

LNAS: Lottery Ticket Neural Architecture Search

Stanford CS231N Project: Theory and Reinforcement Learning

Sergio Charles
Department of Computer Science
Stanford University
sergiocl@stanford.edu

Gordon Chi
Department of Computer Science
Stanford University
gsyochi@stanford.edu

Gil Kornberg
Department of Computer Science
Stanford University
gilk@stanford.edu

Mentors: Yao Liu & Rachel Gardner. Note: Sergio is sharing this project with CS229.

Abstract

Despite the mainstream of usage of neural networks in increasingly difficult tasks in computer vision and natural language processing tasks, optimal network architectures are traditionally attained after countless manual experiments and rounds of redundant training iterations. We propose a different approach that incorporates Neural Architecture Search (NAS) method [10] and the lottery ticket hypothesis [4] – for a randomly initialized neural network, there exists a sub-network architecture that trains and performs at least as efficiently as the super network. Our results show that both variants of this search – the first by pruning model parameters, and the second by pruning architecture parameters – return a model sub-network architecture that achieves state-of-the-art performance without expensive memory or time usage.

1. Introduction

In the past decade, neural networks have been implemented for increasingly difficult tasks in computer vision and natural language processing domains. Yet despite their flexibility and utility, optimal network architectures remain difficult to design; state-of-the-art models are the fruits of countless manual experiments and rounds of trial-and-error training iterations. To bridge this gap, Zoph and Le proposed the use of Neural architecture search (NAS) [10], a

method that automatically learns network architectures and maximizes performance on a specific task in tandem. The most obvious drawback of this work: it uses an RNN meta-controller with proximal policy optimization which is too computationally expensive for most ML development tasks. In fact, the original NAS [10] used 450 GPUs for 3-4 days.

Consequently, there have been recent attempts to decrease the runtime of NAS. Efficient NAS [7] (ENAS) introduces the notion of searching for child models as sub-graphs over a directed acyclic (DAG) super-graph. Local computation nodes have parameters that are shared among child models in the search space, which yields a 1000× improvement over NAS. NASNet [11] defines the architecture of a ConvNet as stacked modules of two types of cells: a *normal cell* which preserves input dimension and a *reduction cell* which halves output dimension by using a stride of 2. They show that a well-chosen cell yields desirable transfer learning on downstream tasks. The Differentiable architecture search algorithm (DARTS) improves upon both of these methods by making the search space over architectures continuous via a softmax over the set of candidate operations [6].

However, DARTS still has prohibitively expensive memory usage due to redundant candidate architecture paths. Hence, we propose a neural architecture search method based on the lottery ticket hypothesis [4] which states that for a randomly initialized neural network, there exists a sub-network architecture that has commensurate accuracy and trains at least as efficiently as the super-network. More specifically, we propose two variants of this search: the first variation conducts a neural architecture search before pruning model parameters during train time based on our lottery

ticket hypothesis; the second variation conducts a neural architecture search that prunes network architecture parameters before performing a standard training loop. We theorize that regardless of the variant chosen, by performing iterative or one-shot pruning when training the architecture and child models, the sparse sub-network neural architecture *winning ticket* can be found in at most as many iterations, and with at least the same accuracy, of the model architecture found by DARTS.

2. Related Work

Neural architecture search (NAS) was proposed in [10] as a method to automatically learn network architectures to maximize performance on a specific task. The original implementation of NAS serialized CNN and RNN network topology representations as a sequence of operations at each layer. That is, they surmised that the "structure and connectivity of a neural network", namely the architectural hyperparameters, can be encoded by a sequence of tokens. Thus, they train an RNN meta controller to produce the sequence $a_{1:T}$ of hyperparameters used to design the architecture, which can be construed as a sequence of actions in the sense of reinforcement learning [10]. The NAS search space is the set of operations that can be performed between layers of a neural network, which induces a strong prior on network design. The controller samples a candidate model called the child model, trains it to convergence, and evaluates the reward R , measured by the accuracy, of the child model on a held-out validation set. The controller uses the objective of maximizing expected reward as an update signal.

