latency-constrained NAS. BigNAS [36] considers FLOP as the efficiency metric, but this is a highly inaccurate proxy since latency of an architecture could differ based on its degree of parallelism and memory access cost, for the same FLOP.

To overcome such limitations, we propose a novel sample-efficient latency predictor, namely Hardware-adaptive Efficient Latency Predictor (HELP), which supports the latency estimation for multiple devices with a **single** model, by allowing it to rapidly adapt to unseen devices with only a few latency measurements collected from each device. Specifically, we formulate the latency prediction problem as a few-shot regression task of estimating the latency given an architecture-device pair, and propose a novel hardware embedding that can embed any devices by utilizing the latencies of the reference architectures on each device as its embeddings. Then, we propose a meta-learning framework which combines amortized meta-learning with gradient-based meta-learning, to learn the latency predictor to generalize across multiple devices, utilizing the proposed hardware embeddings. This allows the model to transfer knowledge learned from known devices to a new device, and thus to achieve high sample efficiency when estimating the latency of unseen devices.

HELP is highly versatile and general, as it is applicable to any hardware platforms and architecture search spaces, thanks to our device-agnostic hardware embeddings. Also, HELP can be coupled with any NAS frameworks to reduce its computational bottleneck in obtaining latency-constrained architectures. Especially, when coupled with rapid NAS methods such as MetaD2A [15] and OFA [4], HELP can perform the entire latency-constrained NAS process for a new device almost instantly. We validate the latency estimation performance of HELP on the NAS-Bench-201 space [8] with various devices from different hardware platforms, utilizing the latency dataset for an extensive device pool in HW-NAS-Bench dataset [17]. The results show that our meta-learned predictor successfully generalize to unseen target devices, largely outperforming baseline latency estimation methods using at least $90\times$ less measurements. Then, we combine HELP with existing NAS methods [4, 32] and show that our meta-latency predictor largely reduce their total search cost. To summarize, the main contributions of this paper are as follows:

- We formulate the latency estimation of a neural architecture for a given device as a few-shot regression problem, which outputs the latency given an architecture-device pair as the input.
- To represent heterogeneous hardware devices, we propose a novel device-agnostic hardware embedding, which embeds a device by its latencies on reference architectures.
- We propose a novel latency predictor, HELP, which meta-learns a few-shot regression model to generalize across hardware devices, that can estimate the latency of an unseen device using only few measurements. HELP obtains significantly higher latency estimation performance over baselines with minimal total time cost.
- We further combine HELP with existing NAS frameworks to show that it leads to find latency-constrained architectures extremely fast, eliminating the latency-estimation bottleneck in the hardware-constrained NAS.

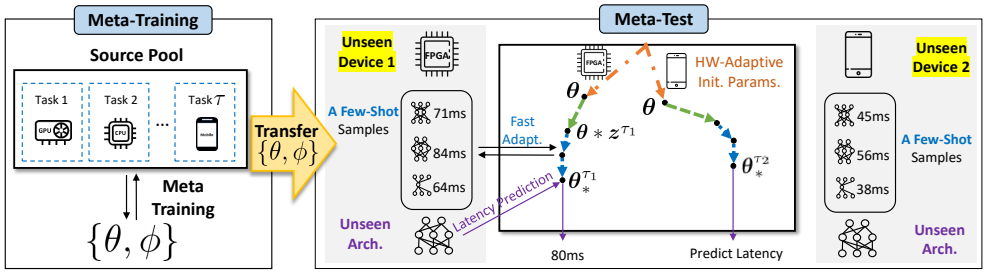

Figure 2: **Overview.** For the hardware-adaptive latency estimation of an unseen device for latency-constrained NAS, we introduce a latency-based hardware embedding and a $\boldsymbol{z}$ modulator of the initial parameters. By formulating the sample-efficient NAS problem as a few-shot regression problem under the meta-learning framework, our meta-learned predictor successfully exploits meta-knowledge $\boldsymbol{\theta}$ from the source device pool, to achieve high sample efficiency on unseen devices.

2 Related Work

Latency Prediction in NAS Hardware-aware NAS [28, 3, 4, 33, 31, 12, 32, 25] aims to design neural architectures that achieves a good trade-off between accuracy and latency for efficient deployment to target devices. While they need to consider the actual latencies of the architectures on a target device, evaluating them while searching for architectures is costly. Thus, most existing works replace it with proxies, such as FLOPs [36], but they are inaccurate measure of latencies. Earlier hardware-aware NAS methods have been used a layer-wise latency predictor (lookup table) [28, 3, 4, 33] which sums up the latencies measured for each operation in the given neural networks. While this is better than FLOPs, they do not accurately capture the complexities of multiple layer execution on real hardware devices. Thus, recent methods use an end-to-end latency predictor [9, 32] that is trained with the latency measurements from the target device, to improve the latency prediction accuracy. BRP-NAS [9], which is a GCN-based model, is currently the state-of-the-art end-to-end latency predictor. However, this method is also limited in that it requires a large number of architecture-latency pairs for each device. Since the latency estimator cannot generalize to a new device, whenever a new device is given, the NAS system needs to build a new latency estimator, which may take hours to finish. The proposed method significantly reduces the total building cost of the latency predictor for each device, by using a single meta-learned predictor and adapting it to a new device with only a few measurements from it.

