MLR-SNET: TRANSFERABLE LR SCHEDULES FOR HETEROGENEOUS TASKS

Anonymous authors
Paper under double-blind review

ABSTRACT

The learning rate (LR) is one of the most important hyper-parameters in stochastic gradient descent (SGD) for deep neural networks (DNN) training and generalization. However, current hand-designed LR schedules need to manually pre-specify a fixed form, which limits their ability to adapt to non-convex optimization problems due to the significant variation of training dynamics. Meanwhile, it always needs to search a proper LR schedule from scratch for new tasks. To address these issues, we propose to parameterize LR schedules with an explicit mapping formulation, called *MLR-SNet*. The learnable structure brings more flexibility for MLR-SNet to learn a proper LR schedule to comply with the training dynamics of DNN. Image and text classification benchmark experiments substantiate the capability of our method for achieving proper LR schedules. Moreover, the meta-learned MLR-SNet is plug-and-play to generalize to new heterogeneous tasks. We transfer our meta-trained MLR-SNet to tasks like different training epochs, network architectures, datasets, especially large scale ImageNet dataset, and achieve comparable performance with hand-designed LR schedules. Finally, MLR-SNet can achieve better robustness when training data are biased with corrupted noise.

1 INTRODUCTION

Stochastic gradient descent (SGD) and its many variants (Robbins & Monro, 1951; Duchi et al., 2011; Zeiler, 2012; Tieleman & Hinton, 2012; Kingma & Ba, 2015), have been served as the cornerstone of modern machine learning with big data. It has been empirically shown that DNN achieves state-of-the-art generalization performance on a wide variety of tasks when trained with SGD (Zhang et al., 2017). Several recent researches observe that SGD tends to select the so-called flat minima (Hochreiter & Schmidhuber, 1997a; Keskar et al., 2017), which seems to generalize better in practice.

Scheduling learning rate (LR) for SGD is one of the most widely studied aspects to help improve the SGD training for DNN. Specifically, it has been experimentally studied how the LR (Jastrzebski et al., 2017) influences mimima solutions found by SGD. Theoretically, Wu et al. (2018a) analyze that LR plays an important role in minima selection from a dynamical stability perspective. He et al. (2019) provide a PAC-Bayes generalization bound for DNN trained by SGD, which is correlated with LR. In a word, finding a proper LR schedule highly influences the generalization performance of DNN, which has been widely studied recently (Bengio, 2012; Schaul et al., 2013; Nar & Sastry, 2018).

There mainly exist three kinds of hand-designed LR schedules: (1) Pre-defined LR policy is mostly used in current DNN training, like decaying or cyclic LR (Gower et al., 2019; Loshchilov & Hutter, 2017), and brings large improvements in training efficiency. Some theoretical works suggested that the decaying schedule can yield faster convergence (Ge et al., 2019; Davis et al., 2019) or avoid strict saddles (Lee et al., 2019; Panageas et al., 2019) under some mild conditions. (2) LR search methods in tranditional convex optimization (Nocedal & Wright, 2006) can be extended to DNN training by searching LR adaptively in each step, such as Polyak's update rule (Rolinek & Martius, 2018), Frank-Wolfe algorithm (Berrada et al., 2019), and Armijo line-search (Vaswani et al., 2019), etc. (3) Adaptive gradient methods like Adam (Duchi et al., 2011; Tieleman & Hinton, 2012; Kingma & Ba, 2015), adapt LR for each parameters separately according to some gradient information.

Although above LR schedules (as depicted in Fig. 1(a) and 1(b)) can achieve competitive results on their learning tasks, they still have evident deficiencies in practice. On the one hand, these policies need to manually pre-specify the form of LR schedules, suffering from the limited flexibility to adapt to non-convex optimization problems due to the significant variation of training dynamics. On the other hand, when solving new heterogeneous tasks, it always needs to search a proper LR

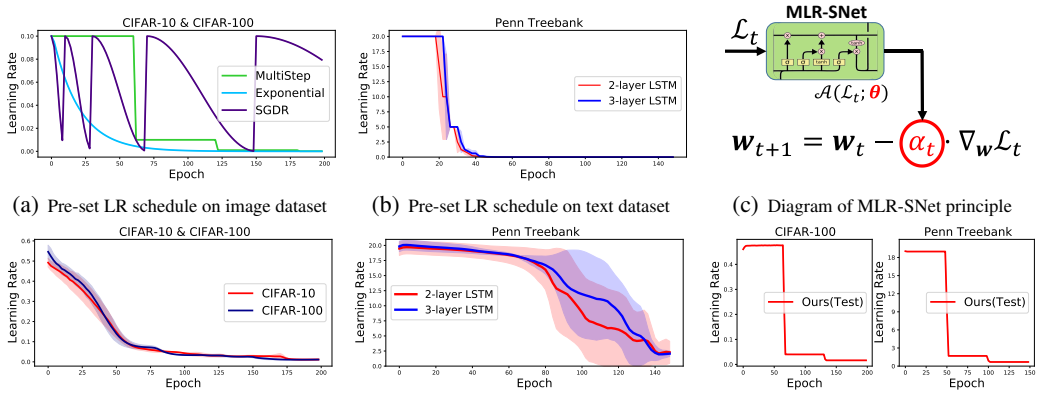

(a) Pre-set LR schedule on image dataset (b) Pre-set LR schedule on text dataset (c) Diagram of MLR-SNet principle

