NAS-Bench-x11 and the Power of Learning Curves

Anonymous Author(s)
Affiliation
Address
email

Abstract

While early research in neural architecture search (NAS) required extreme computational resources, the recent releases of tabular and surrogate benchmarks have greatly increased the speed and reproducibility of NAS research. However, two of the most popular benchmarks do not provide the full training information for each architecture. As a result, on these benchmarks it is not possible to evaluate many types of multi-fidelity algorithms, such as learning curve extrapolation, that require evaluating architectures at arbitrary epochs. In this work, we present a method using singular value decomposition and noise modeling to create surrogate benchmarks, NAS-Bench-111, NAS-Bench-311, and NAS-Bench-NLP11, that output the full training information for each architecture, rather than just the final validation accuracy. We demonstrate the power of using the full training information by introducing a learning curve extrapolation framework to modify single-fidelity algorithms, showing that it leads to improvements over popular single-fidelity algorithms which claimed to be state-of-the-art upon release.

1 Introduction

In the past few years, algorithms for neural architecture search (NAS) have been used to automatically find architectures that achieve state-of-the-art performance on various datasets [77, 49, 35, 13]. In 2019, there were calls for reproducible and fair comparisons within NAS research [32, 52, 71, 34] due to both the lack of a consistent training pipeline between papers and experiments with not enough trials to reach statistically significant conclusions. These concerns spurred the release of tabular benchmarks, such as NAS-Bench-101 [72] and NAS-Bench-201 [11], created by fully training all architectures in search spaces of size $423\,624$ and $6\,466$, respectively. These benchmarks allow researchers to easily simulate NAS experiments, making it possible to run fair NAS comparisons and to run enough trials to reach statistical significance at very little computational cost or carbon emissions [17]. Recently, to extend the benefits of tabular NAS benchmarks to larger, more realistic NAS search spaces which cannot be evaluated exhaustively, it was proposed to construct *surrogate benchmarks* [56]. The first such surrogate benchmark is NAS-Bench-301 [56], which models the DARTS [35] search space of size 10^{18} architectures. It was created by fully training $60\,000$ architectures (both drawn randomly and chosen by top NAS methods) and then fitting a surrogate model that can estimate the performance of all of the remaining architectures. Since 2019, dozens of papers have used these NAS benchmarks to develop new algorithms [63, 51, 69, 60, 54].

An unintended side-effect of the release of these benchmarks is that a significant fraction of new NAS algorithms have been *single fidelity* algorithms: when the NAS algorithm chooses to evaluate an architecture, the architecture is fully trained and only the validation accuracy at the final epoch of training is outputted. This is because NAS-Bench-301 only contains the architectures' accuracy at epoch 100 and NAS-Bench-101 only contains the accuracies at epochs 4, 12, 36, 108 (only allowing single fidelity, or very limited multi-fidelity approaches). NAS-Bench-201 does allow queries on the entire learning curve, but it is smaller in size ($6\,466$) than NAS-Bench-101 ($423\,624$)

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

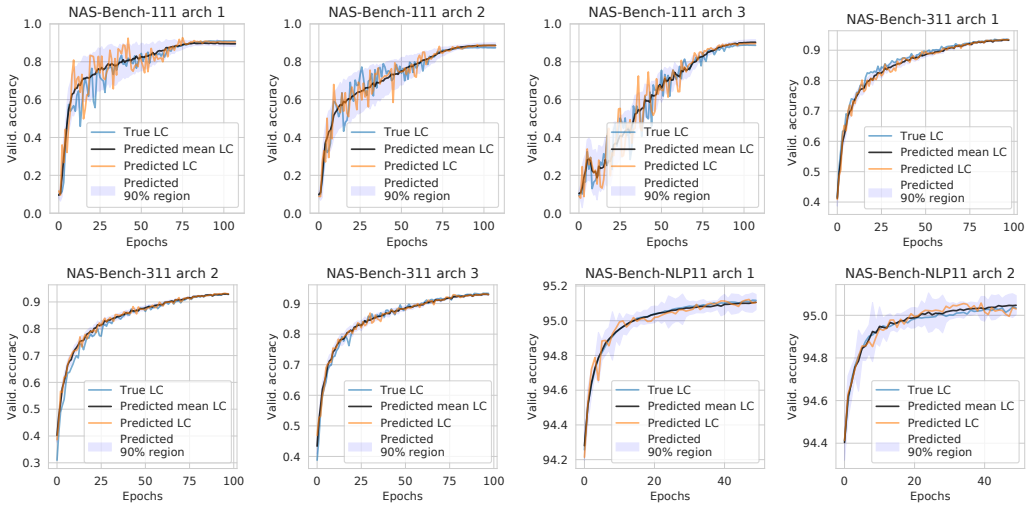

Figure 1: Each image shows a true learning curve vs. a learning curve predicted by one of our surrogate models, with and without predicted noise modeling. We also plot the 90% confidence interval of the predicted noise distribution. For a larger set of predicted learning curves, see Appendix D.

or NAS-Bench-301 (10^{18}). In a real world experiment, since training architectures to convergence is computationally intensive, researchers will often run *multi-fidelity* algorithms: the NAS algorithm can train architectures to any desired epoch. Here, the algorithm can make use of speedup techniques such as learning curve extrapolation (LCE) [59, 8, 1, 28]. Although multi-fidelity techniques are often used in the hyperparameter optimization community [20, 21, 33, 14, 59], they have been under-utilized by the NAS community in the last few years.

In this work, we fill in this gap by releasing NAS-Bench-111, NAS-Bench-311, and NAS-Bench-NLP11, surrogate benchmarks with full learning curve information for train/validation/test loss and accuracy for all architectures, significantly extending NAS-Bench-101, NAS-Bench-301, and NAS-Bench-NLP [30], respectively. With these benchmarks, researchers can easily incorporate multi-fidelity techniques, such as early stopping and LCE into their NAS algorithms. Our technique for creating these benchmarks can be summarized as follows. We use a training dataset of architectures (drawn randomly and chosen by top NAS methods) with good coverage over the search space, along with full learning curves, to fit a model that predicts the full learning curves of the remaining architectures. We employ three techniques to fit the model: *(1)* dimensionality reduction of the learning curves, *(2)* prediction of the top singular value coefficients, and *(3)* noise modeling. These techniques can be used in the future to create new NAS benchmarks as well. To ensure that our surrogate benchmarks are highly accurate, we report statistics such as Kendall Tau rank correlation and Kullback Leibler divergence between ground truth learning curves and predicted learning curves on separate test sets. See Figure 1 for examples of predicted learning curves on the test sets.