Meta-learning and Meta-NAS Meta-learning (learning to learn) [29] aims to learn a model that generalizes over a distribution of tasks rather than a single task, such that the model meta-learned over a large number of tasks rapidly adapts to an unseen task. The performance of existing methods [30, 27, 11, 23, 16] is usually evaluated on few-shot classification tasks, where the model classifies between instances of unseen classes given only a few training examples. Recently, several works have proposed to utilize meta-learning for NAS. Most of them focus on few-shot classification tasks to search for the architectures and parameters that can generalize well to a new task [10, 19, 26] with gradient-based meta-learning. However, they have limited practical applicability since the computational cost of meta-learning is prohibitively large, for NAS under a standard many-shot setting. A recently proposed meta-NAS framework with amortized meta-learning, MetaD2A [15], which utilizes a set encoder to encode a task and uses a graph decoder to generate a task-adaptive architecture, obtained state-of-the-art performance on unseen datasets, with minimal search cost. We also propose a similar amortized meta-learning framework, for hardware-adaptive NAS, based on a novel hardware device embedding. However, after obtaining the device-conditioned initialization parameters, we perform further inner gradient steps for device-specific adaptation, unlike MetaD2A [15].

3 Method

Our goal is to design a prediction model that can accurately predict the efficiency metric for a novel architecture-device pair, using a small number of performance samples from the device. While our method is generally applicable to any efficiency metrics that can be measured from the device (e.g. latency, energy consumption, and memory footprint) we focus on the latency prediction in this work.

3.1 Problem Definition

Assume that we are given a task specification $\tau = \{h^\tau, \mathbf{X}^\tau, \mathbf{Y}^\tau\}$ where $h^\tau \in \mathcal{H}$ is a hardware device, $\mathbf{X}^\tau \subset \mathcal{X}$ is a set of neural architectures, and $\mathbf{Y}^\tau \subset \mathcal{Y}$ is a set of latencies of $\mathbf{X}^\tau$ directly measured on the hardware device h^τ. Then, our goal is to learn a regression model $f(x;\boldsymbol{\theta}) : \mathcal{X} \rightarrow \mathbb{R}$ parameterized by $\boldsymbol{\theta}$ that estimates the latency $y \in \mathbf{Y}^\tau$ of a neural architecture $x \in \mathbf{X}^\tau$ for a given hardware device h^τ by minimizing empirical loss $\mathcal{L}$ (e.g. mean squared error) on the predicted values $f(\mathbf{X}^\tau;\boldsymbol{\theta})$ and actual measurements $\mathbf{Y}^\tau$ as follows:

$$\min_{\boldsymbol{\theta}} \mathcal{L}\big(f^\tau\big(\mathbf{X}^\tau;\boldsymbol{\theta}\big), \mathbf{Y}^\tau\big) \tag{1}$$

Learning such a regression model seems like a simple problem, since we can collect any number of measurements from any devices. However, this is a more challenging problem than it seems:

1. Since we cannot generalize across devices, we need to learn N predictors $\{f^\tau(\cdot;\boldsymbol{\theta}^\tau)\}_{\tau=1}^N$ for N devices, collecting performance samples and training the performance predictor separately for each device, which requires prohibitive computational costs with large number of devices.
2. Even when assuming that we learn a device-specific predictor, in order not to overfit the regression model, we need to collect *a large number of* architecture-latency sample pairs for each device (e.g. 2k [32], 5k [22] samples) to achieve reliable prediction performance.
3. With lack of generalization ability across devices and architectures, the NAS framework needs to repeat the time-consuming sample collection process whenever a new device is given, that may take hours to complete. This will become a computational bottleneck even for a rapid meta-NAS framework such as MetaD2A [15].

To overcome such limitations, we propose a **single** predictor $f(\cdot;\boldsymbol{\theta})$ which can generalize across devices and architectures, by **fast** adapting to a new target device and architecture that are unseen during training, by collecting only a **small number** of architecture-latency pairs from the device ($X^\tau \ll \mathbf{X}^\tau$, $Y^\tau \ll \mathbf{Y}^\tau$). We achieve this goal with a meta-learning framework that can transfer knowledge obtained from the device and architecture pool $p(\tau)$.

3.2 Hardware-adaptive Latency Prediction with Device Embeddings

While the measured latency $y \leftarrow (x, h)$ is dependent on both the device type h and the architecture x, existing latency prediction models takes the form of $f(x;\boldsymbol{\theta})$, ignoring the device-specific constraints, since the latency predictor is learned for each device separately. This results in poor performance when learning a single latency predictor to perform metric estimation on multiple devices, including unseen ones, which is our main objective. Thus, we propose hardware-conditioned prediction model:

$$f(x, h;\boldsymbol{\theta}) : \mathcal{X} \times \mathcal{H} \rightarrow \mathbb{R} \tag{2}$$

that can predict the latency differently depending on the device type h, even for the same architecture x. A crucial challenge of our hardware-conditioned prediction model is how to represent the hardware device h, for all devices regardless of their platform types. This is not a trivial problem since the physical architecture of hardware devices could be very different (e.g. CPU vs FPGA). Thus, we simply consider the device as a black box function which outputs the inference latency given an architecture, instead of directly modeling the hardware devices. Then, we obtain the latencies of the device on a fixed set of reference neural architectures as follows:

$$V_h = \{y_1^*(x_1, h), y_2^*(x_2, h), ..., y_d^*(x_d, h)\} : \mathcal{E} \times \mathcal{H} \rightarrow \mathbb{R}^d \tag{3}$$

where $\mathcal{E}$ is the set of the reference neural architectures $\{x_1, x_2, ..., x_d\} \subset \mathcal{X}$, fixed across all tasks for both meta-training and meta-test, and d is the number of the reference architectures; in our experiments, we set $d = 10$. Further, $y_i^*(x_i, h) = \{y_i(x_i, h) - min(V_h^{(0)})\}/\{max(V_h^{(0)}) - min(V_h^{(0)})\}$ are standardized latency values ranging from 0 to 1, where $V_h^{(0)} = \{y_1(x_1, h), y_2(x_2, h), ..., y_d(x_d, h)\}$. Since this set of reference devices should be representative, we select them to be diverse and heterogeneous. As for the reference architectures, we randomly sample them from the search space. For more detailed descriptions of the reference devices and architectures, please see the supplementary file. The proposed black-box treatment of hardware devices, and the latency-based hardware embedding allows us to embed a new device without considering its detailed hardware specification.