More precisely, the NAS search algorithm iterates over several possible child models in its search space and optimizes expected reward, e.g. model size, accuracy, or latency. The meta-controller RNN optimizes the following objective [10]:

$$J(\theta_C) = \mathbb{E}_{P(a_{1:T}; \theta_C)}[R] \quad (1)$$

In the original NAS paper, they optimized the RNN controller parameters θ_C with REINFORCE [9] on the policy gradient [10]:

$$\nabla_{\theta_C} J(\theta_C) = \sum_{t=1}^T \mathbb{E}[\nabla_{\theta_C} \log P(a_t | a_{(t-1):1}; \theta_C) R]. \quad (2)$$

because R is non-differentiable.

Furthermore, we can represent a child model as a DAG with nodes as "local computations" and edges as the "flow of information" [7]. Efficient NAS (ENAS) [7] introduces the notion of searching for architectures as sub-graphs of the directed acyclic super-graph. That is, this DAG represents a superposition of all possible child models. These local

computation nodes have parameters and, thus, ENAS shares these parameters among child models in the search space, which yields a 1000x improvement over the original NAS method.

Differentiable Architecture Search (DARTS) [6] eschews the need for a proximal policy-based method. In DARTS, they learn a computation cell which can be stacked or recursively concatenated to create a convolutional neural network or recursive neural network, respectively. A computation cell can be represented as a DAG consisting of a topologically ordered set of n nodes. Let $x^{(i)}$ represent a latent representation and each edge (i, j) represent an operation $o^{(i,j)}$ on $x^{(i)}$. Using the topological ordering, we can compute the node based on the operations of preceding nodes [6]:

$$x^{(j)} = \sum_{i < j} o^{(i,j)}(x^{(i)}). \quad (3)$$

If \mathcal{O} denotes a set of our *candidate operations*, DARTS makes the search space continuous by considering a softmax over possible operations:

$$\hat{o}^{(i,j)}(x) = \sum_{o \in \mathcal{O}} \frac{e^{\alpha_o^{(i,j)}}}{\sum_{o' \in \mathcal{O}} e^{\alpha_{o'}^{(i,j)}}} o(x) \quad (4)$$

where $\alpha^{(i,j)}$ are learned parameters between nodes i and j ; finding these variables is equivalent to neural architecture search, which now allows us to employ a standard bi-level optimization. When we have found the neural architecture, we collapse its superposition state by discretizing the mixed operations $\hat{o}^{(i,j)}$ to $o^{(i,j)} = \arg \max_{o \in \mathcal{O}} \alpha_o^{(i,j)}$. Let $\mathcal{L}_{\text{train}}$ denote the training loss and let \mathcal{L}_{val} denote the validation loss. We use a bi-level objective proposed in DARTS:

$$\begin{aligned} & \min_{\alpha} \mathcal{L}_{\text{val}}(\theta^*(\alpha), \alpha) \\ & \text{subject to } \theta^*(\alpha) = \arg \min_{\theta} \mathcal{L}_{\text{train}}(\theta, \alpha). \end{aligned} \quad (5)$$

which determines the optimal α^* to minimize validation loss $\mathcal{L}_{\text{val}}(\theta^*, \alpha^*)$, where the optimal model parameters θ^* is $\theta^* = \arg \min_{\theta} \mathcal{L}_{\text{train}}(\theta, \alpha^*)$. The inner optimization is very computationally expensive and, thus, DARTS uses an unrolled approximation of [6]:

$$\nabla_{\theta} \mathcal{L}_{\text{val}}(\theta^*(\alpha), \alpha) \approx \nabla_{\alpha} \mathcal{L}_{\text{val}}(\theta - \xi \nabla_{\theta}(\theta, \alpha), \alpha) \quad (6)$$

which approximates $\theta^*(\alpha)$ by only using a single step of SGD to avoid performing the inner optimization until convergence.