(d) LR schedule learned on image dataset (e) LR schedule learned on text dataset (f) Predicted LR schedules by MLR-SNet

Figure 1: Pre-set LR schedules for (a) image and (b) text classification. (c) Visualization of how we input current loss $\mathcal{L}_t$ to MLR-SNet, which then outputs a proper LR α_t to help SGD find a better minima. LR schedules learned by MLR-SNet on (d) image and (e) text classification. (f) We transfer LR schedules learned on CIFAR-10 to image (CIFAR-100) and text (Penn Treebank) classification, and the subfigure shows the predicted LR during training.

schedule from scratch, as well as to tune their involving hyper-parameters. This process is time and computation expensive, which tends to further raise their application difficulty in real problems.

To alleviate the aforementioned issues, this paper presents a model to learn a plug-and-play LR schedule. The main idea is to parameterize the LR schedule as a LSTM network (Hochreiter & Schmidhuber, 1997b), which is capable of dealing with such a long-term information dependent problem. As shown in Fig. 1(c), the proposed Meta-LR-Schedule-Net (***MLR-SNet***) learns an explicit loss-LR dependent relationship. In a nutshell, this paper makes the following three-fold contributions.

(1) We propose a MLR-SNet to learn an adaptive LR schedule, which can adjust LR based on current training loss as well as the information delivered from past training histories stored in the MLR-SNet. Due to the parameterized form of the MLR-SNet, it can be more flexible than hand-designed policies to find a proper LR schedule for the specific learning task. Fig.1(d) and 1(e) show our learned LR schedules, which have similar tendency as pre-defined policies, but more variations at their locality. This validates the efficacy of our method for adaptively adjusting LR according to training dynamics.

(2) With an explicit parameterized structure, the meta-trained MLR-SNet can be transferred to new heterogeneous tasks (meta-test stage), including different training epochs, network architectures and datasets. Experimental results verify that our plug-and-play LR schedules can achieve comparable performance, while do not have any hyper-parameters compared with tranditional LR schedules. This potentially saves large labor and computation cost in real world applications.

(3) The MLR-SNet is meta-learned to improve generalization performance on unseen data. We validate that with the guidance of clean data, our MLR-SNet can achieve better robustness when training data are biased with corrupted noise than hand-designed LR schedules.

2 Related Work

Meta learning for optimization. Meta learning has a long history in psychology (Ward, 1937; Lake et al., 2017). Meta learning for optimization can date back to 1980s-1990s (Schmidhuber, 1992; Bengio et al., 1991), aiming to meta-learn the optimization process of learning itself. Recently, Andrychowicz et al. (2016); Ravi & Larochelle (2017); Chen et al. (2017); Wichrowska et al. (2017); Li & Malik (2017); Lv et al. (2017) have attempted to scale this idea to larger DNN optimization problems. The main idea is to construct a meta-learner as the optimizer, which takes the gradients as input and outputs the whole updating rules. These approaches tend to make selecting appropriate training algorithms, scheduling LR and tuning other hyper-parameters in an automatic way. Except for solving continuous optimization problems, some works employ these ideas to other optimization problems, such as black-box functions (Chen et al., 2017), model's curvature (Park & Oliva, 2019), evolution strategies (Houthooft et al., 2018), combinatorial functions (Rosenfeld et al., 2018), etc.

Though faster in decreasing training loss than the traditional optimizers in some cases, the learned optimizers may not always generalize well to diverse problems, especially longer horizons (Lv et al., 2017) and large scale optimization problems (Wichrowska et al., 2017). Moreover, they can not be guaranteed to output a proper descent direction in each iteration for DNN training, since they assume

all parameters share one small net and ignore the relationship between each parameters. Our proposed method attempts to learn an adaptive LR schedule rather than the whole update rules. This makes it easy to learn and the meta-learned LR schedule can be transferred to new heterogeneous tasks.

HPO and LR schedule adaptation. Hyper-parameter optimization (HPO) was historically investigated by selecting proper values for algorithm hyper-parameters to obtain better performance on validation set (see (Hutter et al., 2019) for an overview). Typical methods include grid search, random search (Bergstra & Bengio, 2012), Bayesian optimization (Snoek et al., 2012), gradient-based methods (Franceschi et al., 2017; Shu et al., 2020), etc. Recently, some works attempt to find a proper LR schedule under the framework of gradient-based HPO, which can be solved by bilevel optimization (Franceschi et al., 2017; Baydin et al., 2018). However, most HPO techniques tend to fall into short-horizon bias and easily find a bad minima (Wu et al., 2018b). Our MLR-SNet has an explicit function form, which makes the optimization of the LR schedules more robust and effective.