3.3 Meta-Learning the Hardware-adaptive Latency Predictor

To tackle the few-shot regression problem for multiple devices by utilizing the collected pool of devices and architectures $p(\tau)$, we propose a novel hardware-adaptive meta-learning framework of the latency predictor that meta-learns $f(x, h; \boldsymbol{\theta})$ across a task distribution $p(\tau)$ to rapidly adapt the predictor $f(x, h; \boldsymbol{\theta}^{\tau})$ to an unseen neural architecture x given the task specification $\tau = \{h^{\tau}, V_h, \mathbf{X}^{\tau}, \mathbf{Y}^{\tau}\}$. During the meta-training phase, we leverage the episodic training strategy, in which we simulate a few-shot regression task for each iteration by randomly sampling task τ from the device-architecture pool $p(\tau)$ and splitting τ as training set $\mathcal{D} = \{h^{\tau}, X^{\tau}, Y^{\tau}\}$ and test set $\tilde{\mathcal{D}} = \{h^{\tau}, \tilde{X}^{\tau}, \tilde{Y}^{\tau}\}$ where $X^{\tau} \subset \mathbf{X}$ and $\tilde{X}^{\tau} \subset \mathbf{X}$ denote sets of neural architecture samples, X^{τ} is the set of few-shot samples $|X^{\tau}| \ll |\mathbf{X}|$. Note that there is no overlap between them; that is, $X^{\tau} \cap \tilde{X}^{\tau} = \emptyset$. $Y^{\tau} \subset \mathbf{Y}^{\tau}$ and $\tilde{Y}^{\tau} \subset \mathbf{Y}^{\tau}$ denote the sets of corresponding latency values of neural architectures X^{τ} and $\tilde{X}^{\tau}$, measured on device h^{τ}, respectively. Basically, for each task τ, we obtain the hardware-adaptive prediction model $f(X, V_h^{\tau}; \boldsymbol{\theta}^{\tau})$ as a function of V_h^{τ}. Formally, we meta-train the latency predictor to minimize the test loss $\mathcal{L}(\cdot; \tilde{\mathcal{D}}^{\tau})$ by optimizing the following task-adaptive meta-learning objective:

$$\min_{\boldsymbol{\theta}} \sum_{\tau \sim p(\tau)} \mathcal{L}\big(f\big(\tilde{X}^{\tau}, V_h^{\tau}; \boldsymbol{\theta}^{\tau}\big), \tilde{Y}^{\tau}\big) \tag{4}$$

We can simply use the task embedding V_h^{τ} to obtain a task-conditioned latency predictor, in which case we are using an amortized meta-learning framework similar to one proposed in Lee et al. [15], which aims to meta-learn a dataset-adaptive performance predictor and a NAS framework. However, we can further perform few-shot adaptation with the few latency measurements collected from the target device, by conducting inner gradient updates with them as follows:

$$\boldsymbol{\theta}^{\tau}_{(t+1)} = \boldsymbol{\theta}^{\tau}_{(t)} - \alpha \nabla_{\boldsymbol{\theta}_{(t)}} \mathcal{L}(f(X^{\tau}, V_h^{\tau}; \boldsymbol{\theta}_{(t)}), Y^{\tau})) \text{for } \ t = 1, \ldots, T \tag{5}$$

where t denotes the t_{th} inner gradient step, T is the total number of inner gradient steps, and α denotes the multi-dimensional global learning rate vector [18]. This meta-learning formulation allows us to adapt to a new device rapidly, by using the knowledge of the devices used for meta-training. However, when we are encountered with a completely new device that has little relatedness to any devices from the meta-device pool, it will be helpful to deviate more from the meta-knowledge captured by $\theta^{\tau}_{(0)}$ [14]. To this end, we introduce a hardware-adaptive modulator $\boldsymbol{z}^{\tau} = g(V_h^{\tau}; \boldsymbol{\phi}) : \mathbb{R}^{d} \rightarrow \mathbb{R}^{d_{\boldsymbol{\theta}}}$ parameterized by ϕ to modulate the initial parameter as $\boldsymbol{\theta}_{(0)} = \boldsymbol{\theta} * \boldsymbol{z}^{\tau}$, where $\boldsymbol{\theta}_{(0)}$ is the new initialization for hardware h^{τ}. Following [14], we set $\boldsymbol{\theta}_{(0)} \leftarrow \boldsymbol{\theta} \circ \boldsymbol{z}^{\tau}$ for weights and $\boldsymbol{\theta}_{(0)} \leftarrow \boldsymbol{\theta} + \boldsymbol{z}^{\tau}$ for biases with an element-wise multiplication operator $\circ$. This leads to the following update rule:

$$\boldsymbol{\theta}^{\tau}_{(0)} = \boldsymbol{\theta} * \boldsymbol{z}^{\tau} \tag{6}$$

$$\boldsymbol{\theta}^{\tau}_{(t+1)} = \boldsymbol{\theta}^{\tau}_{(t)} - \alpha \nabla_{\boldsymbol{\theta}^{\tau}_{(t)}} \mathcal{L}(f(X^{\tau}, V_h^{\tau}; \boldsymbol{\theta}^{\tau}_{(t)}), Y^{\tau})) \quad \text{for } \ t = 1, \ldots, T \tag{7}$$

where T is the number of inner gradient steps. Then, the final meta-learning objective is as follows:

$$\min_{\boldsymbol{\theta}, \boldsymbol{\phi}} \sum_{\tau \sim p(\tau)} \mathcal{L}\big(f\big(\tilde{X}^{\tau}, V_h^{\tau}; \boldsymbol{\theta}^{\tau}_{(T+1)}\big), \tilde{Y}^{\tau}\big) \tag{8}$$

Thus, we meta-learn both the model parameters for the hardware-adaptive latency predictor, and the modulator for the shared initial parameters.