However, the memory requirements of DARTS are still cumbersome because GPU consumption is proportional to $|\mathcal{O}|$. Methods like MnasNet [8] employ optimization with a multi-objective of accuracy on a target task and inference

time latency. Likewise, ProxylessNAS, proposed by Cai et. al [1], redresses the problem of high memory consumption by using path-level pruning which binarizes the associated softmax probabilities of the real-valued architecture parameters α . They also present "hardware-aware neural network specialization", which optimizes the architecture for the specific hardware device [1].

We make use of the results of the lottery ticket hypothesis [4], which states that for a randomly initialized neural network, there exists a sub-network architecture that has at least as good test accuracy and trains at least as efficiently as the super-network. Such a winning ticket sub-network is identified and trained as follows [4]:

1. Initialize a neural network $f_{\theta_0}(x)$ for parameters θ_0 distributed according to the training data $\theta \sim \mathcal{D}_{\text{train}}$.
2. Train the network for k iterations until we have parameters θ_j .
3. Prune $p\%$ of the parameters in the network, according to magnitude, forming a mask \mathcal{M} .
4. Reset the remaining parameters to their respective θ_0 values, which yields the winning ticket $f_{\mathcal{M} \odot \theta_0}(x)$.

3. Methods

In light of the lottery ticket hypothesis [4], we posit that by performing iterative or one-shot pruning during neural architecture search with DARTS, the sparse sub-network neural architecture *winning ticket* can be found in *commensurate time* (at most as many iterations as the architecture found by search) and with *commensurate accuracy* (at least the same accuracy of the architecture found by search).

This is distinct from NAS methods that optimize on a multi-objective of model size and accuracy on a target task, as seen in [3], because these methods do not leverage the notion of a winning ticket from the lottery ticket hypothesis. That is, they are not searching for sparse child neural architectures, but instead small neural architectures; this would generally improve inference time latency but does not address the computational expensiveness of modern NAS.

We search for an architecture of a ConvNet as stacked computation cells, like in NASNet [11]. We use two types of cells: a *normal cell* which preserves input dimension and a *reduction cell* which halves output dimension by using a stride of 2. NASNet showed a well-chosen cell yields desirable transfer learning on downstream tasks. We use DARTS [6] as the NAS method, which eschews the meta-controller configuration. Operations are limited to convolutions, average pooling, max pooling, and the zero operation.

Following the notation of ProxylessNAS [1], let $\mathcal{N}(e_1, \dots, e_n)$ denote the over-parameterized convolutional neural network, where e_i is an edge in its DAG representation of a topologically ordered set of its n nodes. Let

$x^{(i)}$ represent a latent representation and each edge $e_{i,j}$ represents the operation $o^{(i,j)}$ on $x^{(i)}$. For our N candidate operations \mathcal{O} , we make the search space continuous by defining a mixed operation as a softmax over possible operations [6]: $\hat{o}^{(i,j)}(x) = \sum_{o \in \mathcal{O}} \frac{e^{\alpha_{o}^{(i,j)}}}{\sum_{o' \in \mathcal{O}} e^{\alpha_{o'}^{(i,j)}}} o(x)$ where

$\alpha^{(i,j)}$ are learned architecture parameters between node i and node j . Since each edge is a mixed operation, i.e. a linear combination of the $N = |\mathcal{O}|$ candidate operations $\hat{o}(x) = \sum_{i=1}^N p_i o_i(x)$, we can intuitively represent a mixed operation by N paths where each path is weighted with probability p_i . When performing the neural ticket architecture search to find the winning neural ticket α^* , we apply an architecture-dependent mask \mathcal{M}_α :

$$\hat{o}(x) = \mathbb{E}_{o \sim \mathcal{O} \odot \mathcal{M}_\alpha} [o(x)]. \quad (7)$$

to prune the architecture parameters α . When we retrain the model f_θ with optimal architecture α^* , we use a model-dependent mask \mathcal{M}_θ to prune the parameters θ of the model f_θ . Since we effectively hard-prune the choice of candidate operations when selecting the architecture with highest candidate probability, this will induce significant regressions in performance. Therefore, it is important to retrain the model f_θ with the sparse architecture. Let $\mathcal{L}_{\text{train}}$ and \mathcal{L}_{val} and denote the training and validation loss, respectively. We use a bi-level objective proposed in DARTS:

$$\begin{aligned} \min_{\alpha} \mathcal{L}_{\text{val}}(\theta^*(\alpha), \alpha) \\ \text{subject to } \theta^*(\alpha) = \arg \min_{\theta} \mathcal{L}_{\text{train}}(\theta, \alpha). \end{aligned} \quad (8)$$