Transfer to heterogeneous tasks. Transfer learning (Pan & Yang, 2009) aims to transfer knowledge obtained from source task to help the learning on the target task. Most transfer learning methods assume the source and target tasks consist of the same instance, feature or model spaces (Yang et al., 2020), which greatly limits their applications. Recently, meta learning (Finn et al., 2017) aims to learn common knowledge shared over a distribution of tasks, such that the learned knowledge can transfer to unseen heterogeneous tasks. Most meta learning approaches focus on few shot learning framework, while we attempt to extend it into a standard learning framwork. The hand-designed LR schedules and HPO methods just try to find a proper LR schedule for given tasks, and need to be learned from scratch for new tasks. However, our meta-learned MLR-SNet is plug-and-play, which can directly transfer how to schedule LR for SGD to heterogeneous tasks without additional learning.

3 The Proposed Meta-LR-Schedule-Net (MLR-SNet) Method

The problem of training DNN can be formulated as the following non-convex optimization problem,

$$\min_{w\in\mathbb{R}^n} \mathcal{L}_{Tr}(D_{Tr};w) := \frac{1}{N}\sum_{i=1}^{N}\mathcal{L}_i^{Tr}(w), \tag{1}$$

where $\mathcal{L}_i^{Tr}$ is the training loss function for data samples $i \in D_{Tr} = \{1, 2, \cdots, N\}$, which characters the deviation of the model prediction from the data, and $w \in \mathbb{R}^n$ represents the parameters of the model (e.g., the weight matrices in DNN) to be optimized. SGD (Robbins & Monro, 1951; Polyak, 1964) and its variants, including Momentum (Tseng, 1998), Adagrad (Duchi et al., 2011), Adadelta (Zeiler, 2012), RMSprop (Tieleman & Hinton, 2012), Adam (Kingma & Ba, 2015), are often used for training DNN. In general, these algorithms can be summarized as the following formulation,

$$w_{t+1} = w_t + \Delta w_t, \Delta w_t = \mathcal{O}_t(\nabla\mathcal{L}^{Tr}(w_t), \mathcal{H}_t; \Theta_t), \tag{2}$$

where w_t is t-th updating model parameters, $\nabla\mathcal{L}^{Tr}(w_t)$ denotes the gradient of $\mathcal{L}^{Tr}$ at w_t, $\mathcal{H}_t$ represents the historical gradient information, and Θ_t is the hyperparameter of the optimizer $\mathcal{O}$, e.g., LR. To present our method's efficiency, we focus on the following vanilla SGD formulation,

$$w_{t+1} = w_t - \alpha_t\left(\frac{1}{|B_t|}\sum_{i\in B_t}\nabla\mathcal{L}_i^{Tr}(w_t)\right), \tag{3}$$

where $B_t \subset D_{Tr}$ denotes the batch samples randomly sampled from the training dataset, $|B_t|$ denotes the number of the sampled batch samples, and $\nabla\mathcal{L}_i^{Tr}(w_t)$ denotes the gradient of sample i computed at w_t and α_t is the LR at t-th iteration.

3.1 Existing LR schedule strategies

As Bengio (2012) demonstrated, the choice of LR remains central to effective DNN training with SGD. As mentioned in Section 1, a variety of hand-designed LR schedules have been proposed. Though they achieve competitive results on some learning tasks, they share several drawbacks: (1) The pre-defined LR schedules suffer from the limited flexibility to adapt to the significantly changed training dynamics for the non-convex optimization problems. (2) It needs to be learned from scratch to find a proper LR schedule for the new tasks, which raises their application difficulty in real problems.

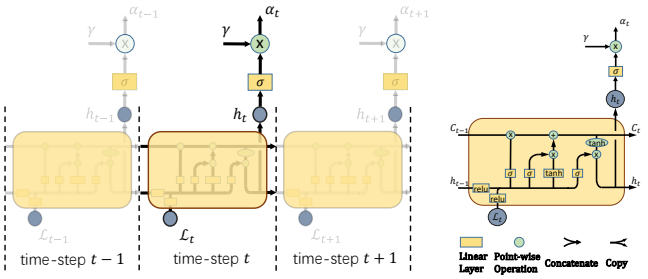

(a) Computational graph of MLR-SNet (b) One step of MLR-SNet

Figure 2: The structure of our proposed MLR-SNet.

