
Learning an Adaptive Learning Rate Schedule

Zhen Xu, Andrew M. Dai, Jonas Kemp, Luke Metz*, Mimi Sun
Google Health, *Google Research
zhenxu, adai, jonasbkemp, lmetz, mimisun@google.com

Abstract

The learning rate is one of the most important hyper-parameters for model training and generalization. However, current hand-designed parametric learning rate schedules offer limited flexibility and the predefined schedule may not match the training dynamics of high dimensional and non-convex optimization problems. In this paper, we propose a reinforcement learning based framework that can automatically learn an adaptive learning rate schedule by leveraging the information from past training histories. The learning rate dynamically changes based on the current training dynamics. To validate this framework, we conduct experiments with different neural network architectures on the Fashion MNIST and CIFAR10 datasets. Experimental results show that the auto-learned learning rate controller can achieve better test results. In addition, the trained controller network is generalizable – able to be trained on one data set and transferred to new problems.

1 Introduction

The learning rate is often regarded as the single most important hyper-parameter to tune and highly influences model training using gradient decent algorithms [1, 2]. Researchers have developed several learning rate schedules such as linear decay, cosine decay, exponential decay, inverse square root decay, etc., sometimes with warm up steps, for different optimization problems [3, 4]. However, there is limited intuition about which learning rate schedule best suits a given problem. In practice, researchers adopt a trial-and-error approach for different learning rate schedules along with different hyper-parameters, which is very time consuming [5]. In this paper, we would like to automatically learn a controller that adapts learning rate schedule by incorporating information from past training dynamics. In addition, current learning rate schedules assume predefined parametric learning rate changes, which are fixed irregardless of actual training dynamics. The optimization landscape can be very complex [6] and these parametric schedules have limited flexibility and may not be optimized for the training dynamics of different high dimensional and non-convex optimization problems. In comparison, our framework offers an auto-learned or meta-learned adaptive learning rate schedule that adapts dynamically based on current training dynamics.

There are several related works proposing better update schedules for gradient descent algorithms. Andrychowicz et al. [7] propose to directly learn the gradient descent updates using a long short-term memory (LSTM) network. Our work only learns the learning rate without gradients and so it is more efficient. Hypergradient takes the derivative of the learning rate and updates the learning rate based on its gradient [8]. In addition to the current state, our approach also considers the entire training history and has a more comprehensive view. Daniel et al. [9] propose to use reinforcement learning (RL) to adapt the learning rate. In comparison, we use validation loss as the reward signal and a learning rate scaling function as the action. They improve the generalization capability and stability. There are a family of widely used optimizers that dynamically adapt the learning rate on a per-parameter basis. For example, Adagrad adapts the learning rate per weight based on the sums of the squares of the gradients [10], while Adam uses an exponentially decayed average of past gradients [11]. However, these optimizers still require a global learning rate which is important to tune. Our work is complementary to these works.

This paper makes three main contributions: First, we propose a reinforcement learning based framework to automatically learn an adaptive learning rate schedule based on past training histories. This schedule can adjust the learning rate dynamically to adapt to current training dynamics. Second, we present an effective set of state observation features, reward functions, and actions for the learning rate decision problem. Specifically, different from the previous work, we use validation loss as the reward signal and a learning rate scaling function as the action. Third, we conduct experiments on Fashion MNIST and CIFAR10 datasets with convolutional neural networks (CNN) [12] and residual networks (ResNet) [13, 14] to show the effectiveness and generalization capability of our framework. The auto-learned learning rate schedule can achieve better results and generalize to different datasets.

2 An Auto-learned Adaptive Learning Rate

2.1 Controller and Trainee Network

We use RL to train a learning rate controller, which proposes learning rates using features from the training dynamics of the trainee network. The trainee network is trained for a certain number of steps using a proposed learning rate, reports the observations of training dynamics to the controller which then returns a new learning rate. The whole process keeps running until reaching a stopping criterion.