Iterative pruning is performed to identify the *neural winning ticket*. In the generalized method, the cell architecture is trained in tandem with the pruning program of the lottery ticket hypothesis until early stopping, and accuracy is measured on a held-out validation set. Then we retrain the model with the architecture found during the search using a similar pruning program. We developed two approaches to performing the lottery ticket pruning, which are as follows:

Variation 1: We perform DARTS neural architecture search to find the optimal architecture α^* . Then, when retraining the model f_θ with architecture α^* , with a sparsity of $p \in (0, 1)$, we prune $p^{1/n}$ of the lowest magnitude param-

ters for n rounds. The algorithm is enumerated as follows:

Algorithm 1: Variant 1 of LNAS

Randomly initialize $\theta_0 \sim \mathcal{D}_{\text{train}}, \alpha_0 \sim \mathcal{D}_{\text{val}}$;
 Search for optimal architecture α^*
while not converged do
 1. Perform the α gradient update
 $\alpha \leftarrow \alpha - \eta \nabla_{\alpha} \mathcal{L}_{\text{val}}(\theta - \xi \nabla_{\theta} \mathcal{L}_{\text{train}}(\theta, \alpha), \alpha)$ for k iterations;
 2. Perform the model θ gradient update
 $\theta \leftarrow \theta - \xi \nabla_{\theta} \mathcal{L}_{\text{train}}(\theta, \alpha)$ for k iterations;
end
 Re-train model parameters θ of optimal architecture α^* to find (α^*, θ^*)
for n pruning iterations do
 1. Perform the θ gradient update
 $\theta \leftarrow \theta - \xi \nabla_{\theta} \mathcal{L}_{\text{train}}(\theta, \alpha)$ for k iterations;
 2. Prune $p^{1/n}$ of smallest magnitude model parameters by applying mask: $\theta \odot \mathcal{M}_{\theta}$;
end
return winning neural ticket (α^*, θ^*)

For fixed θ , we optimize the architecture parameters α with SGD on the validation set, using the unrolled approximation of Equation 6. Then during retraining, we optimize the model parameters θ with SGD on the training set until convergence, while pruning $p^{1/n}$ of model parameters every k iterations of SGD for a total of n rounds of training.

Variant 2: We perform path-level pruning, in the sense of ProxylessNAS [1], whereby with a sparsity of $p \in (0, 1)$, we prune $p^{1/n}$ of the lowest-magnitude architecture parameters α during the neural architecture search. This corresponds to pruning candidate operations or, equivalently, paths for each edge of the computational graph. As before, during this step, we use the unrolled approximation in Equation 8. Then we re-train the model f_{θ} , with the optimal architecture α^* found during the previous step, via SGD until

convergence. The algorithm is as follows:

Algorithm 2: Variant 2 of LNAS

Randomly initialize $\theta_0 \sim \mathcal{D}_{\text{train}}, \alpha_0 \sim \mathcal{D}_{\text{val}}$;
 Search for optimal architecture α^*
for n pruning iterations do
 1. Perform the α gradient update
 $\alpha \leftarrow \alpha - \eta \nabla_{\alpha} \mathcal{L}_{\text{val}}(\theta - \xi \nabla_{\theta} \mathcal{L}_{\text{train}}(\theta, \alpha), \alpha)$ for k iterations;
 2. Perform the model θ gradient update
 $\theta \leftarrow \theta - \xi \nabla_{\theta} \mathcal{L}_{\text{train}}(\theta, \alpha)$ for k iterations;
 3. Prune $p^{1/n}$ of smallest magnitude architecture parameters by applying mask: $\alpha \odot \mathcal{M}_{\alpha}$.
end
 Re-train model parameters θ of optimal architecture α^* to find optimal (α^*, θ^*)
while not converged do
 1. Perform the θ gradient update
 $\theta \leftarrow \theta - \xi \nabla_{\theta} \mathcal{L}_{\text{train}}(\theta, \alpha)$;
end
return winning neural ticket (α, θ)