To demonstrate the power of using the full learning curve information, we present a framework for converting single-fidelity NAS algorithms into multi-fidelity algorithms using LCE. We apply our framework to popular single-fidelity NAS algorithms, such as regularized evolution [49], local search [64], and BANANAS [63], all of which claimed state-of-the-art upon release, showing that they can be further improved across four search spaces. Finally, we also benchmark multi-fidelity algorithms such as Hyperband [33] and BOHB [14] alongside single-fidelity algorithms. Overall, our work bridges the gap between different areas of AutoML and will allow researchers to easily develop effective multi-fidelity and LCE techniques in the future. We release all of our pretrained models and code. To promote reproducibility, we followed both the NeurIPS 2021 paper checklist and the specialized NAS best practices checklist [34].

Our contributions. We summarize our main contributions below.

- We develop a technique to create surrogate NAS benchmarks that include the full training information for each architecture, including train/validation/test loss and accuracy learning curves. This technique can be used to create future NAS benchmarks on any search space.

Table 1: Overview of existing NAS benchmarks. We introduce NAS-Bench-111, -311, and -NLP11. Note: NAS-Bench-NLP11 is based on a subset of NAS-Bench-NLP, as explained in Section 3.

Benchmark	Size	Queryable	Based on	Full train info
NAS-Bench-101	423k	✓	DARTS	✗
NAS-Bench-201	6k	✓		✓
NAS-Bench-NLP	10^{53}	✗		✗
NAS-Bench-301	10^{18}	✓		✗
NAS-Bench-ASR	8k	✓		✓
NAS-Bench-111	423k	✓	NAS-Bench-101	✓
NAS-Bench-311	10^{18}	✓	DARTS	✓
NAS-Bench-NLP11	10^{22}	✓	NAS-Bench-NLP	✓

- We apply our technique to create NAS-Bench-111, NAS-Bench-311, and NAS-Bench-NLP11, which allow researchers to easily develop multi-fidelity NAS algorithms that achieve higher performance than single-fidelity techniques.
- We present a framework for converting single-fidelity NAS algorithms into multi-fidelity NAS algorithms using learning curve extrapolation, and we show that our framework allows popular state-of-the-art NAS algorithms to achieve further improvements.

2 Related Work

NAS has been studied since at least the late 1980s [40, 26, 58] and has recently seen a resurgence [77, 41, 48, 22, 49, 18]. Weight sharing algorithms have become popular due to their computational efficiency [2, 35, 10, 75, 73, 47, 74]. Recent advances in performance prediction [61, 43, 55, 70, 36, 65, 51, 63] and other iterative techniques [14, 42] have reduced the runtime gap between iterative and weight sharing techniques. For the detailed surveys on NAS, we suggest referring to [13, 67].

Learning curve extrapolation. Several methods have been proposed to estimate the final validation accuracy of a neural network by extrapolating the learning curve of a partially trained neural network. Techniques include, fitting the partial curve to an ensemble of parametric functions [8], predicting the performance based on the partial trained neural network configurations [1], summing the training losses [50], using the basis functions as the output layer of a Bayesian neural network [28], using previous learning curves as basis function extrapolators [4], using the positive-definite covariance kernel to capture a variety of training curves [59], or using a Bayesian recurrent neural network [15]. While in this work we focus on multi-fidelity optimization utilizing learning curve-based extrapolation, another main category of methods lie in bandit-based algorithm selection [33, 14, 29, 19, 37], and the fidelities can be further adjusted according to the previous observations or a learning rate scheduler [20, 21, 27].

NAS benchmarks. NAS-Bench-101 [72], a tabular NAS benchmark, was created by defining a search space of size $423\,624$ unique architectures and then training all architectures from the search space on CIFAR-10 until 108 epochs. However, the train, validation, and test accuracies are only reported for epochs 4, 12, 36, and 108, and the train/valid/test losses are not reported. NAS-Bench-1shot1 [76] defines a subset of the NAS-Bench-101 search space that allows one-shot algorithms to be run. NAS-Bench-201 [11] contains $15\,625$ architectures, of which $6\,466$ are unique up to isomorphisms. It comes with full learning curve information on three datasets: CIFAR-10 [31], CIFAR-100 [31], and ImageNet16-120 [7]. Recently, NAS-Bench-201 was extended to NATS-Bench [9] which searches over architecture size as well as architecture topology.

Virtually every published NAS method for image classification in the last 3 years evaluates on the DARTS search space with CIFAR-10 [57]. The DARTS search space [35] consists of 10^{18} neural architectures, making it computationally prohibitive to create a tabular benchmark. To overcome this fundamental limitation and query architectures in this much larger search space, NAS-Bench-301 [56] evaluates various regression models trained on a sample of $60\,000$ architectures that is carefully created to cover the whole search space. The surrogate models allow users to query the validation

accuracy (at epoch 100) and training time for any of the 10^{18} architectures in the DARTS search space. However, since the surrogates do not predict the entire learning curve, it is not possible to run multi-fidelity algorithms.

NAS-Bench-NLP [30] is a search space for language modeling tasks. The search space consists of 10^{53} LSTM-like architectures, of which $14\,000$ are evaluated on Penn Tree Bank [39], containing the training, validation, and test losses/accuracies from epochs 1 to 50. Since only $14\,000$ of 10^{53} architectures can be queried, this dataset cannot be directly used for NAS experiments. NAS-Bench-ASR [38] is a recent tabular NAS benchmark for speech recognition. The search space consists of $8\,000$ architectures with full learning curve information. For an overview of NAS benchmarks, see Table 1. We give a more detailed discussion of all related work in Appendix B.

3 Generating Learning Curves

In this section, we describe our techniques to create a surrogate model that outputs realistic learning curves. We apply these techniques to create NAS-Bench-111, NAS-Bench-311, and NAS-Bench-NLP11. Our techniques apply to any type of learning curve, including train/test losses and accuracies. For simplicity, the following presentation assumes validation accuracy learning curves.

Given a search space $\mathcal{D}$, let $(\boldsymbol{x}_i, \boldsymbol{y}_i) \sim \mathcal{D}$ denote one datapoint, where $\boldsymbol{x}_i \in \mathbb{R}^d$ is the architecture encoding (e.g., one-hot adjacency matrix [72, 62]), and $\boldsymbol{y}_i \in [0,1]^{E_{\text{max}}}$ is a learning curve of validation accuracies drawn from a distribution $Y(\boldsymbol{x}_i)$ based on training the architecture for E_{max} epochs on a fixed training pipeline with a random initial seed. Each learning curve $\boldsymbol{y}_i$ can be decomposed into two parts: one part that is deterministic and depends only on the architecture encoding, and another part that is based on the inherent noise in the architecture training pipeline. Formally, $\boldsymbol{y}_i = \mathbb{E}[Y(\boldsymbol{x}_i)] + \boldsymbol{\epsilon}_i$, where $\mathbb{E}[Y(\boldsymbol{x}_i)] \in [0,1]^{E_{\text{max}}}$ is fixed and depends only on $\boldsymbol{x}_i$, and $\boldsymbol{\epsilon}_i \in [0,1]^{E_{\text{max}}}$ comes from a noise distribution Z_i with expectation 0 for all epochs. In practice, $\mathbb{E}[Y(\boldsymbol{x}_i)]$ can be estimated by averaging a large set of learning curves produced by training architecture $\boldsymbol{x}_i$ with different initial seeds. We represent such an estimate as $\bar{\boldsymbol{y}}_i$.