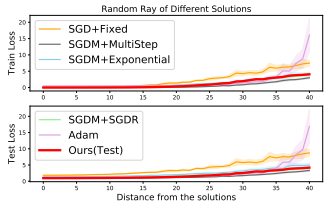

Figure 3: (**Above**) Train loss and (**Below**) test loss as a function of a point on a random ray starting at the solutions for different methods on CIFAR-100 with ResNet-18.

Inspired by current meta-learning developments (Finn et al., 2017; Shu et al., 2018; 2019), some researches proposed to learn a generic optimizer from data (Andrychowicz et al., 2016; Ravi & Larochelle, 2017; Chen et al., 2017; Wichrowska et al., 2017; Li & Malik, 2017; Lv et al., 2017). The main idea is to learn a meta-learner as the optimizer to guide the learning of the whole updating rules. For example, Andrychowicz et al. (2016) try to replace Eq.(2) with the following formulation,

$$w_{t+1} = w_t + g_t, [g_t, h_{t+1}]^T = m(\nabla_t, h_t; \phi), \tag{4}$$

where g_t is the output of a LSTM net m, parameterized by ϕ, whose state is h_t.This strategy can make selecting appropriate training algorithms, scheduling LR and tuning other hyper-parameters in a unified and automatic way. Though faster in decreasing training loss than the traditional optimizers in some cases, the learned optimizer may not always generalize well to more variant and diverse problems, like longer horizons (Lv et al., 2017) and large scale optimization problems (Wichrowska et al., 2017). Moreover, it can not guarantee to output a proper descent direction in each iteration for network training. This tends to further increase their application difficulty in real problems.

Recently, some methods (Franceschi et al., 2017; Baydin et al., 2018) consider the following constrained optimization problem to search the optimal LR schedule α^* such that the produced models are associated with small validation error,

$$\min_{\alpha=\{\alpha_0,\cdots,\alpha_{T-1}\}} \mathcal{L}_{Val}(D_{Val}, w_T),\ s.t.\ w_{t+1} = \phi_t(w_t, \alpha_t),\ t = 0, 1, \cdots, T-1, \tag{5}$$

where $\mathcal{L}_{Val}$ denotes the validation loss function, $D_{Val} = \{1, 2, \cdots, M\}$ denotes hold-out validation set, α is to-be-solved hyper-parameter, $\phi_t : \mathbb{R}^n \times \mathbb{R}_+ \to \mathbb{R}^n$ is a stochastic weight update dynamics, like the updating rule in Eq.(2) or the vanilla SGD in Eq.(3), and T is the maximum iteration step. Though achieving comparable results on some tasks with hand-designed LR schedules, they can not directly transfer to new tasks, since they do not have an explict transferable structure form.

3.2 Proposed Meta-LR-Schedule-Net (MLR-SNet) Method

To address aforementioned issues, the main idea is to design a meta-learner with an explicit mapping formulation to parameterize LR schedules as shown in Fig.1(c), called MLR-SNet. The parameterized structure can bring two benefits: 1) It gives more flexibility to learn a proper LR schedule to comply with the significantly changed training dynamics of DNN; 2) It makes the meta-learned LR schedules be transferable and plug-and-play, which can be applied to new heterogeneous tasks.

Formulation of MLR-SNet. The computational graph of MLR-SNet is depicted in Fig.2(a). Let $\mathcal{A}(\cdot;\theta)$ denote the MLR-SNet, and then the updating equation of SGD in Eq.(3) can be rewritten as

$$w_{t+1} = w_t - \mathcal{A}(\mathcal{L}_t; \theta_t)\left(\frac{1}{|B_t|}\sum_{i\in B_t}\nabla\mathcal{L}_i^{Tr}(w_t)\right), \mathcal{L}_t = \frac{1}{|B_t|}\sum_{i\in B_t}\mathcal{L}_i^{Tr}(w_t), \tag{6}$$

where θ_t is the parameter of MLR-SNet at t-th iteration ($t = 0, \cdots, T-1$). At any iteration steps, $\mathcal{A}(\cdot;\theta)$ can learn an explicit loss-LR dependent relationship, such that the net can adaptively predict LR according to the current input loss $\mathcal{L}_t$, as well as the historical information stored in the net. For every iteration step, the whole forward computation process is (as shown in Fig. 2(b))

$$\begin{pmatrix} i_t \\ f_t \\ o_t \\ g_t \end{pmatrix} = \begin{pmatrix} \sigma \\ \sigma \\ \sigma \\ \tanh \end{pmatrix} W_2 \begin{pmatrix} \text{ReLU} \\ \text{ReLU} \end{pmatrix} W_1 \begin{pmatrix} h_{t-1} \\ \mathcal{L}_t \end{pmatrix}, \begin{array}{l} c_t = f_t \odot c_{t-1} + i_t \odot g_t \\ h_t = o_t \odot \tanh(c_t) \\ p_t = \sigma(W_3 h_t) \\ \alpha_t = \gamma \cdot p_t \end{array}, \tag{7}$$

where i_t, f_t, o_t denote the Input, Forget and Output gates, respectively. Different from vanilla LSTM, the input h_{t-1} and the training loss $\mathcal{L}_t$ are preprocessed by a fully-connected layer W_1 with ReLU activation function. Then it works as LSTM and obtains the output h_t. After that, the predicted value p_t is obtained by a linear transform W_3 on the h_t with a Sigmoid activation function. Finally, we introduce a scale factor γ [1] to guarantee the final predicted LR located in the interval of $[0, \gamma]$. Albeit simple, this net is known for dealing with such long-term information dependent problem, and thus capable of finding a proper LR schedule to comply with the complex variations of training dynamics.