Although we did not complete a full grid search of pruning both model parameters and architecture parameters at various sparsity levels, to determine the relative utility of each variant in a wider scope, it is useful to consider the more generalized version of our algorithm:

Algorithm 3: Generalized LNAS

Randomly initialize $\theta_0 \sim \mathcal{D}_{\text{train}}, \alpha_0 \sim \mathcal{D}_{\text{val}}$;
 Search for optimal architecture α^*
for n pruning iterations do
 1. Perform the α gradient update
 $\alpha \leftarrow \alpha - \eta \nabla_{\alpha} \mathcal{L}_{\text{val}}(\theta - \xi \nabla_{\theta} \mathcal{L}_{\text{train}}(\theta, \alpha), \alpha)$ for k iterations;
 2. Perform the model θ gradient update
 $\theta \leftarrow \theta - \xi \nabla_{\theta} \mathcal{L}_{\text{train}}(\theta, \alpha)$ for k iterations;
 3. Prune $p^{1/n}$ of smallest magnitude architecture parameters by applying mask: $\alpha \odot \mathcal{M}_{\alpha}$.
end
 Re-train model parameters θ of optimal architecture α^* to find optimal (θ^*, α^*)
for n pruning iterations do
 1. Perform the θ gradient update
 $\theta \leftarrow \theta - \xi \nabla_{\theta} \mathcal{L}_{\text{train}}(\theta, \alpha)$ for k iterations;
 2. Prune $p^{1/n}$ of smallest magnitude model parameters by applying mask: $\theta \odot \mathcal{M}_{\theta}$.
end
return winning neural ticket (α, θ)

That is, we perform k SGD update steps for the DARTS neural architecture search, using the unrolled approximation in Equation 6. Then, with a sparsity of $p \in (0, 1)$, we prune $p^{1/n}$ of the smallest magnitude architecture pa-

rameters, which is a form of path-level pruning of the N possible operation candidate paths [1]. We do this for n rounds until we have found the optimal neural architecture α^* . Next, we optimize the model parameters θ with SGD on the training set for k steps. Then we prune $p^{1/n}$ of the smallest magnitude model parameters. We repeat for n rounds to converge to an optimal θ^* and α^* , which is the theorized winning neural ticket.

4. Data

For all experiments, we train and validate all experiments on the CIFAR-10 dataset [5], which consists of 60,000 32×32 colored images of 10 different classes. We utilize a 50,000 to 10,000 *train-test* split of the images; all final models were validated on Top-1 accuracy performance on the test dataset. No pre-processing or filtering was applied prior to the training of our networks.

5. Experiments

5.1. Lottery Ticket Neural Architecture Search

For each variant, we train models on different sparsity levels and evaluate their relative performance on the CIFAR-10 dataset to a baseline model constructed via differential architectural search (DARTS). Our goal is determine how the number of parameters remaining affects model accuracy and training time of the optimal *neural winning ticket*.

Baseline DARTS Model: We reimplemented the DARTS search on the CIFAR-10 dataset using the `nni` library as proposed in the original paper [6]. Our search space \mathcal{O} of candidate operations consisted of 3×3 and 5×5 separable convolutions, 3×3 and 5×5 dilated separable convolutions, 3×3 max and average pooling, identity mapping, and a zero mapping. All operations were stride one and padded to preserve spatial resolution. The convolutional cell consisted of 7 nodes, and was placed after a ReLU activation block and before batch normalization.

Our neural network was optimized with respect to cross-entropy loss, a metric known for its strong performance in classification tasks. We perform 10 epochs of neural architecture search on the \mathcal{O} search space described above, followed by 20 epochs of a standard training loop. Both the search and train loops utilized a stochastic gradient descent (SGD) + Momentum optimizer with 0.9 momentum and 3.0×10^{-4} weight decay; the initial learning rate of 0.025 was adjusted according to a cosine annealing scheduler. Batch size was set at 64.