Our goal is to create a surrogate model that takes as input any architecture encoding $\boldsymbol{x}_i$ and outputs a distribution of learning curves that mimics the ground truth distribution. We assume that we are given two datasets, $\mathcal{D}_{\text{train}}$ and $\mathcal{D}_{\text{test}}$, of architecture and learning curve pairs. We use $\mathcal{D}_{\text{train}}$ (often size $> 10\,000$) to train the surrogate, and we use $\mathcal{D}_{\text{test}}$ for evaluation. We describe the process of creating $\mathcal{D}_{\text{train}}$ and $\mathcal{D}_{\text{test}}$ for specific search spaces in the next section. In order to predict a learning curve distribution for each architecture, we split up our approach into two separate processes: we train a model $f : \mathbb{R}^d \to [0,1]^{E_{\text{max}}}$ to predict the deterministic part of the learning curve, $\bar{\boldsymbol{y}}_i$, and we train a noise model $p_\phi(\boldsymbol{\epsilon} \mid \bar{\boldsymbol{y}}, \boldsymbol{x})$, parameterized by ϕ, to simulate the random draws from Z_i.

Surrogate model training. Training a model f to predict mean learning curves is a challenging task, since the training datapoints $\mathcal{D}_{\text{train}}$ consist only of a single (or few) noisy learning curve(s) $\boldsymbol{y}_i$ for each $\boldsymbol{x}_i$. Furthermore, E_{max} is typically length 100 or larger, meaning that f must predict a high-dimensional output. We propose a technique to help with both of these challenges: we use the training data to learn compression and decompression functions $c_k : [0,1]^{E_{\text{max}}} \to [0,1]^k$ and $d_k : [0,1]^k \to [0,1]^{E_{\text{max}}}$, respectively, for $k \ll E_{\text{max}}$. The surrogate is trained to predict *compressed* learning curves $c_k(\boldsymbol{y}_i)$ of size k from the corresponding architecture encoding $\boldsymbol{x}_i$, and then each prediction can be reconstructed to a full learning curve using d_k. A good compression model should not only cause the surrogate prediction to become significantly faster and simpler, but should also reduce the noise in the learning curves, since it would only save the most important information in the compressed representations. That is, $(d_k \circ c_k)(\boldsymbol{y}_i)$ should be a less noisy version of $\boldsymbol{y}_i$. Therefore, models trained on $c_k(\boldsymbol{y}_i)$ tend to have better generalization ability and do not overfit to the noise in individual learning curves.

We test two compression techniques: singular value decomposition (SVD) [16] and variational autoencoders (VAEs) [25], and we show later that SVD performs better. We give the details of SVD here and describe the VAE compression algorithm in Appendix C. Formally, we take the singular value decomposition of a matrix S of dimension $(|\mathcal{D}_{\text{train}}|, E_{\text{max}})$ created by stacking together the learning curves from all architectures in $\mathcal{D}_{\text{train}}$. Performing the truncated SVD on the learning curve matrix S allow us to create functions c_k and d_k that correspond to the optimal linear compression of S. In Figure 2 (left), we see that for architectures in the NAS-Bench-101 search space, there

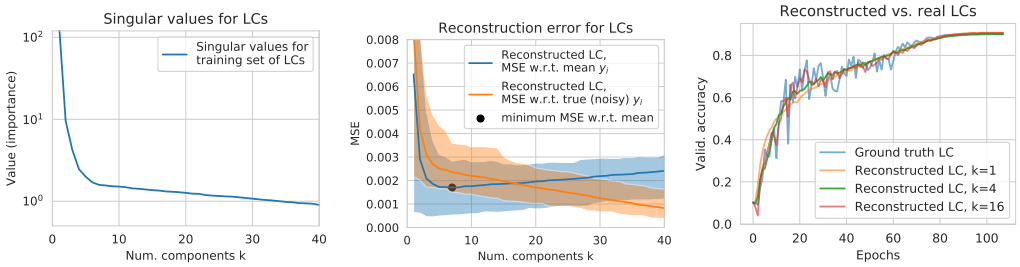

Figure 2: The singular values from the SVD decomposition of the learning curves (LC) (left). The MSE of a reconstructed LC, showing that $k = 6$ is closest to the true mean LC, while larger values of k overfit to the noise of the LC (middle). An LC reconstructed using different values of k (right).

is a steep dropoff of importance after the first six singular values, which intuitively means that most of the information for each learning curve is contained in its first six singular values. In Figure 2 (middle), we compute the mean squared error (MSE) of the reconstructed learning curves $(d_k \circ c_k)(\boldsymbol{y}_i)$ compared to a test set of ground truth learning curves averaged over several initial seeds (approximating $\mathbb{E}[Y(\boldsymbol{x}_i)]$). The lowest MSE is achieved at $k = 6$, which implies that $k = 6$ is the sweet spot where the compression function minimizes reconstruction error without overfitting to the noise of individual learning curves. We further validate this in Figure 2 (right) by plotting $(d_k \circ c_k)(\boldsymbol{y}_i)$ for different values of k. Now that we have compression and decompression functions, we train a surrogate model with $\boldsymbol{x}_i$ as features and $c_k(\boldsymbol{y}_i)$ as the label, for architectures in $\mathcal{D}_{\text{train}}$. We test LGBoost [23], XGBoost [5], and MLPs for the surrogate model.