2.2 State Observation, Reward, Action

Observation In order to characterize the training dynamics, we design a set of state observation features s , which is an extension of the features used in Daniel et al. [9]. The proposed features include the current train loss, validation loss, variance of network predictions, variance of network prediction changes, mean and variance of the weight matrix of the final dense layer, and the previous step learning rate.

Reward We use the per step validation loss as the reward r . Even though the final validation loss is what we really care about, empirically, we find that providing a reward r_t for each training step can achieve better results compared with just using the final reward. These intermediate rewards provide more direct feedback and make credit assignment much easier.

Action The most direct action of the controller could be proposing a new learning rate. However, the learning rate is very sensitive and could be in the 10^{-6} scale or even smaller. It would be very unstable to directly use the network output as the learning rate. Another choice of action could be proposing the log of the learning rate. However, we want our action to be generalizable to different data sets since they may require different learning rate scales. Instead, we propose a learning rate scaling action a_t . At the first step, we provide a default learning rate. In the following steps, we use the network output as the scaling factor for the previous step learning rate, which can scaling it up or down. In this case, the controller can provide both warm up and decay capabilities in a stable way. This action provides a better inductive bias keeping learning consistent across steps.

2.3 Auto-learned Adaptive Learning Rate Schedule with Proximal Policy Optimization

In this section, we present the RL-based learning rate controller, which is trained to propose a better learning rate schedule in order to reduce the validation loss in the model training process. We use proximal policy optimization (PPO) algorithm [15, 16] to learn the controller parameters θ_c as it has better sample complexity. PPO optimizes a clipped surrogate objective function:

$$J(\theta_c) = E_t \left[\min(w_t(\theta_c)A_t, \text{clip}(w_t(\theta_c), 1 - \epsilon, 1 + \epsilon)A_t) \right] \quad (1)$$

to penalize large policy updates. $w_t(\theta_c)$ is the importance weighting probability ratio. A_t is the advantage function at time t . ϵ is a hyper-parameter.

3 Experiment Setting

3.1 Data Set and Model Architecture

In our experiments, we use two data sets: Fashion MNIST [17] and CIFAR-10 [18]. For Fashion MNIST, we use 50k, 10k, and 10k images as the training, validation and test sets. For CIFAR10, we use 40k, 10k, and 10k images as the training, validation and test sets. We test on two model architectures: CNN and ResNet. CNN follows the structure of LeNet [12]. ResNet follows the structures from [13, 14]. We use the open source TensorFlow implementation [19].

Table 1: Performance comparison between baseline and auto-learned learning rate schedules. We report the mean and standard deviation of the test set loss and accuracy derived from 10 runs with the same set of hyper-parameters. Parentheses denote standard deviations. We bold the best numbers for each model/dataset pair. Asterisk denotes the improvement is statistical significant under independent two-sample t-test with p-value threshold 0.05. Fa. stands for Fashion.

Dataset	Model	Baseline		Auto-learned	
		Test Loss	Test Accuracy	Test Loss	Test Accuracy
Fa. MNIST	CNN	0.2497 (0.0042)	0.9102 (0.0019)	0.2351* (0.0038)	0.9201* (0.0022)
Fa. MNIST	ResNet	0.2346 (0.0074)	0.9188 (0.0029)	0.2296 (0.0069)	0.9192 (0.0028)
CIFAR10	CNN	0.9539 (0.0140)	0.6759 (0.0048)	0.9361* (0.0104)	0.6787 (0.0041)
CIFAR10	ResNet	0.8317 (0.0155)	0.7395 (0.0206)	0.6288* (0.0196)	0.8181* (0.0069)




Figure 1: Training trajectory comparison between baseline and auto-learned learning rate schedules. The y axis is the log of validation loss. The x axis is the number of training steps.