Initial Architectural Search: For **Variation 1**, observe that all experiments focus only on pruning model parameters of

the baseline models given different sparsity levels. Hence, our initial model architecture is identical to the baseline model, where we performed 10 epochs of neural architecture search on the \mathcal{O} search space consisting of 3×3 and 5×5 separable convolutions, 3×3 and 5×5 dilated separable convolutions, 3×3 max and average pooling, identity mapping, and a zero mapping. Learning rates and optimizers were also kept constant.

For **Variation 2**, we performed a redesigned DARTS search on the CIFAR-10 dataset that iteratively prunes architecture parameters after each epoch of our neural architecture search, in conjunction with lottery ticket hypothesis. Otherwise, all other hyperparameters of the search were kept constant with the baseline model; our search space \mathcal{O} of candidate operations again consisted only of 3×3 and 5×5 separable convolutions, 3×3 and 5×5 dilated separable convolutions, 3×3 max and average pooling, identity mapping, and a zero mapping. Likewise, all operations are stride one and padded to preserve spatial resolution; the convolutional cell consisted of 7 nodes, and was placed after a ReLU activation block and before batch normalization. Learning rates and optimizers were also kept constant.

Training Loop: For **Variation 1** and a given sparsity level p , we run a modified training loop for 20 epochs, where on the 10th epoch we perform one-shot pruning that zeroes out p parameters of the model, before re-initializing the model parameters to θ_0 . For **Variation 2** and a given sparsity level p , a standard training loop for 20 epochs is implemented; there is no need to prune out parameters during train-time since the initial neural-architecture-search leaves us with a sufficiently pruned neural architecture.

Models trained from both variants were optimized with respect to cross-entropy loss, and all other hyperparameters were held constant. Both The training loop utilized a stochastic gradient descent (SGD) + Momentum optimizer with 0.9 momentum and 3.0×10^{-4} weight decay; the initial learning rate of 0.025 was adjusted according to a cosine annealing scheduler. Batch size was set at 64.

Sparsity Levels: For each variant, we experimented with 8 different sparsity levels: 0.125, 0.25, 0.375, 0.5, 0.675, 0.75, 0.875, 0.95. This leaves us with 16 trained models alongside our baseline model for analysis.

Evaluation Metrics: We evaluate our models on three metrics:

- First, we evaluate the prediction accuracy of the models across sparsity levels p and both variants, in order to determine which variant and sparsity levels leave us with a strong-performing model.

- Second, we evaluate test inference time by reporting the floating point operations per second (FLOPS) required; although the relative speed of models may be apparent given the chosen sparsity levels, this analysis presents with a better estimate of runtime.
- Finally, we evaluate train inference time by reporting the total number of hours required for the 10 epochs of neural architecture search and 20 epochs of the concurrent training loop. While we had attempted to keep all hyperparameters regarding batch size, learning rate and number of epochs constant across both variants, it is important to acknowledge that the lottery ticket hypothesis performs fundamentally different operations in the two variants, which may result in drastically different training times.