Noise modeling. The final step for creating a realistic surrogate benchmark is to add a noise model, so that the outputs are *noisy* learning curves. We first create a new dataset of predicted $\boldsymbol{\epsilon}_i$ values, which we call residuals, by subtracting the reconstructed mean learning curves from the real learning curves in $\mathcal{D}_{\text{train}}$. That is, $\hat{\boldsymbol{\epsilon}}_i = \boldsymbol{y}_i - (d_k \circ c_k)(\boldsymbol{y}_i)$ is the residual for the ith learning curve. Since the training data only contains one (or few) learning curve(s) per architecture $\boldsymbol{x}_i$, it is not possible to accurately estimate the distribution Z_i for each architecture without making further assumptions. We assume that the noise comes from an isotropic Gaussian distribution, and we consider two other noise assumptions: *(1)* the noise distribution is the same for all architectures, and *(2)* for each architecture, the noise in a small window of epochs are iid. In Appendix C, we estimate the extent to which all of these assumptions hold true. Assumption *(1)* suggests two potential noise models: *(i)* a simple sample standard deviation statistic, $\boldsymbol{\sigma} \in \mathbb{R}_+^{E_{\max}}$, where

$$\sigma_j = \sqrt{\frac{1}{|\mathcal{D}_{\text{train}}| - 1} \sum_{i=1}^{|\mathcal{D}_{\text{train}}|} \hat{\epsilon}_{i,j}^2}.$$

To sample the noise using this model, we sample from $\mathcal{N}(\mathbf{0}, \text{diag}(\boldsymbol{\sigma}))$. *(ii)* The second model is a Gaussian kernel density estimation (GKDE) model [53] trained on the residuals to create a multivariate Gaussian kernel density estimate for the noise. Assumption *(2)* suggests one potential noise model: *(iii)* a model that is trained to estimate the distribution of the noise over a window of epochs of a specific architecture. For each architecture and each epoch, the model is trained to estimate the sample standard deviation of the residuals within a window centered around that epoch.

Evaluation. Overall, we test two compression methods, three surrogate models, and three noise models. For each approach, we evaluate both the predicted mean learning curves and the predicted noisy learning curves using held-out test sets $\mathcal{D}_{\text{test}}$. To evaluate the mean learning curves, we measure the coefficient of determination (R^2) [66] and Kendall Tau (KT) rank correlation [24], both at the final epoch and averaged over all epochs. To measure KT rank correlation for a specific epoch n, we find the number of concordant, P, and discordant, Q, pairs of predicted and true learning curve values for that epoch. The number of concordant pairs is given by the number of pairs, $((\hat{y}_{i,n}, y_{i,n}), (\hat{y}_{j,n}, y_{j,n}))$, where either both $\hat{y}_{i,n} > \hat{y}_{j,n}$ and $y_{i,n} > y_{j,n}$, or both $\hat{y}_{i,n} < \hat{y}_{j,n}$ and $y_{j,n} < y_{i,n}$. We can then calculate $KT = \frac{P-Q}{P+Q}$. While this metric can be used to compare surrogate predictions and has been used in prior work [56], the KT value is affected by the inherent noise in the learning curves (even an

oracle would achieve a KT value smaller than 1.0 which depends on the amount of noise). Finally, we evaluate the Kullback Leibler (KL) divergence between the ground truth distribution of noisy learning curves, and the predicted distribution of noisy learning curves on a test set. Since we can only estimate the ground truth distribution, we assume the ground truth is an isotropic Gaussian distribution. Then we measure the KL divergence between the true and predicted learning curves for architecture i by the following formula:

$$D_{KL}(\boldsymbol{y}_i||\hat{\boldsymbol{y}}_i) = \frac{1}{2E_{\text{max}}}\left[\log\frac{|\Sigma_{\hat{\boldsymbol{y}}_i}|}{|\Sigma_{\boldsymbol{y}_i}|} - E_{\text{max}} + (\boldsymbol{\mu}_{\boldsymbol{y}_i} - \boldsymbol{\mu}_{\hat{\boldsymbol{y}}_i})^T\Sigma_{\hat{\boldsymbol{y}}_i}^{-1}(\boldsymbol{\mu}_{\boldsymbol{y}_i} - \boldsymbol{\mu}_{\hat{\boldsymbol{y}}_i}) + tr\left\{\Sigma_{\hat{\boldsymbol{y}}_i}^{-1}\Sigma_{\boldsymbol{y}_i}\right\}\right]$$

where $\Sigma_{\{\boldsymbol{y}_i,\hat{\boldsymbol{y}}_i\}}$ is a diagonal matrix with the entries $\Sigma_{\{\boldsymbol{y}_i,\hat{\boldsymbol{y}}_i\}_{k,k}}$ representing the sample variance of the k-th epoch, and $\boldsymbol{\mu}_{\{\boldsymbol{y}_i,\hat{\boldsymbol{y}}_i\}}$ is the sample mean for either learning curves $\{\boldsymbol{y}_i, \hat{\boldsymbol{y}}_i\}$.

Surrogate Benchmark Creation. Now we describe the creation of NAS-Bench-111, NAS-Bench-311, and NAS-Bench-NLP11. As described above, we test two different compression methods (SVD, VAE), three different surrogate models (LGB, XGB, MLP), and three different noise models (stddev, GKDE, sliding window) for a total of eighteen distinct approaches. See Appendix C for a full ablation study, and Table 2 for a summary using the best techniques for each search space. See Figure 1 for a visualization of predicted learning curves from the test set of each search space using these models.

First, we describe the creation of NAS-Bench-111. Since the NAS-Bench-101 tabular benchmark [72] consists only of accuracies at epochs 4, 12, 36, and 108 (and without losses), we train a new set of architectures and save the full learning curves. Similar to prior work [12, 56], we sample a set of architectures with good overall coverage while also focusing on the high-performing regions exploited by NAS algorithms. Specifically, we sample 861 architectures generated uniformly at random, 149 architectures generated by 30 trials of BANANAS, local search, and regularized evolution, and all 91 architectures which contain fewer than five nodes, for a total of 1101 architectures. We kept our training pipeline as close as possible to the original pipeline. See Appendix C for the full training details. We find that SVD-LGB-GKDE achieves the best performance. Because the tabular benchmark already exists, we can substantially improve the accuracy of our surrogate by using the accuracies from the tabular benchmark (at epochs 4, 12, 36, 108) as additional features along with the architecture encoding. This substantially improves the performance of the surrogate, as shown in Table 2. Note that the large difference between the average KT and last epoch KT values show that the learning curves are very noisy (which is also evidenced in Figure 1).

Next, we create NAS-Bench-311 by using the training data from NAS-Bench-301, which consists of $40\,000$ random architectures along with $26\,000$ additional architectures generated by evolution [49], Bayesian optimization [3, 63, 44], and one-shot [35, 68, 10, 6] techniques in order to achieve good coverage over the search space. Again, SVD-LGB-GKDE achieves the best performance, which achieves an average and last epoch KT values of 0.728 and 0.788, respectively. This is comparable to the final KT of 0.817 reported by NAS-Bench-301 [56], despite optimizing for the full learning curve rather than only the final epoch. Furthermore, our KL divergences in Table 2 surpasses the top value of 16.4 reported by NAS-Bench-301 (for KL divergence, lower is better).

Finally, we create NAS-Bench-NLP11 by using the NAS-Bench-NLP dataset consisting of $14\,000$ architectures drawn uniformly at random. Due to the extreme size of the search space (10^{53}), we restrict architectures to a maximum of 12 nodes (reducing the size to 10^{22}), and we achieve an average and final epoch KT of 0.505 and 0.475, respectively. To create a stronger surrogate, we add the first three epochs of the learning curve as features in the surrogate. This improves the average and final epoch KT values to 0.878 and 0.844, respectively. To use this surrogate, any architecture to be predicted with the surrogate must be trained for three epochs. Note that the small difference between the average and last epoch KT values indicates that the learning curves have very little noise, which can also be seen in Figure 1. Since there are no architectures trained multiple times on the NAS-Bench-NLP dataset [30], we cannot compute the KL divergence.