Meta-Train: adapting to the training dynamics of DNN. The MLR-SNet can be meta-trained to improve the generalization performance on unseen validation data for DNN training as follows:

$$\min_{\theta} \mathcal{L}_{Val}(D_{Val}, w_T),\ s.t.\ w_{t+1} = \phi_t(w_t, \mathcal{A}(\mathcal{L}_t; \theta)),\ t = 0, 1, \cdots, T-1. \tag{8}$$

Now the important question is how to efficiently meta-learn the parameter θ for the MLR-SNet. We employ the online approximation technique in (Shu et al., 2019) to jointly update θ and model parameter w to explore a proper LR schedule with better generalization for DNNs training. However, the step-wise optimization for θ is still expensive to handle large-scale datasets and DNN. Furthermore, we attempt to update θ after updating w several steps (T_{val}) as summarized in Algorithm 1.

Updating θ. When it does not satisfy the updating conditions, θ is fixed; otherwise, θ is updated using the model parameter w_t and MLR-SNet parameter θ_t obtained in the last step by minimizing the validation loss defined in Eq.(8). Adam can be employed to optimize the validation loss, i.e.,

$$\theta_{t+1} = \theta_t + Adam(\nabla_\theta \mathcal{L}_{Val}(D_m, \hat{w}_{t+1}(\theta)); \beta_t), \tag{9}$$

where $Adam$ denotes the Adam algorithm, whose input is the gradient of validation loss with respect to MLR-SNet parameter θ on m mini-batch samples D_m from D_{Val}. β_t denotes the LR of Adam. $\hat{w}_{t+1}(\theta)$ [2] is formulated on a mini-batch training samples D_n from D_{Tr} as follows:

$$\hat{w}_{t+1}(\theta) = w_t - \mathcal{A}(\mathcal{L}_{Tr}(D_n, w_t); \theta) \cdot \nabla_w \mathcal{L}_{Tr}(D_n, w)\big|_{w_t}. \tag{10}$$

Updating w. Then, the updated θ_{t+1} is employed to ameliorate the model parameter w, i.e.,

$$w_{t+1} = w_t - \mathcal{A}(\mathcal{L}_{Tr}(D_n, w_t); \theta_{t+1}) \cdot \nabla_w \mathcal{L}_{Tr}(D_n, w)\big|_{w_t}. \tag{11}$$

The whole meta-train learning algorithm can be summarized in Algorithm 1. All computations of gradients can be efficiently implemented by automatic differentiation libraries, like PyTorch (Paszke et al., 2019), and generalized to any DNN architectures. It can be seen that the MLR-SNet can be gradually optimized during the learning process and adjust the LR dynamically based on the training dynamics of DNNs.

Algorithm 1 The Meta-Train Algorithm of MLR-SNet

Input: Training data D_{Tr}, validation set D_{Val}, batch size n, m, max iterations T, updating period T_{val}.
Output: Model parameter w_T and MLR-SNet parameter θ_T

```
Initialize model parameter w_0 and MLR-SNet parameter θ_0.
for t = 0 to T − 1 do
    D_n ← SampleMiniBatch(D_Tr, n).
    if t % T_val = 0, then
        D_m ← SampleMiniBatch(D_Val, m).
        Update θ_{t+1} by Eq. (9).
    end if
    Update w_{t+1} by Eq. (12).
end for
```

Meta-Test: transferring to heterogeneous tasks. When we obtain the meta-learned MLR-SNet, it can be easily applied to new tasks. Now the new model parameter u for the new task is updated by,

$$u_{t+1} = u_t - \mathcal{A}(\mathcal{L}_{Tr}(D_n, u_t); \theta^*) \cdot \nabla_u \mathcal{L}_{Tr}(D_n, u)\big|_{u_t}, \tag{12}$$

where θ^* is the parameter of the meta-learned MLR-SNet, which is fixed in the meta-test stage.

4 Experimental Results

To evaluate the proposed MLR-SNet, we firstly conduct experiments to show our method is capable of finding proper LR schedules compared with baseline methods. Then we transfer the learned LR schedules to various tasks to show its superiority in generalization. Finally, we show our method behaves robust and stable when training data contain different data corruptions.