Few-shot Adaptation to Unseen Devices (Meta-Test) We now describe how to use our meta-learned latency prediction model $f(\cdot; \boldsymbol{\theta})$ to estimate the latency of an architecture on a novel device h^{υ}that is unseen during meta-training. The task-specific predictors [32, 4, 22] need a large amount of latency measurements from an unseen device h^{υ}, over diverse architectures in order not to overfit, which may take an excessive time to collect. However, our model is able to measure the latency values y^{υ} of a new architecture $\tilde{x}^{\upsilon}$ by collecting only few latency measurements from it (we use 10 or 20), using the device-conditioned meta-learning. Given a latency prediction task of an architecture $\tilde{x}^{\upsilon}$ on a novel device h^{υ}, $\upsilon = \{h^{\upsilon}, X^{\upsilon}\}$, we first obtain its hardware device embedding V_h by obtaining its latencies on a fixed set of reference architectures, following Equation (3). Then we use the device embedding V_h to obtain the device-optimized parameters of the latency predictor $\boldsymbol{\theta}^{\upsilon}_{(T+1)}$, following Equation (6) and (7). Then, we use the device-optimized latency predictor $f(\cdot, V_{h^u}; \boldsymbol{\theta}^{\upsilon}_{(T+1)})$ to measure the latency of $\tilde{x}^{\upsilon}$. We can further combine this meta-latency predictor with a NAS method to perform latency-constrained NAS for a novel device.

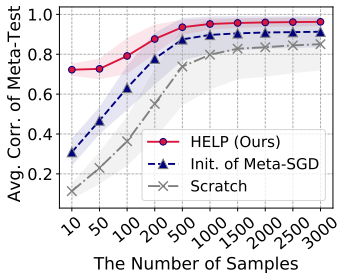

Figure 3: Latency estimation performance as a function of the number of samples collected.

Method		Unseen Platform		
		Raspi4	ASIC	FPGA
MAML [11]		0.568	0.602	0.541
ANP [13]		0.801	0.657	0.884
Meta-SGD [18]		0.844	0.831	0.882
HELP (Ours)	Amortization	0.568	0.604	0.539
	+ HW-Condition	0.853	0.904	0.861
	+ Few-Shot Adapt	0.872	0.913	0.866
	+ z Modulator	**0.885**	**0.942**	**0.889**

Table 2: The correlation of the estimated latency with HELP to the actual latency, on unseen devices with 10 measurement samples from each device.

Method	Architecture Search Cost	Latency Estimator Building Cost
Task-specific NAS	O(D)	O(N)
MetaD2A [15]	O(1)	O(N)
MetaD2A + HELP	**O(1)**	**O(1)**

Table 1: Cost of NAS and latency estimation.

Computational Complexity of HELP. The meta-training of the latency predictor is done only once, and once done, we can adapt the meta-latency estimator for the latency estimation of **any number of devices**. Since conventional approaches require to collect a large number of samples and train a device-specific latency estimator for each target device, while HELP only needs to collect 10 samples per device, HELP reduces the time complexity of obtaining latency estimations on N devices from $O(N)$ to $O(1)$.

4 Experiment

In this section, we first verify the efficacy of our meta-learning scheme on the latency prediction of architectures on unseen devices, in Section 4.1. In this section, we also validate the sample-efficiency and the performance of HELP against existing latency prediction models, and a predictor trained with conventional meta-learning. Then, in Section 4.2, we validate HELP's effectiveness and efficiency on the full latency-constrained NAS for novel devices, by combining it with existing NAS methods.

Search Space. Following the evaluation procedure of HW-NAS-Bench [17], we consider two search spaces, **NAS-Bench-201** [8] and **FBNet** [33]. Additionally, we consider **MobileNetV3** [12, 4] search space for end-to-end NAS in Section 4.2. For detailed description, refer to the supplementary file.

Meta-Training Pool/Meta-Test Pool. To construct the **Meta-Training Pool**, we collect the latency measurements from **18 heterogeneous devices**, including GPUs, CPUs, mobile devices (NVIDIA 1080ti, Titan X, Titan XP, RTX 2080ti, Xeon Silver 4114, Silver 4210r, Samsung A50, S7, Pixel3, Essential Ph 1). For GPUs, we consider three different batch sizes [1, 32, 256(64)] and for all other hardware devices, we use the batch size of 1. We collect 900/4000 (architecture, latency) pairs of each training device for NAS-Bench-201 and FBNet search space, respectively. As for the **Meta-Test Pool**, we consider **Unseen Devices** and **Unseen Platforms**. Unseen Devices include NVIDIA GPU Titan RTX, Intel CPU Xeon Gold 6226, and Pixel2, which are different from the devices in the meta-training pool but belong to the same categories (GPU, CPU, mobile device). On the other hand, Unseen Platforms include Raspi4, ASIC-Eyeriss, and FPGA, which are completely unseen categories of devices. For Unseen Platforms and Pixel2, we use the latency measurements provided in the HW-NAS-Bench [17]. For the implementation details of our model, please refer to the Suppl.

Baselines. We compare our framework against relevant baselines. 1) **MAML** [11]: A few-shot regression baseline which meta-learns the initial parameters over multiple tasks via bi-level optimization. 2) **Meta-SGD** [18]: An extension of MAML with the meta-learned learning rate for the inner-gradient step. 3) **ANP** [13]: A few-shot regression model implemented with Attentional Neural Processes, which uses differentiable attentions to attend to the relevant contexts for the given query. 4) **BRP-NAS** [9]: A predictor-based NAS method with a graph convolution neural network-based latency predictor. This baseline achieves the previous state-of-the-art performance on latency prediction in the NAS-Bench-201 search space. 5) **MetaD2A** [15]: This is a meta-NAS framework without a latency predictor which enables rapid architecture search (a few GPU seconds) for unseen datasets, which obtains state-of-the-art performance on mulitple datasets in the NAS-Bench-201 search space.