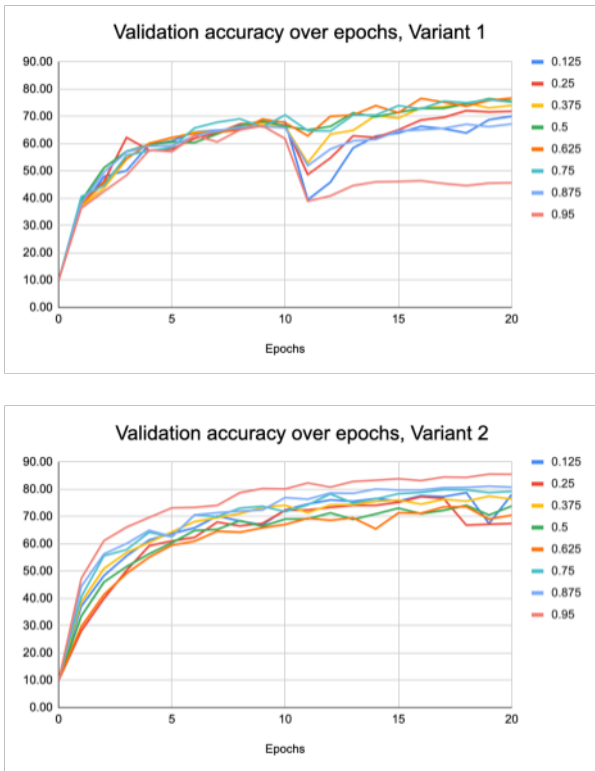


Figure 1. A visualization of model validation accuracy over epochs for both variants. Note that Variant 1 is less robust to sparsity levels; the sparsity levels of 0.95 is unable to reach past 50% accuracy in the aftermath of the one-shot pruning. Meanwhile, Variant 2 performs particularly well with the sparsity level of 0.95, reaching a validation accuracy of near 86.00% after the twenty epochs.

5.2. Prediction Accuracy and FLOPS

Overall, although not all sparsity levels and variant choices were able to obtain the winning lottery ticket af-

ter 20 epochs, many of our models returned comparable accuracy on the validation dataset. For Variant 1, the LNAS performed best on medium to high sparsity levels, with the models pruned on 0.5, 0.625, and 0.75 sparsity levels returning a validation accuracy of over 76% compared to just an average of around 70% for models trained on higher or lower sparsity levels. An opposite trend is observed in Variant 2, where the LNAS performed stronger overall and best on high and low sparsity levels. Surprisingly, we were able to obtain the three *winning tickets* under Variant 2, with the sparsity level of 0.95 obtaining a state-of-the-art validation accuracy of 85.6% in 20 epochs. This outperforms even the super baseline model we constructed with DARTS. More specific results across all the sparsity levels can be seen in Table 1 below.

We also present prediction accuracy on the validation set over time in Figure 1. Due to the pruning and re-initializing of weights to θ_0 in Variant 1, the drop in accuracy following Epoch 10 may result in models with a given sparsity level p being unable to attain high validation accuracy in the successive train steps. More importantly, our experiment suggests that if a model’s validation accuracy is immediately robust to the re-initializing of weights, then it is likely that the given sparsity level will be closer to a winning lottery ticket.

In contrast, the validation accuracy curves for Variant 2 do not have large jumps in performance, resembling a standard training loop.

5.3. Variant Train Time Comparison

Though all experiments were run with the same number of epochs in both the search and train loops, Variant 1 has the lottery ticket hypothesis implemented in only the training loop while Variant 2 has the same hypothesis implemented in the search loop. This results in large differences in train time: as outlined in the Methods section, the neural architecture search requires a longer runtime when implemented alongside the lottery ticket hypothesis while the training loop only requires a one-shot pruning of model parameters over 20 epochs.

Hence, with all experiments done on a Tesla K80 GPU, the initial neural architecture search took around 2 hours to run for Variant 1 compared to about 5.2 hours for Variant 2; differences in training time between the two variants were negligible.

5.4. Discussion

The analysis on the differences in train time across variants (Section 5.3) suggests that the relative validation accuracy sparsity levels cannot be compared across variants, since one approach has the advantage of having a relatively longer train time. However, general discussions on sparsity levels and training robustness can still give us a strong

Variant	Sparsity Level	Validation Accuracy (%)		
		Epoch 10	Epoch 20	FLOPS
<i>Baseline</i>	N/A	74.04	78.90	752.355M
<i>1</i> <i>(model parameter pruning)</i>	0.125	68.03	70.05	549.204M
	0.25	67.50	72.13	473.034M
	0.375	67.43	74.02	391.527M
	0.5	68.32	76.52	315.356M
	0.625	69.03	76.74	233.849M
	0.75	70.54	76.15	157.678M
	0.875	66.32	67.23	76.171M
<i>2</i> <i>(model architecture pruning)</i>	0.95	66.58	46.43	84.050M
	0.125	72.45	78.77	549.204M
	0.25	72.28	78.04	473.034M
	0.375	74.13	77.47	391.527M
	0.5	69.11	74.12	315.356M
	0.625	67.12	73.60	233.849M
	0.75	71.80	79.86	157.678M
0.875	79.86	81.15	76.171M	
	0.95	80.14	85.59	84.050M