4 The Power of Learning Curve Extrapolation

Now we describe a simple framework for converting single-fidelity NAS algorithms to multi-fidelity NAS algorithms using learning curve extrapolation techniques. We show that this framework is able

Table 2: Evaluation of the surrogate benchmarks on test sets. For NAS-Bench-111 and NAS-Bench-NLP11, we use architecture accuracies as additional features to improve performance.

Benchmark	**Avg.** R^2	**Final** R^2	**Avg. KT**	**Final KT**	**Avg. KL**	**Final KL**
NAS-Bench-111	0.529	0.630	0.531	0.645	2.016	1.061
NAS-Bench-111 (w. accs)	0.630	0.853	0.611	0.794	1.710	0.926
NAS-Bench-311	0.779	0.800	0.728	0.788	0.905	0.600
NAS-Bench-NLP11	0.326	0.314	0.505	0.475	-	-
NAS-Bench-NLP11 (w. accs)	0.878	0.895	0.878	0.844	-	-

to substantially improve the performance of popular algorithms such as regularized evolution [49], BANANAS [63], and local search [64, 45].

A single-fidelity algorithm is an algorithm which iteratively chooses an architecture based on its history, which is then fully trained to E_{max} epochs. To exploit parallel resources, many single-fidelity algorithms iteratively output several architectures at a time, instead of just one. In Algorithm 1, we present pseudocode for a generic single-fidelity algorithm. For example, for local search, `gen_candidates` would output the neighbors of the architecture with the best accuracy in `history`.

Our framework makes use of learning curve extrapolation (LCE) techniques [8, 1] to predict the final validation accuracies of all architecture choices after only training for a small number of epochs. See Algorithm 2. After each iteration of `gen_candidates`, only the architectures predicted by `LCE()` to be in the top percentage of validation accuracies of `history` are fully trained. For example, when the framework is applied to local search, in each iteration, all neighbors are trained to E_{few} epochs, and only the most promising architectures are trained up to E_{max} epochs. This simple modification can substantially improve the runtime efficiency of popular NAS algorithms by weeding out unpromising architectures before they are fully trained.

Any LCE technique can be used, and in our experiments in Section 5, we use weighted probabilistic modeling (WPM) [8] and learning curve support vector regressor (LcSVR) [1]. The first technique, WPM [8], is a function that takes a partial learning curve as input, and then extrapolates it by fitting the learning curve to a set of parametric functions, using MCMC to sample the most promising fit. The second technique, LcSVR [1], is a model-based learning curve extrapolation technique: after generating an initial set of training architectures, a support vector regressor is trained to predict the final validation accuracy from the architecture encoding and partial learning curve.

Algorithm 1 Single-Fidelity Algorithm

```
initialize history
while t < t_max :
  arches = gen_candidates(history)
  accs = train(arches, epoch=E_max)
  history.update(arches, accs)
Return arch with the highest acc
```

Algorithm 2 LCE Framework

```
initialize history
while t < t_max :
  arches = gen_candidates(history)
  accs = train(arches, epoch=E_few)
  pred_arch = LCE(arches, accs)
  arches = pred_arch[:top_n]
  accs = train(arches, epoch=E_max)
  history.update(arches, accs)
Return arch with the highest acc
```

5 Experiments

In this section, we benchmark single-fidelity and multi-fidelity NAS algorithms, including popular existing single-fidelity and multi-fidelity algorithms, as well as algorithms created using our framework defined in the previous section. In the experiments, we use our three surrogate benchmarks defined in Section 3, as well as NAS-Bench-201.

NAS algorithms. For single-fidelity algorithms, we implemented random search (RS) [32], local search (LS) [64, 45], regularized evolution (REA) [49], and BANANAS [63]. For multi-fidelity bandit-based algorithms, we implemented Hyperband (HB) [33] and Bayesian optimization Hyperband (BOHB) [14]. For all methods, we use the original implementation whenever possible. See

Appendix D for a description, implementation details, and hyperparameter details for each method. Finally, we use our framework from Section 4 to create six new multi-fidelity algorithms: BANANAS, LS, and REA are each augmented using WPM and LcSVR. This gives us a total of 12 algorithms.

Experimental setup. For each search space, we run each algorithm for a total wall-clock time that is equivalent to running 500 iterations of the single-fidelity algorithms for NAS-Bench-111 and NAS-Bench-311, and 100 iterations for NAS-Bench-201 and NAS-Bench-NLP11. For example, the average time to train a NAS-Bench-111 architecture to 108 epochs is roughly 10^3 seconds, so we set the maximum runtime on NAS-Bench-111 to roughly $5 \cdot 10^5$ seconds. We run 30 trials of each NAS algorithm and compute the mean and standard deviation.

Results. We evaluate BANANAS, LS, and REA compared to their augmented WPM and SVR versions in Figure 3 (NAS-Bench-311) and Appendix D (all other search spaces). Across four search spaces, we see that WPM and SVR improve all algorithms in almost all settings. The improvements are particularly strong for the larger NAS-Bench-111 and NAS-Bench-311 search spaces. We also see that for each single-fidelity algorithm, the LcSVR variant often outperforms the WPM variant. This suggests that model-based techniques for extrapolating learning curves are more reliable than extrapolating each learning curve individually, which has also been noted in prior work [65].

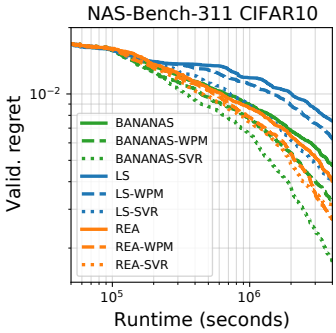

Figure 3: LCE Framework applied to single-fidelity algorithms on NAS-Bench-311.

In Figure 4, we compare single- and multi-fidelity algorithms on four search spaces, along with the three SVR-based algorithms from Figure 3. Across all search spaces, an SVR-based algorithm is the top-performing algorithm. Specifically, BANANAS-SVR performs the best on NAS-Bench-111, NAS-Bench-311, and NAS-Bench-NLP11, and LS-SVR performs the best on NAS-Bench-201. Note that HB and BOHB may not perform well on search spaces with low correlation between the relative rankings of architectures using low fidelities and high fidelities (such as NAS-Bench-201 [11]) since HB-based methods will predict the final accuracy of partially trained architectures directly from the last trained accuracy (i.e., extrapolating the learning curve as a constant after the last seen accuracy). On the other hand, SVR-based approaches use a model that can learn more complex relationships between accuracy at an early epoch vs. accuarcy at the final epoch, and are therefore more robust to this type of search space.