4.1 Efficacy of HELP on Few-shot Latency Estimation for Novel Devices

We first validate whether the transferring the meta-knowledge obtained over a meta-training pool to an unseen meta-test device helps to improve the sample efficiency and prediction performance of the latency predictor. We adapt the meta-latency predictor on 10 architecture-latency pairs of 6

Table 3: Comparison of the latency estimators on unseen devices and platforms on NAS-Bench-201.

Method	Transfer	Sample	Unseen Device			Unseen Platform			Mean
			GPU	CPU	Pixel2	Raspi4	ASIC	FPGA	
FLOPS		-	0.950	0.826	0.865	0.846	0.437	0.900	0.795
Layer-wise Predictor		-	0.667	0.866	-	-	-	-	0.767
BRP-NAS [9]		900	0.814	0.796	0.757	0.847	0.811	0.801	0.804
BRP-NAS(+extra samples)		3200	0.822	0.805	0.793	0.853	0.830	0.828	0.822
HELP (Ours)	✓	**10**	**0.996**	**0.983**	**0.884**	**0.873**	**0.971**	**0.987**	**0.949**

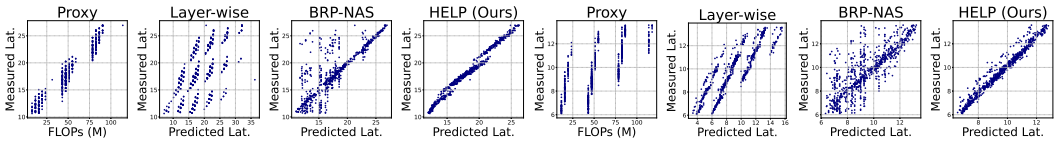

(a) Titan RTX GPU (b) Intel Xeon Gold CPU

Figure 4: Comparison of the **estimated** and **measured** latencies on a Titan RTX GPU and Intel Xeon Gold CPU. While BRP-NAS requires 900 samples to train the latency predictor, our meta-latency predictor requires only 10 samples for adaptation, and significantly outperforms it in the estimation performance.

unseen meta-test devices, and report the Spearman's rank correlation (higher the better) between the estimated latencies and actual latencies on 1,000 neural architectures from the test set in FBNet search space, over 5 random runs. (Figure 3, Table 2, and Figure 5).

Latency Estimation Performance for Unseen Devices. Figure 3 reports the average value of correlation scores of different predictors as a function of the number of architecture-latency pairs from meta-test devices. HELP and Meta-SGD uses the initial parameters that is meta-learned over a large meta-training pool, and fine-tunes its parameter with given training samples. On the other hand, Scratch means to simply train a regression model from scratch with a given samples. The result shows that using meta-knowledge (HELP and Init. of Meta-SGD) consistently outperforms the model trained from scratch. Specifically, HELP's latency predictor with hardware-adaptive initial parameter $\boldsymbol{\theta}_0$ achieves significantly larger performance gain over baselines when the number of samples is smaller (e.g., 10 and 50). Such sample-efficiency allows HELP to search for architectures that satisfy the latency constraints with orders of magnitudes shorter time compared to existing methods.

Effect of the Hardware-adaptive Meta-learning. We analyze the effect of hardware-adaptive meta-learning in Table 2. We observe that the meta-latency predictor using the proposed modules largely outperforms the other ones trained with other meta-learning baselines that are hardware-independent, which shows the effectiveness of our hardware-adaptive meta-learning on unseen platforms. This is due to the heterogeneity of the tasks (devices) in the meta-training dataset, in which the task-conditioning becomes more important.

Effect of the Size of the Device Pool. We further analyze the effect of the size of the meta-training pool, on the performance of the meta-latency predictor. Figure 5 reports the performance of our latency predictor over different sizes of the randomly sampled device pool. In particular, when the number of devices in the meta-training pool is 10 or more, our model achieves over 0.9 correlation on unseen devices using only 10 samples of the unseen target device, regardless of the device types of the meta-training pool. Contrarily, Meta-SGD does not yield meaningful performance gains even with large meta-training pools.

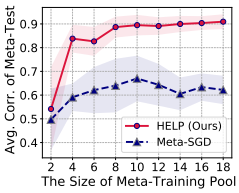

Figure 5: Effect of the meta-training pool size.

Sample-efficiency of HELP To demonstrate the sample-efficiency of our meta-learned latency predictor, we compare against three baselines: 1) a proxy predictor using number of FLOPs; 2) a layer-wise predictor and 3) the latency predictor from BRP-NAS [9] (Table 3 and Figure 4). The results show that FLOPs, although easy to compute, is an inaccurate proxy for latency estimation. The layer-wise predictor also achieves poor performance, because it cannot reflect the complexity and the holistic effect of the network architecture. The latency predictor from BRP-NAS [9], which is a 4-layer GCN with 600 hidden units followed by a fully connected layer to produce a scalar output, achieves significantly better estimation compared to the first two baselines. However, this model requires 900 samples from each architecture-device pair, as decribed in [9]. Finally, our HELP predictor achieves the best performance, achieving the Spearman's rank correlations of 0.996 on GPU

Table 4: Performance comparison of different latency estimators combined with MetaD2A for latency-constrained NAS, on CIFAR-100 dataset with NAS-Bench-201 search space. We use Google Pixel2 phone (top) and NVIDIA Titan RTX GPU (bottom) as the target devices. We exclude the layer-wise predictor for Pixel2 since we use the latency measurements from HW-NAS-Bench [17] for this device, and it does not provide block information for the layer-wise predictor to use.