Table 1. Test inference time and validation accuracy across variants and multiple sparsity levels for both variants, compared alongside the baseline model. Winning tickets are shown in **bold**. The result reported for validation accuracy at epochs 10 and 20 were based off of highest validation accuracy over the course of training up until the stated epoch. For Variant 1, 20-epoch performance was consistently strongest on sparsity levels between 0.375-0.75, with performance teetering off at the higher sparsity levels. Variant 2 observes an opposite trend, with performance strongest at low and high sparsity levels.

indication on which variant performs better; there is sufficient evidence to hypothesize how the proposed Algorithm 3 (generalized LNAS) would perform on the CIFAR-10 dataset.

Overall, Variant 1 is more sensitive to sparsity level. We theorize that this is because the initial neural architecture search is done just once and is focused on searching for a general model with a certain range of parameters; as a result, only mid to high sparsity levels such as 0.5, 0.625 and 0.75 were close to reaching the lottery ticket hypothesis. This hypothesis is supported by the sharp drop-off once the sparsity level hits 0.875 and 0.95, when the validation accuracy never eclipses the validation accuracy obtained prior to the pruning and re-initializing of weights. In contrast, Variant 2 does not suffer from the same issue, because the architecture outputted from our lottery ticket hypothesis is optimized on performing well with a lower number of parameters.

With this in mind, it is most likely that the generalized LNAS will perform better than variants 1 and 2 individually. By pruning network architecture parameters (Variant 2) before pruning model architecture parameters (Variant 1), our concern with the sensitivity to sparsity level is mitigated; the model architecture may overall be better suited for the pruning process.

6. Conclusion

Our work shows that with the implementation of the lottery ticket hypothesis alongside a differential architectural search (DARTS) algorithm, we are able to obtain sub-network architecture that achieve state-of-the-art performance without expensive memory or time usage. Furthermore, we show that both parts of a generalized LNAS algorithm – the first part of pruning model parameters, and the second part of pruning architecture parameters – have strong performances in isolation and may be combined in future work to create an even better sub-model.

While this work is a good start in the field of neural network optimization, our future experiments aim to further improve upon training efficiency. We aim to explore the intersectionality of DARTS in other machine learning subfields. We hope that by implementing algorithms such as MetaNAS [2] that combine gradient-based neural architecture search (NAS) methods with gradient-based meta-learning methods, we will be able to solve more challenging problems with just a fraction of the computational and time-related resources required when compared to more traditional approaches.

References

- [1] H. Cai, L. Zhu, and S. Han. Proxylesnas: Direct neural architecture search on target task and hardware, 2019.
- [2] T. Elsken, B. Staffler, J. H. Metzen, and F. Hutter. Meta-learning of neural architectures for few-shot learning, 2020.
- [3] I. Fedorov, R. P. Adams, M. Mattina, and P. N. Whatmough. Sparse: Sparse architecture search for cnns on resource-constrained microcontrollers, 2019.
- [4] J. Frankle and M. Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks, 2019.
- [5] A. Krizhevsky, V. Nair, and G. Hinton. Cifar-10 (canadian institute for advanced research).
- [6] H. Liu, K. Simonyan, and Y. Yang. Darts: Differentiable architecture search, 2019.
- [7] H. Pham, M. Y. Guan, B. Zoph, Q. V. Le, and J. Dean. Efficient neural architecture search via parameter sharing, 2018.
- [8] M. Tan, B. Chen, R. Pang, V. Vasudevan, M. Sandler, A. Howard, and Q. V. Le. Mnasnet: Platform-aware neural architecture search for mobile, 2019.
- [9] R. J. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8:229–256, 1992.
- [10] B. Zoph and Q. V. Le. Neural architecture search with reinforcement learning, 2017.
- [11] B. Zoph, V. Vasudevan, J. Shlens, and Q. V. Le. Learning transferable architectures for scalable image recognition, 2018.