In Appendix D, we perform an ablation study on the epoch at which the SVR and WPM methods start extrapolating in our framework (i.e., we ablate E_{few} from Algorithm 2). We find that for most search spaces, running SVR and WPM based NAS methods by starting the LCE at roughly 20% of the total number of epochs performs the best. Any earlier, and there is not enough information to accurately extrapolate the learning curve. Any later, and the LCE sees diminishing returns because less time is saved by early stopping the training.

6 Societal Impact

Our hope is that our work will make it quicker and easier for researchers to run fair experiments and give reproducible conclusions. In particular, the surrogate benchmarks allow AutoML researchers to develop NAS algorithms directly on a CPU, as opposed to using a GPU, which may decrease the carbon emissions from GPU-based NAS research [46, 17]. In terms of our proposed NAS speedups, these techniques are a level of abstraction away from real applications, but they can indirectly impact the broader society. For example, this work may facilitate the creation of new high-performing NAS techniques, which can then be used to improve various deep learning methods, both beneficial (e.g. algorithms that reduce CO_2 emissions), or harmful (e.g. algorithms that produce deep fakes).

7 Conclusions, Limitations, and Guidelines

In this work, we released three benchmarks for neural architecture search based on three popular search spaces, which substantially improve the capability of existing benchmarks due to the avail-

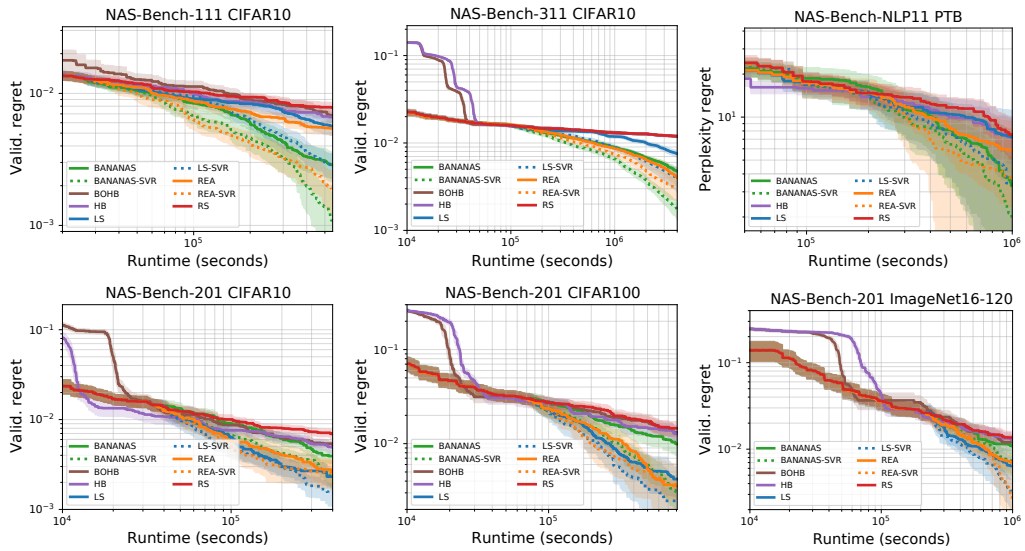

Figure 4: NAS results on six different combinations of search spaces and datasets. For every setting, an SVR augmented method performs best.

ability of the full learning curve for train/validation/test loss and accuracy for each architecture. Our techniques to generate these benchmarks, which includes singular value decomposition of the learning curve and noise modeling, can be used to model the full learning curve for future surrogate NAS benchmarks as well.

Furthermore, we demonstrated the power of the full learning curve information by introducing a framework that converts single-fidelity NAS algorithms into multi-fidelity NAS algorithms that make use of learning curve extrapolation techniques. This framework improves the performance of recent popular single-fidelity algorithms which claimed to be state-of-the-art upon release.

While we believe our surrogate benchmarks will help advance scientific research in NAS, a few guidelines and limitations are important to keep in mind. As with prior surrogate benchmarks [56], we give the following two caveats. *(1)* We discourage evaluating NAS methods that use the same internal techniques as those used in the surrogate model. For example, any NAS method that makes use of variational autoencoders or XGBoost should not be benchmarked using our VAE-XGB surrogate benchmark. *(2)* As the surrogate benchmarks are likely to evolve as new training data is added, or as better techniques for training a surrogate are devised, we recommend reporting the surrogate benchmark version number whenever running experiments.

We also note the following strengths and limitations for specific benchmarks. Our NAS-Bench-111 surrogate benchmark gives strong performance even with just $1\,101$ architectures used as training data, due to the existence of the four extra validation accuracies from NAS-Bench-101 that can be used as additional features. While these features help to achieve high predictive power over the entire search space, the only limitation is that this technique likely cannot be used to create future surrogate benchmarks: we hope that all future surrogate works will save the full learning curve information from the start so that the creation of an after-the-fact extended benchmark is not necessary.

Since the NAS-Bench-NLP surrogate benchmark achieves significantly stronger performance when the accuracy of the first three epochs are added as features, we recommend using this benchmark by training architectures for three epochs before querying the surrogate. Therefore, benchmarking NAS algorithms are slower than for NAS-Bench-111 and NAS-Bench-311, but NAS-Bench-NLP11 still offers a $15\times$ speedup compared to a NAS experiment without this benchmark. We still release the fast version of NAS-Bench-NLP11 that does not use the first three accuracies as features but with a warning that the observed NAS trends may differ from the true NAS trends. We expect that the performance of these benchmarks will improve over time, as data for more trained architectures become available.

References

[1] Bowen Baker, Otkrist Gupta, Ramesh Raskar, and Nikhil Naik. Accelerating neural architecture search using performance prediction. In *ICLR Workshop*, 2018.

[2] Gabriel Bender, Pieter-Jan Kindermans, Barret Zoph, Vijay Vasudevan, and Quoc Le. Understanding and simplifying one-shot architecture search. In *ICML*, 2018.

[3] James S Bergstra, Rémi Bardenet, Yoshua Bengio, and Balázs Kégl. Algorithms for hyper-parameter optimization. In *NeurIPS*, 2011.

[4] Akshay Chandrashekaran and Ian R Lane. Speeding up hyper-parameter optimization by extrapolation of learning curves using previous builds. In *ECML-PKDD*, 2017.

[5] Tianqi Chen and Carlos Guestrin. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, pages 785–794, 2016.

[6] Xiangning Chen, Ruochen Wang, Minhao Cheng, Xiaocheng Tang, and Cho-Jui Hsieh. Drnas: Dirichlet neural architecture search. In *ICLR*, 2021.