3.2 Baseline Learning Rate Schedule

We use a popular step decay learning rate schedule as our baseline [20], which is used in the open source ResNet implementation [19]. It contains three components: initial learning rate, discount step, and discount factor. The baseline schedule starts from the initial learning rate, then it decreases by the discount factor every discount steps. In the baseline experiments, we test all combinations from the initial learning rate in $[0.1, 0.01, 0.001, 0.0001]$, the discount step in $[10, 20, 50, 100]$, and the discount factor in $[0.99, 0.9, 0.88]$. After choosing the best baseline schedule, we run it 10 times with the same set of hyper-parameters and report mean and standard deviation of test loss and accuracy.

3.3 Training Details

We use validation loss as the reward signal and compare the hold-out test set loss and accuracy between baseline learning rate schedules and our auto-learned schedules. For trainee network training, the initial learning rate is in $[0.01, 0.001, 0.0001]$, the train batch size is 1k, and we run 1k training steps for all model architectures on all data sets. Given 50k training set size, 50 train steps equal to 1 train epoch. In total, we train 20 epochs for Fashion MNIST and 25 epochs for CIFAR10. We will use the number of train steps for future plots. In experiments, the controller proposes a new learning rate every 10 training steps. We refer to the whole 1k train steps of the trainee network as one training episode. The controller network is trained after every training episode of the trainee network. In total, it is trained on 1k episodes. The actor of the controller is a multilayer perceptron (MLP) that contains one hidden layer with size 32. We also test a LSTM actor with hidden size 32. The critic is a MLP that contains one hidden layer with size 32. The learning rate is 0.001 for actor and 0.005 for critic. All networks are trained using Adam optimizer. Note that the main goal of our experiments is to show the effectiveness of the auto-learned learning rate schedule instead of outperforming the state-of-the-art classification accuracy on the target task. In our experiments, we restrict the training epoch to 25 and use ResNet with 18 total layers for computational reasons.

4 Experiment Results

4.1 Test Set Results And Training Trajectories

Table 1 shows the performance comparison between baseline and auto-learned learning rate schedules on test loss and accuracy. We choose the checkpoint based on the best validation loss and evaluate it on the test set once to get the single run test loss and accuracy. From this table, we can see the




Figure 2: Auto-learned learning rate schedule. The x axis is the number of training steps. The y axis is the learning rate.

Table 2: Performance comparison between transferred baseline learning rate schedule and transferred/pre-trained controller network from CIFAR10 to Fashion MNIST. We report the mean and standard deviation of the test set loss and accuracy derived from 10 runs with the same set of hyper-parameters. Asterisk denotes the improvement is statistical significant under independent two-sample t-test with p-value threshold 0.05.

Model	Transferred Baseline		Transferred Controller Network	
	Test Loss	Test Accuracy	Test Loss	Test Accuracy
CNN	0.2730 (0.0031)	0.9021 (0.0013)	0.2598* (0.0071)	0.9074* (0.0030)
ResNet	0.2443 (0.0040)	0.9166 (0.0025)	0.2315* (0.0072)	0.9212* (0.0034)

auto-learned learning rate schedule achieves better results on all tasks. We hypothesize because it does not follow the predefined parametric learning rate changes and has the higher flexibility to adapt the learning rate based on the training dynamics. Figure 1 shows the training dynamics comparison between baseline and auto-learned learning rate schedules in terms of the log of validation loss. For CNN on Fashion MNIST and ResNet on CIFAR10 in Figure 1(a) and Figure 1(d), the auto-learned schedule achieves lower validation loss faster.

4.2 Auto-learned Learning Rate Schedule

In this section, we present the auto-learned learning rate schedules as shown in Figure 2. For the CNN model on Fashion MNIST in Figure 2(a) and ResNet model on CIFAR10 in Figure 2(d), the controller learns to warm up first and then decay the learning rate. This fits human’s intuition when designing some existing learning rate schedules, but it is automatically learned and does not follow a predefined trajectory. For ResNet model on Fashion MNIST in Figure 2(b), the auto-learned learning rate schedule is similar to an exponential decay learning rate schedule. For CNN model on CIFAR10 in Figure 2(c), the auto-learned learning rate keeps flat at first, then warms up and decays later.