Device	Model	Const (ms)	Latency (ms)	Accuracy (%)	MACs (M)	Latency Model Sample	Latency Model Building Time	Total NAS Cost (Wall Clock)	Speed Up
Pixel2	MetaD2A + BRP-NAS [9]	14	14	66.9	79	900	1120s	1220s	1.0×
	MetaD2A + HELP (Ours)		13	67.4	47	**10**	**25s**	**125s**	**9.8**×
	MetaD2A + BRP-NAS [9]	22	34	73.5	185	900	1120s	1220s	1.0×
	MetaD2A + HELP (Ours)		19	70.6	55	**10**	**25s**	**125s**	**9.8**×
	MetaD2A + BRP-NAS [9]	34	34	73.5	185	900	1120s	1220s	1.0×
	MetaD2A + HELP (Ours)		34	73.5	185	**10**	**25s**	**125s**	**9.8**×
Titan RTX (Batch 256)	MetaD2A + Layer-wise Pred.	18	37	73.2	121	900	998s	1098s	1.0×
	MetaD2A + BRP-NAS [9]		21	67.0	86	900	940s	1040s	1.1×
	MetaD2A + HELP (Ours)		18	69.3	51	**10**	**11s**	**111s**	**9.9** ×
	MetaD2A + Layer-wise Pred.	21	41	73.5	184	900	998s	1098s	1.0×
	MetaD2A + BRP-NAS [9]		19	71.5	55	900	940s	1040s	1.1×
	MetaD2A + HELP (Ours)		19	71.6	55	**10**	**11s**	**111s**	**9.9**×
	MetaD2A + Layer-wise Pred.	25	41	73.5	184	900	998s	1098s	1.0×
	MetaD2A + BRP-NAS [9]		23	70.7	82	900	940s	1040s	1.1×
	MetaD2A + HELP (Ours)		25	71.8	86	**10**	**11s**	**111s**	**9.9**×

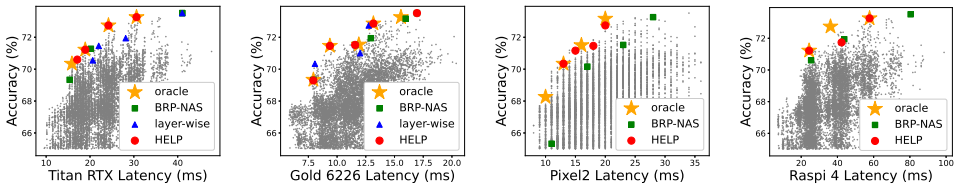

Figure 6: The accuracy-latency trade-off of the NAS framework with an oracle accuracy predictor, combined with the oracle latency estimator (yellow star), BRP-NAS (green square), layer-wise predictor (blue triangle), and HELP (Ours - red circle), on various devices in NAS-Bench-201 space. HELP, with its accurate latency estimation, obtains Pareto-frontier models, while baselines yield sub-optimal architectures.

and 0.983 on CPU, using only 10 latency measurements of the architecture on each device. This shows the clear advantage of our method, in terms of the estimation accuracy and sample-efficiency.

4.2 End-to-end Latency-constrained NAS with HELP

To show that HELP does help NAS frameworks rapidly obtain latency-constrained/optimal architectures for a novel device, we combine HELP with existing NAS methods, namely MetaD2A [15] and OFA [4], and validate the performance on the latency-constrained NAS tasks.

In Table 4 and Table 5, besides latency and accuracy, we additionally report the time-efficiency of the latency estimators with three different measures. First, we use the **number of samples** which are the number of architecture-latency pairs obtained from the device, that are used to build or train the latency estimator. We also report the **building time**, which is the total wall clock time to build the latency estimator, including the time required for sample collection, architecture compilation on the target devices, transmitting the architecture to the target device, and measuring the latency on the device. Finally, we report **the total NAS cost**, which is the sum of both the estimator building time and the architecture search time, on a target task. After obtaining the architecture, we measure the actual latency of the architecture on the target device and report it as **latency** (ms). For the building time and the total NAS cost, we exclude the cost of any procedures that are not done during the meta-test time, such as the meta-training of MetaD2A model (46 GPU hours) and HELP (25/18 hours for MetaD2A and OFA), as well as the time to train the supernet for OFA (1,200 GPU hours).

We first combine HELP with MetaD2A and compare it with the MetaD2A combined with other latency predictors. We conduct NAS on NAS-Bench-201 [8] benchmark for the CIFAR-100 dataset, using Google Pixel2 mobile phone and NVIDIA Titan RTX GPU as the target devices*. The results in Table 4 show that HELP largely outperforms BRP-NAS [9], the previous state-of-the art latency predictor, with 90× sample efficiency and 9.8×, 9.9× computational efficiency on Pixel2 and Titan RTX, respectively. Specifically, HELP + MetaD2A can efficiently retrieve an optimal latency-constrained architecture in 125s/111s for the given dataset on Pixel2/Titan RTX, respectively,

*For more results on various devices, please check the supplementary file.

Table 5: The results of the latency constrained-NAS experiment for ImageNet-1k with MobileNetV3 search space, with NVIDIA Titan RTX GPU and Xeon CPU as the target device. OFA + HELP reduces the total NAS cost to 26 GPU seconds, from 15 GPU hours required by the OFA + layer-wise predictor.