[7] Patryk Chrabaszcz, Ilya Loshchilov, and Frank Hutter. A downsampled variant of imagenet as an alternative to the cifar datasets. In *arXiv:1707.08819*, 2017.

[8] 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.

[9] Xuanyi Dong, Lu Liu, Katarzyna Musial, and Bogdan Gabrys. Nats-bench: Benchmarking nas algorithms for architecture topology and size. In *PAMI*, 2021.

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

[11] Xuanyi Dong and Yi Yang. Nas-bench-201: Extending the scope of reproducible neural architecture search. In *ICLR*, 2020.

[12] Katharina Eggensperger, Frank Hutter, Holger Hoos, and Kevin Leyton-Brown. Efficient benchmarking of hyperparameter optimizers via surrogates. In *AAAI*, 2015.

[13] Thomas Elsken, Jan Hendrik Metzen, and Frank Hutter. Neural architecture search: A survey. In *JMLR*, 2019.

[14] Stefan Falkner, Aaron Klein, and Frank Hutter. Bohb: Robust and efficient hyperparameter optimization at scale. In *ICML*, 2018.

[15] Matilde Gargiani, Aaron Klein, Stefan Falkner, and Frank Hutter. Probabilistic rollouts for learning curve extrapolation across hyperparameter settings. *arXiv preprint arXiv:1910.04522*, 2019.

[16] Gene Golub and William Kahan. Calculating the singular values and pseudo-inverse of a matrix. *Journal of the Society for Industrial and Applied Mathematics, Series B: Numerical Analysis*, 2(2):205–224, 1965.

[17] Karen Hao. Training a single ai model can emit as much carbon as five cars in their lifetimes. *MIT Technology Review*, 2019.

[18] Hanzhang Hu, John Langford, Rich Caruana, Saurajit Mukherjee, Eric Horvitz, and Debadeepta Dey. Efficient forward architecture search. In *NeurIPS*, 2019.

[19] Yimin Huang, Yujun Li, Hanrong Ye, Zhenguo Li, and Zhihua Zhang. An asymptotically optimal multi-armed bandit algorithm and hyperparameter optimization. *arXiv preprint arXiv:2007.05670*, 2020.

[20] Kirthevasan Kandasamy, Gautam Dasarathy, Junier B. Oliva, Jeff Schneider, and Barnabas Poczos. Multi-fidelity gaussian process bandit optimisation. In *NeurIPS*, 2016.

[21] Kirthevasan Kandasamy, Gautam Dasarathy, Jeff Schneider, and Barnabas Poczos. Multi-fidelity bayesian optimisation with continuous approximations. In *JMLR*, 2017.

[22] Kirthevasan Kandasamy, Willie Neiswanger, Jeff Schneider, Barnabas Poczos, and Eric P Xing. Neural architecture search with bayesian optimisation and optimal transport. In *NeurIPS*, 2018.

[23] Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. Lightgbm: A highly efficient gradient boosting decision tree. In *NeurIPS*, 2017.

[24] Maurice G Kendall. A new measure of rank correlation. *Biometrika*, 30(1/2):81–93, 1938.

[25] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. In *ICLR*, 2014.

[26] Hiroaki Kitano. Designing neural networks using genetic algorithms with graph generation system. *Complex systems*, 4(4):461–476, 1990.

[27] Aaron Klein, Stefan Falkner, Simon Bartels, Philipp Hennig, and Frank Hutter. Fast Bayesian Optimization of Machine Learning Hyperparameters on Large Datasets. In *AISTATS*, 2017.

[28] Aaron Klein, Stefan Falkner, Jost Tobias Springenberg, and Frank Hutter. Learning curve prediction with bayesian neural networks. In *ICLR*, 2017.

[29] Aaron Klein, Louis Tiao, Thibaut Lienart, Cedric Archambeau, and Matthias Seeger. Model-based asynchronous hyperparameter and neural architecture search. *arXiv preprint arXiv:2003.10865*, 2020.

[30] Nikita Klyuchnikov, Ilya Trofimov, Ekaterina Artemova, Mikhail Salnikov, Maxim Fedorov, and Evgeny Burnaev. Nas-bench-nlp: neural architecture search benchmark for natural language processing. *arXiv preprint arXiv:2006.07116*, 2020.

[31] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.

[32] Liam Li and Ameet Talwalkar. Random search and reproducibility for neural architecture search. In *UAI*, 2019.

[33] Lisha Li, Kevin Jamieson, Giulia DeSalvo, Afshin Rostamizadeh, and Ameet Talwalkar. Hyperband: A novel bandit-based approach to hyperparameter optimization. In *JMLR*, 2018.

[34] Marius Lindauer and Frank Hutter. Best practices for scientific research on neural architecture search. In *JMLR*, 2020.

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

[36] Renqian Luo, Xu Tan, Rui Wang, Tao Qin, Enhong Chen, and Tie-Yan Liu. Semi-supervised neural architecture search. In *NeurIPS*, 2020.

[37] Neeratyoy Mallik and Noor Awad. Dehb: Evolutionary hyperband for scalable, robust and efficient hyperparameter optimization. In *IJCAI*, 2021.

[38] Abhinav Mehrotra, Alberto Gil C. P. Ramos, Sourav Bhattacharya, Łukasz Dudziak, Ravichander Vipperla, Thomas Chau, Mohamed S Abdelfattah, Samin Ishtiaq, and Nicholas Donald Lane. Nas-bench-asr: Reproducible neural architecture search for speech recognition. In *ICLR*, 2021.

[39] Tomáš Mikolov, Martin Karafiát, Lukáš Burget, Jan Černockỳ, and Sanjeev Khudanpur. Recurrent neural network based language model. In *Annual conference of the international speech communication association*, 2010.

[40] Geoffrey F Miller, Peter M Todd, and Shailesh U Hegde. Designing neural networks using genetic algorithms. In *ICGA*, volume 89, pages 379–384, 1989.

[41] Renato Negrinho and Geoff Gordon. Deeparchitect: Automatically designing and training deep architectures. *arXiv preprint arXiv:1704.08792*, 2017.

[42] Vu Nguyen, Sebastian Schulze, and Michael Osborne. Bayesian optimization for iterative learning. In *NeurIPS*, 2020.

[43] Xuefei Ning, Wenshuo Li, Zixuan Zhou, Tianchen Zhao, Yin Zheng, Shuang Liang, Huazhong Yang, and Yu Wang. A surgery of the neural architecture evaluators. In *ECCV*, 2020.

[44] Changyong Oh, Jakub M Tomczak, Efstratios Gavves, and Max Welling. Combinatorial bayesian optimization using the graph cartesian product. *arXiv preprint arXiv:1902.00448*, 2019.

[45] T Den Ottelander, Arkadiy Dushatskiy, Marco Virgolin, and Peter AN Bosman. Local search is a remarkably strong baseline for neural architecture search. In *International Conference on Evolutionary Multi-Criterion Optimization*, 2021.