[1] As we know that the performance of hand-designed LR schedules and HPO methods is very sensitive to the initial LR. To avoid carefully tuning the initial LR, we learn the LR schedules from an interval $[0, \gamma]$, and now the initial LR is determined by the output of the MLR-SNet. We set $\gamma = 1$ for image tasks, and $\gamma = 40$ for text tasks in all our experiments to eliminate the influence of loss magnitude between two different tasks.

[2] Notice that $\hat{w}_{t+1}(\theta)$ here is a function of θ to guarantee the gradient in Eq.(9) to be able to compute.

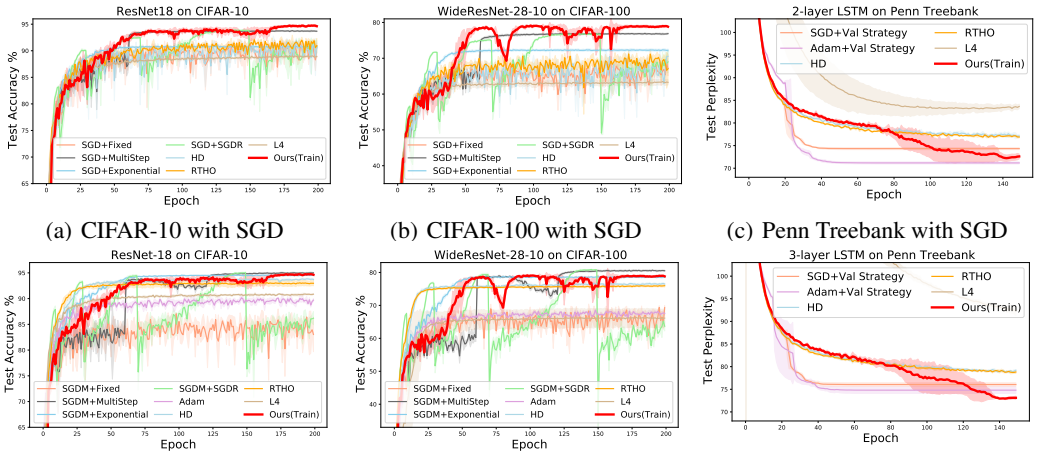

(a) CIFAR-10 with SGD (b) CIFAR-100 with SGD (c) Penn Treebank with SGD

(d) CIFAR-10 with SGDM (e) CIFAR-100 with SGDM (f) Penn Treebank with SGD

Figure 4: Test accuracy of our methods (train) and compared baselines on different datasets.

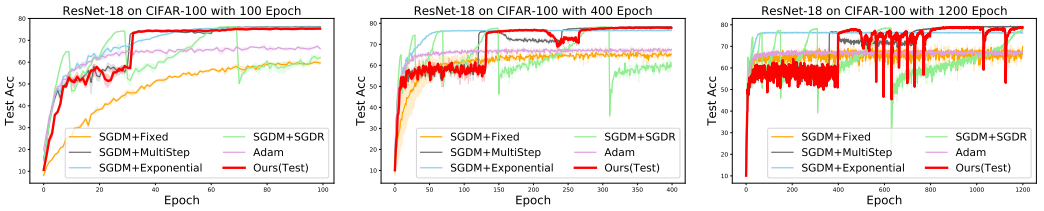

(a) Epoch 100 (b) Epoch 400 (c) Epoch 1200

Figure 5: Test accuracy on CIFAR-100 of ResNet-18 with varying epochs.

4.1 META-TRAIN: EVALUATION ON THE LR SCHEDULE LEARNED BY MLR-SNET

Datasets and models. To verify general effectiveness of our method, we respectively train different models on four benchmark data, including ResNet-18 (He et al., 2016) on CIFAR-10, WideResNet-28-10 (Zagoruyko & Komodakis, 2016) on CIFAR-100 (Krizhevsky, 2009), 2-layer LSTM and 3-layer LSTM on Penn Treebank (Marcus & Marcinkiewicz).

Baselines. For image classification tasks, the compared methods include SGD with hand-designed LR schedules: 1) **Fixed** LR, 2) **Exponential** decay, 3) **MultiStep** decay, 4) SGD with restarts (**SGDR**) (Loshchilov & Hutter, 2017). Also, we compare with SGD with Momentum (SGDM) with above four LR schedules. The momentum is fixed as 0.9. Meanwhile, we compare with adaptive gradient method: 5)**Adam**, LR search method: 6) **L4** (Rolinek & Martius, 2018), and current LR schedule adaptation methods: 7) hyper-gradient descent (**HD**) (Baydin et al., 2018), 8) real-time hyper-parameter optimization (**RTHO**) (Franceschi et al., 2017). For text classification tasks, we compare with 1) SGD and 2) Adam with LR tuned using a validation set. They drop the LR by a factor of 4 when the validation loss stops decreasing. Also, we compared with 3) **L4**, 4) **HD**, 5) **RTHO**. We run all experiments with 3 different seeds reporting accuracy. The detailed illustrations of experimental setting, and more experimental results are presented in supplementary material.