4.3 Transferability

We also investigate the transferability of a pre-trained controller network when applied to different datasets. In the experiments, we load the controller checkpoints of CNN and ResNet models on CIFAR10, and let them propose the learning rate schedule for CNN and ResNet models on Fashion MNIST accordingly without training the learning rate controller. In comparison, we also apply the best baseline learning rate schedules of models on CIFAR10 to Fashion MNIST. From Table 2 we can see the trained controller is transferable between two data sets. Our learning rate controller does not simply memorize, but is able to learn a transferable procedure to tune learning rate that generalizes.

5 Conclusion

In this paper, we propose a reinforcement learning based framework which can auto-learn an adaptive learning rate schedule based on the information from past training histories. In order to achieve this goal, we introduce an effective set of features to characterize the dynamic training process, meaningful reward function, action space, and a sample-efficient RL algorithm to adapt the learning rate dynamically. Experimental results on Fashion MNIST and CIFAR10 data sets with CNN and ResNet models show our framework can learn a better learning rate schedule compared with step decay baseline schedules and can be transferred to new datasets never seen during meta-training.

References

- [1] Yoshua Bengio. Practical recommendations for gradient-based training of deep architectures. *Neural networks: Tricks of the trade*, pages 437–478, 2012.
- [2] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep learning*. MIT press, 2016.
- [3] Tom Schaul, Sixin Zhang, and Yann LeCun. No more pesky learning rates. In *International Conference on Machine Learning*, pages 343–351, 2013.
- [4] Matthew D Zeiler. Adadelata: an adaptive learning rate method. *arXiv preprint arXiv:1212.5701*, 2012.
- [5] James Bergstra and Yoshua Bengio. Random search for hyper-parameter optimization. *Journal of Machine Learning Research*, 13(Feb):281–305, 2012.
- [6] H. Li, Z. Xu, G. Taylor, C. Studer, and T. Goldstein. Visualizing the loss landscape of neural nets. In *Advances in Neural Information Processing Systems*, 2018.
- [7] Marcin Andrychowicz, Misha Denil, Sergio Gomez, Matthew W Hoffman, David Pfau, Tom Schaul, Brendan Shillingford, and Nando De Freitas. Learning to learn by gradient descent by gradient descent. In *Advances in Neural Information Processing Systems*, pages 3981–3989, 2016.
- [8] Atilim Gunes Baydin, Robert Cornish, David Martinez Rubio, Mark Schmidt, and Frank Wood. Online learning rate adaptation with hypergradient descent. *arXiv preprint arXiv:1703.04782*, 2017.
- [9] Christian Daniel, Jonathan Taylor, and Sebastian Nowozin. Learning step size controllers for robust neural network training. In *Thirtieth AAAI Conference on Artificial Intelligence*, 2016.
- [10] John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12(Jul):2121–2159, 2011.
- [11] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *The International Conference on Learning Representations (ICLR)*, 2015.
- [12] Yann LeCun, Léon Bottou, Yoshua Bengio, Patrick Haffner, et al. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European conference on computer vision*, pages 630–645. Springer, 2016.
- [15] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [16] John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In *International Conference on Machine Learning*, pages 1889–1897, 2015.
- [17] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.
- [18] Alex Krizhevsky and Geoff Hinton. Convolutional deep belief networks on cifar-10. *Unpublished manuscript*, 40(7), 2010.
- [19] Tensorflow. Resnet in tensorflow. From <https://github.com/tensorflow/models/tree/master/official/resnet>, 2019.
- [20] Rong Ge, Sham M Kakade, Rahul Kidambi, and Praneeth Netrapalli. The step decay schedule: A near optimal, geometrically decaying learning rate procedure. *arXiv preprint arXiv:1904.12838*, 2019.