[46] David Patterson, Joseph Gonzalez, Quoc Le, Chen Liang, Lluis-Miquel Munguia, Daniel Rothchild, David So, Maud Texier, and Jeff Dean. Carbon emissions and large neural network training. *arXiv preprint arXiv:2104.10350*, 2021.

[47] Houwen Peng, Hao Du, Hongyuan Yu, Qi Li, Jing Liao, and Jianlong Fu. Cream of the crop: Distilling prioritized paths for one-shot neural architecture search. In *NeurIPS*, 2020.

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

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

[50] Binxin Ru, Clare Lyle, Lisa Schut, Mark van der Wilk, and Yarin Gal. Revisiting the train loss: an efficient performance estimator for neural architecture search. *arXiv preprint arXiv:2006.04492*, 2020.

[51] Binxin Ru, Xingchen Wan, Xiaowen Dong, and Michael Osborne. Neural architecture search using bayesian optimisation with weisfeiler-lehman kernel. In *ICLR*, 2021.

[52] Christian Sciuto, Kaicheng Yu, Martin Jaggi, Claudiu Musat, and Mathieu Salzmann. Evaluating the search phase of neural architecture search. In *ICLR*, 2020.

[53] David W Scott. *Multivariate density estimation: theory, practice, and visualization*. John Wiley & Sons, 2015.

[54] Han Shi, Renjie Pi, Hang Xu, Zhenguo Li, James Kwok, and Tong Zhang. Bridging the gap between sample-based and one-shot neural architecture search with bonas. *Advances in Neural Information Processing Systems*, 33, 2020.

[55] Han Shi, Renjie Pi, Hang Xu, Zhenguo Li, James Kwok, and Tong Zhang. Bridging the gap between sample-based and one-shot neural architecture search with bonas. In *NeurIPS*, 2020.

[56] Julien Siems, Lucas Zimmer, Arber Zela, Jovita Lukasik, Margret Keuper, and Frank Hutter. Nas-bench-301 and the case for surrogate benchmarks for neural architecture search. *arXiv preprint arXiv:2008.09777*, 2020.

[57] Julien Siems, Lucas Zimmer, Arber Zela, Jovita Lukasik, Margret Keuper, and Frank Hutter. Nas-bench-301 and the case for surrogate benchmarks for neural architecture search: Openreview response, 2021.

[58] Kenneth O Stanley and Risto Miikkulainen. Evolving neural networks through augmenting topologies. *Evolutionary computation*, 10(2):99–127, 2002.

[59] Kevin Swersky, Jasper Snoek, and Ryan Prescott Adams. Freeze-thaw bayesian optimization. *arXiv preprint arXiv:1406.3896*, 2014.

[60] Chen Wei, Chuang Niu, Yiping Tang, and Jimin Liang. Npenas: Neural predictor guided evolution for neural architecture search. *arXiv preprint arXiv:2003.12857*, 2020.

[61] Wei Wen, Hanxiao Liu, Hai Li, Yiran Chen, Gabriel Bender, and Pieter-Jan Kindermans. Neural predictor for neural architecture search. In *ECCV*, 2020.

[62] Colin White, Willie Neiswanger, Sam Nolen, and Yash Savani. A study on encodings for neural architecture search. In *NeurIPS*, 2020.

[63] Colin White, Willie Neiswanger, and Yash Savani. Bananas: Bayesian optimization with neural architectures for neural architecture search. In *AAAI*, 2021.

[64] Colin White, Sam Nolen, and Yash Savani. Local search is state of the art for nas benchmarks. In *UAI*, 2021.

[65] Colin White, Arber Zela, Binxin Ru, Yang Liu, and Frank Hutter. How powerful are performance predictors in neural architecture search? *arXiv preprint arXiv:2104.01177*, 2021.

[66] Sewall Wright. Correlation and causation. *Journal of Agricultural Research*, 20:557–580, 1921.

[67] Lingxi Xie, Xin Chen, Kaifeng Bi, Longhui Wei, Yuhui Xu, Zhengsu Chen, Lanfei Wang, An Xiao, Jianlong Chang, Xiaopeng Zhang, et al. Weight-sharing neural architecture search: A battle to shrink the optimization gap. *arXiv preprint arXiv:2008.01475*, 2020.

[68] 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 *ICLR*, 2019.

[69] Shen Yan, Kaiqiang Song, Fei Liu, and Mi Zhang. Cate: Computation-aware neural architecture encoding with transformers. In *ICML*, 2021.

[70] Shen Yan, Yu Zheng, Wei Ao, Xiao Zeng, and Mi Zhang. Does unsupervised architecture representation learning help neural architecture search? In *NeurIPS*, 2020.

[71] Antoine Yang, Pedro M Esperança, and Fabio M Carlucci. Nas evaluation is frustratingly hard. In *ICLR*, 2020.

[72] Chris Ying, Aaron Klein, Esteban Real, Eric Christiansen, Kevin Murphy, and Frank Hutter. Nas-bench-101: Towards reproducible neural architecture search. In *ICML*, 2019.

[73] Shan You, Tao Huang, Mingmin Yang, Fei Wang, Chen Qian, and Changshui Zhang. Greedynas: Towards fast one-shot nas with greedy supernet. In *CVPR*, 2020.

[74] Kaicheng Yu, Rene Ranftl, and Mathieu Salzmann. Landmark regularization: Ranking guided super-net training in neural architecture search. In *CVPR*, 2021.

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

[76] Arber Zela, Julien Siems, and Frank Hutter. Nas-bench-1shot1: Benchmarking and dissecting one-shot neural architecture search. In *ICLR*, 2020.

[77] Barret Zoph and Quoc V. Le. Neural architecture search with reinforcement learning. In *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] Our abstract and introduction accurately reflect our paper.
   (b) Did you describe the limitations of your work? [Yes] See Section 7
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] See Section 6.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes] We discuss the ethics guidelines in Section 6
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A] We did not include theoretical results.
   (b) Did you include complete proofs of all theoretical results? [N/A] We did not include theoretical results.
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] We included all code, data, and instructions in the supplementary material.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] All training details are specified in the supplementary material.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] Our experiments (Section 5 report error bars with respect to random seeds.
   (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] We report the compute time and resources used in Section 5
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] We cite the creators of all code, data, and models used.
   (b) Did you mention the license of the assets? [Yes] We mention the licenses.
   (c) Did you include any new assets either in the supplemental material or as a URL? [Yes] We include new surrogate benchmarks in the supplementary material.
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [Yes] We did not use or release any datasets with personal data.
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A] The data we are using does not contain personal information or offensive content.
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] We did not run experiments with human subjects.
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A] We did not run experiments with human subjects.
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A] We did not run experiments with human subjects.