Image tasks. Fig.4(a) and 4(b) show the classification accuracy on CIFAR-10 and CIFAR-100 test sets, respectively. It can be observed that: 1) our algorithm outperforms all other competing methods, and the learned LR schedules by MLR-SNet are presented in Fig.1(d), which have similar shape as the hand-designed policies, while with more elaborate variation details in locality for adapting training dynamics. 2) The Fixed LR has similar performance to other baselines at the early training, while falls into fluctuations at the later training. This implies that the Fixed LR can not finely adapt to such DNN training dynamics. 3) The MultiStep LR drops the LR at some epochs, and such elegant strategy overcomes the issue of Fixed LR and obtains higher and stabler performance at the later training. 3) The Exponential LR improves test performance faster at the early training than other baselines, while makes a slow progress due to smaller LR at the later training. 4) SGDR uses the cyclic LR, which needs more epochs to obtain a stable result. 5) Though Adam has an adaptive coordinate-specific LR, it behaves worse than MultiStep and Exponential LR as demonstrated in Wilson et al. (2017). An extra tuning is necessary for better performance. 6) L4 greedily searches LR locally to decrease loss, while the complex DNN training dynamics can not guarantee it to obtain a good minima. 7) HD and RTHO are able to achieve similar performance to hand-designed LR schedules. The LR schedules learned by L4, HD and RTHO can be found in supplementary material.

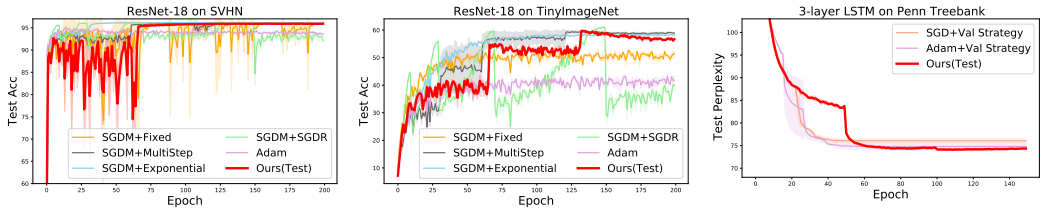

(a) SVHN (b) TinyImageNet (c) Penn Treebank

Figure 6: Test accuracy of transferred LR schedules on different datasets.

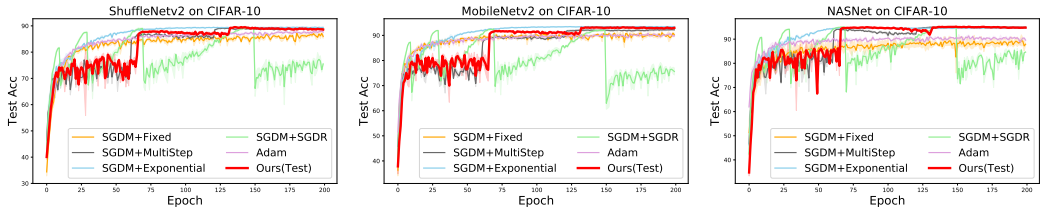

(a) ShuffleNetV2 (b) MobileNetV2 (c) NASNet

Figure 7: Test accuracy on CIFAR-10 of different network architectures

Since image tasks often use SGDM to train DNNs, Fig.4(d) and 4(e) show the results of baseline methods trained with SGDM, and they obtain a remarkable improvement than SGD. Though not using extra historical gradient information to help optimization, our method achieves comparable results with baselines by finding a proper LR schedule for SGD.

Text tasks. Fig.4(c) and 4(f) show the test perplexity on the Penn Treebank with 2-layer and 3-layer LSTM, respectively. Adam and SGD heuristically drops LR when the validation loss stops decreasing. However, our MLR-SNet predicts LR according to training dynamics by minimizing the validation loss, which is a more intelligent way to employ the validation dataset. Thus our method achieves comparable or even better performance than Adam and SGD. The learned LR schedules of the MLR-SNet are presented in Fig.1(b), which have similar shape as the hand-designed policies. L4 often falls into a bad minima since it greedily searches LR locally. HD and RTHO directly optimize LR to improve the performance on validation dataset, obtaining the similar results as Adam and SGD. With an explicit strcuture, our method behaves more robust and efficient than HD and RTHO.

4.2 Meta-Test: Transferability of Plug-and-Play LR Schedules

The learned MLR-SNet is transferable and plug-and-play. Here we validate if MLR-SNet can transfer to new heterogeneous tasks. Since the methods L4,HD,RTHO in Section 4.1 are not able to generalize, we do not compare them here. Actually our results show superiority on image tasks beyond baseline methods when trained with SGD, here we present stronger baseline results in which compared methods are trained with SGDM. We use the MLR-SNet meta-learned on CIFAR-10 with ResNet-18 in Section 4.1 as the plug-and-play LR schedules for the following experiments.