Device	Model	MACs (M)	Latency (ms)	Accuracy (%)	Latency Model Sample	Latency Model Building Time	Total NAS Cost (Wall Clock)	Speed Up
Unseen Device Titan RTX (Batch 64)	MobileNetV3-Large [12]	219M	22.1	75.2	-	-	-	-
	MnasNet-A1 [28]	312M	20.0	75.2	8k	4.5h	40,004.5h	1.0×
	FBNet-C [33]	375M	27.5	74.9	7.5k	4.2h	580.2h	69×
	ProxylessNAS-GPU [3]	465M	22.0	75.1	5k	2.8h	502.8h	80×
	OFA+Layer-wise Pred. [4]	397M	21.5	76.4	27k	15h	15h	2667×
	OFA + HELP (20ms)	230M	20.3	76.0	**10**	**26s**	**0.007h (26s)**	**5.7M×**
	OFA + HELP (23ms)	268M	23.1	76.8				
	OFA + HELP (28ms)	346M	28.6	77.9				
Unseen Platform Xeon CPU Gold 6226	MobileNetV3-Large [12]	219M	132	75.2	-	-	-	-
	MnasNet-A1 [28]	312M	212	75.2	8k	35.6h	40,035.6h	1.0×
	FBNet-B [33]	295M	212	74.1	7.5k	33.3h	609.3h	66×
	ProxylessNAS-CPU [3]	465M	200	75.1	5k	22.2h	522.2h	77×
	OFA+Layer-wise Pred. [4]	301M	167	74.6	27k	120h	120h	334×
	OFA + HELP (170ms)	336M	147	77.6	**20**	**300s**	**0.08h (300s)**	**0.5M×**
	OFA + HELP (190ms)	375M	171	78.1				

while BRP-NAS [9]'s predictor with large building time (1120s/940s) becomes a bottleneck for MetaD2A's rapid NAS process. Further, we validate the accuracy-efficiency trade-off of HELP against the baseline latency estimators, by combining them with the oracle accuracy predictor on NAS-Bench-201 space (in Figure 6). With HELP, oracle NAS obtains near Pareto-optimal models in most cases, while combining it with other latency estimators yield sub-optimal models.

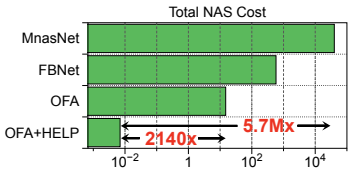

Figure 7: HELP reduces the total NAS cost by 2140× on Titan RTX. The total NAS cost is represented on a log-scale.

In Table 5, we validate HELP on a large scale dataset, ImageNet [7], against state-of-the-art NAS methods. In this experiment, we only consider GPUs with various batch sizes (1, 32, 64) as the meta-training devices, and Titan RTX GPU and Xeon CPU as unseen devices (for results on more devices, please refer to the supplementary file). We combine HELP with OFA by training the accuracy predictor on the pre-trained OFA network, and conducting predictor-guided evolutionary architecture search to fit the latency constraint. The result shows that, while searching for competitive architectures, HELP significantly reduces the total NAS cost of the baseline NAS methods which builds a new latency estimator for each target device. Figure 7 shows that OFA+HELP reduces the total NAS cost on a target device by 2140 times, when compared with the OFA + layer-wise predictor. This allows us to benefit from the rapid search speed of OFA, since the total NAS cost is only tens or hundreds of seconds, while the OFA + layer-wise predictor takes 15 hours, which is impractical. The details of the searched architectures are provided in supplementary file.

ProxylessNAS has roughly 42 unique blocks (7 different operations per 6 different input shapes) and use 5k architecture samples to build a layer-wise latency predictor. Thus, we proportionally estimate the number of samples to train its latency estimator for the FBNet and OFA supernet as 7.5k and 27k, which have 63 and 225 unique blocks, respectively. MnasNet does not have a latency estimator, but directly measures the latency of every architecture in the search process (8k architectures).

5 Conclusion

We proposed a novel meta-learned latency predictor, that can estimate the latency of an architecture on a novel device, using only a few measurements from it. While conventional latency prediction methods are inefficient since they cannot generalize across devices, and require a large number of latency measurements for each device, our latency predictor is meta-learned to rapidly adapt to an unseen device by utilizing the meta-knowledge accumulated over a device pool. Using a novel hardware embedding function that embeds each device based on its latencies on a set of reference architectures, we conducted hardware-conditioned meta-learning to obtain a device-specific initial parameters, and further took inner gradient steps to adapt to a new device. We validated our meta-latency predictor by measuring its latency estimation performance on unseen devices, on which it outperforms baselines, using only 10 to 20 samples per device. Further, we combined our latency predictor with two rapid NAS methods, to show that it performs latency-constrained NAS on unseen devices extremely fast and accurately, while baselines become computational bottlenecks with large time cost. We discuss the **societal impact** and **limitation** of this work in the supplementary file.

References

[1] Bowen Baker, Otkrist Gupta, Nikhil Naik, and Ramesh Raskar. Designing neural network architectures using reinforcement learning. In *In International Conference on Learning Representations (ICLR)*, 2017.

[2] Maxim Berman, Leonid Pishchulin, Ning Xu, Matthew B Blaschko, and Gérard Medioni. Aows: Adaptive and optimal network width search with latency constraints. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11217–11226, 2020.

[3] Han Cai, Ligeng Zhu, and Song Han. ProxylessNAS: Direct neural architecture search on target task and hardware. In *International Conference on Learning Representations*, 2019.

[4] Han Cai, Chuang Gan, Tianzhe Wang, Zhekai Zhang, and Song Han. Once for all: Train one network and specialize it for efficient deployment. In *International Conference on Learning Representations*, 2020. URL https://arxiv.org/pdf/1908.09791.pdf.

[5] Xiangning Chen, Ruochen Wang, Minhao Cheng, Xiaocheng Tang, and Cho-Jui Hsieh. Dr{nas}: Dirichlet neural architecture search. In *International Conference on Learning Representations*, 2021.

[6] Xiaoliang Dai, Peizhao Zhang, Bichen Wu, Hongxu Yin, Fei Sun, Yanghan Wang, Marat Dukhan, Yunqing Hu, Yiming Wu, Yangqing Jia, et al. Chamnet: Towards efficient network design through platform-aware model adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11398–11407, 2019.

[7] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009.

[8] Xuanyi Dong and Yi Yang. Nas-bench-201: Extending the scope of reproducible neural architecture search. In *International Conference on Learning Representations (ICLR)*, 2020.

[9] Łukasz Dudziak, Thomas Chau, Mohamed S Abdelfattah, Royson Lee, Hyeji Kim, and Nicholas D Lane. Brp-nas: Prediction-based nas using gcns. *Advances in neural information processing systems (NeurIPS)*, 2020.

[10] Thomas Elsken, Benedikt Staffler, Jan Hendrik Metzen, and Frank Hutter. Meta-learning of neural architectures for few-shot learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, 2020.

[11] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International Conference on Machine Learning (ICML)*, 2017.

[12] 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1314–1324, 2019.

[13] Hyunjik Kim, Andriy Mnih, Jonathan Schwarz, Marta Garnelo, Ali Eslami, Dan Rosenbaum, Oriol Vinyals, and Yee Whye Teh. Attentive neural processes. In *International Conference on Learning Representations*, 2019.

[14] Hae Beom Lee, Hayeon Lee, Donghyun Na, Saehoon Kim, Minseop Park, Eunho Yang, and Sung Ju Hwang. Learning to balance: Bayesian meta-learning for imbalanced and out-of-distribution tasks. In *International Conference on Learning Representations (ICLR)*, 2020.

[15] Hayeon Lee, Eunyoung Hyung, and Sung Ju Hwang. Rapid neural architecture search by learning to generate graphs from datasets. In *International Conference on Learning Representations*, 2021.

[16] Kwonjoon Lee, Subhransu Maji, Avinash Ravichandran, and Stefano Soatto. Meta-learning with differentiable convex optimization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.

[17] Chaojian Li, Zhongzhi Yu, Yonggan Fu, Yongan Zhang, Yang Zhao, Haoran You, Qixuan Yu, Yue Wang, Cong Hao, and Yingyan Lin. {HW}-{nas}-bench: Hardware-aware neural architecture search benchmark. In *International Conference on Learning Representations*, 2021.

[18] Zhenguo Li, Fengwei Zhou, Fei Chen, and Hang Li. Meta-sgd: Learning to learn quickly for few-shot learning. *arXiv preprint arXiv:1707.09835*, 2017.

[19] Dongze Lian, Yin Zheng, Yintao Xu, Yanxiong Lu, Leyu Lin, Peilin Zhao, Junzhou Huang, and Shenghua Gao. Towards fast adaptation of neural architectures with meta learning. In *International Conference on Learning Representations (ICLR)*, 2019.

[20] 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 *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018.

[21] Hanxiao Liu, Karen Simonyan, and Yiming Yang. Darts: Differentiable architecture search. In *In International Conference on Learning Representations (ICLR)*, 2019.

[22] Renqian Luo, Fei Tian, Tao Qin, Enhong Chen, and Tie-Yan Liu. Neural architecture optimization. In *Advances in neural information processing systems (NeurIPS)*, 2018.

[23] Alex Nichol, Joshua Achiam, and John Schulman. On first-order meta-learning algorithms. *arXiv preprint arXiv:1803.02999*, 2018.

[24] Hieu Pham, Melody Y Guan, Barret Zoph, Quoc V Le, and Jeff Dean. Efficient neural architecture search via parameter sharing. In *International Conference on Machine Learning (ICML)*, 2018.

[25] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4510–4520, 2018.

[26] Albert Shaw, Wei Wei, Weiyang Liu, Le Song, and Bo Dai. Meta architecture search. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.

[27] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *Advances in neural information processing systems (NIPS)*, 2017.

[28] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2820–2828, 2019.

[29] Sebastian Thrun and Lorien Pratt, editors. *Learning to Learn*. Kluwer Academic Publishers, Norwell, MA, USA, 1998. ISBN 0-7923-8047-9.

[30] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. In *Advances in neural information processing systems (NIPS)*, 2016.

[31] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12965–12974, 2020.

[32] Hanrui Wang, Zhanghao Wu, Zhijian Liu, Han Cai, Ligeng Zhu, Chuang Gan, and Song Han. Hat: Hardware-aware transformers for efficient natural language processing. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 2020.

[33] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10734–10742, 2019.

[34] 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. In *International Conference on Learning Representations (ICLR)*, 2020.

[35] Yuhui Xu, Lingxi Xie, Xiaopeng Zhang, Xin Chen, Bowen Shi, Qi Tian, and Hongkai Xiong. Latency-aware differentiable neural architecture search. *arXiv preprint arXiv:2001.06392*, 2020.

[36] Jiahui Yu, Pengchong Jin, Hanxiao Liu, Gabriel Bender, Pieter-Jan Kindermans, Mingxing Tan, Thomas Huang, Xiaodan Song, Ruoming Pang, and Quoc Le. Bignas: Scaling up neural architecture search with big single-stage models. In *European Conference on Computer Vision*, pages 702–717. Springer, 2020.

[37] Barret Zoph and Quoc V Le. Neural architecture search with reinforcement learning. In *In International Conference on Learning Representations (ICLR)*, 2017.

[38] Barret Zoph and Quoc V. Le. Neural architecture search with reinforcement learning. In *International Conference on Learning Representations (ICLR)*, 2017.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes]
   (c) Did you discuss any potential negative societal impacts of your work? [Yes]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] We report Section 4.1 averaged with 5 random runs. The results of Section 4.2 are conducted with single seed. For the experimental results with multiple seeds are in supplementary.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [Yes]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [Yes]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [Yes]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]

(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]