Transfer to different epochs. The plug-and-play MLR-SNet is meta-trained with epoch 200, and we transfer it to other different training epochs, e.g., 100, 400,1200. As shown in Fig.5, our MLR-SNet has the ability to train for longer horizons and achieves almost same performance as MultiStep LR. The slight shakes for epoch 1200 may due to that our MLR-SNet can learn the LR similar to SGDR locally. The Exponential LR has a little performance decreased for the longer epochs.

Transfer to different datasets. We transfer the LR schedules meta-learned on CIFAR-10 to SVHN (Netzer et al., 2011), TinyImageNet [3], and Penn Treebank (Marcus & Marcinkiewicz). As shown in Fig.6, though datasets vary from image to text, our method can still obtain a relatively stable and comparable generalization performance for different tasks with baseline method.

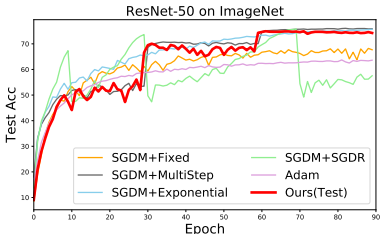

Figure 8: Test accuracy on ImageNet.

Transfer to different net architectures. We also transfer the LR schedules meta-learned on ResNet-18 to light-weight nets ShuffleNetV2 (Ma et al., 2018), MobileNetV2 (Sandler et al., 2018) or NASNet (Zoph et al., 2018) [4]. As shown in Fig.7, our method achieves almost similar results to SGDM with MultiStep or Exponential LR.

[3] It can be downloaded at https://tiny-imagenet.herokuapp.com.

[4] The pytorch code of all these networks can be found on https://github.com/weiaicunzai/pytorch-cifar100.

Table 1: Test accuracy (%) on CIFAR-10 and CIFAR-100 training set of different methods trained on CIFAR-10-C and CIFAR-100-C. Best and Last denote the results of the best and the last epoch.

Datasets/Methods		Fixed	MultiStep	Exponential	SGDR	Adam	Ours(Train)
CIFAR-10-C	Best	79.78±3.95	85.52±1.72	83.48±1.45	85.94±1.52	81.45±1.42	**86.04±1.51**
	Last	77.88±3.91	85.36±1.71	83.32±1.43	78.21±2.01	80.29±1.64	**85.87±1.54**
CIFAR-100-C	Best	46.74±3.03	52.26±2.58	49.72±1.97	52.54±2.49	45.45±1.94	**52.56±2.26**
	Last	44.79±3.91	52.16±2.59	49.58±1.98	41.58±3.24	43.76±2.22	**52.42±2.34**

Transfer to large scale optimization problem. To our best knowledge, only Wichrowska et al. (2017) attempted to use the learned optimizers to train DNN on ImageNet dataset (Deng et al., 2009) among existing learning-to-optimize literatures. However, it can only be executed for thousands of steps, and then its loss begins to increase dramatically, far from the optimization process in practice. We transfer the LR schedule meta-trained on CIFAR-10 with ResNet-18 to ImageNet dataset with ResNet-50 [5]. As shown in Fig.8, the validation accuracy of our method is competitive with those hand-designed LR schedules methods. This implies our method is capable of dealing with such large scale optimization problem, making learning-to-optimize ideas towards more practical applications.

4.3 Robustness on Different Data Corruptions

In this section, we validate that whether our MLR-SNet behaves robust against corrupted training data. To this aim, we design experiments as follows: we take CIFAR-10-C and CIFAR-100-C (Hendrycks & Dietterich, 2019) as our training set, consisting of 15 types of different generated corruptions on test images data of CIFAR-10/CIFAR-100, and the original training set of CIFAR-10/100 as test set. Though the original images of CIFAR-10/100-C are the same with the CIFAR-10/100 test set, different corruptions have changed the data distributions. To guarantee the calculated models finely generalize to test set, we choose the validation set as 10 clean images for each class. Each corruption can be roughly regarded as a task, and thus we obtain 15 models trained on CIFAR-10/100-C. Table 1 shows the mean test accuracy of 15 models ($\pm$std). As can be seen, our proposed MLR-SNet is capable of achieving better generalization performance on clean test data than baseline methods, which implies that our method behaves more robust and stable than the pre-set LR schedules when the learning tasks are changed. This is due to that our MLR-SNet has more flexibility to adapt the variation of the data distribution than the pre-set LR schedules, and it can find a proper LR schedule through minimizing the generalization error which is based on the knowledge specifically conveyed from the given validation data. The detailed illustrations of experimental setting, and the transferrablity experiment of meta-learned MLR-SNet are presented in supplementary material.

5 Conclusion and Discussion

In this paper, we have proposed to learn an adaptive and transferrable LR schedule in a meta learning manner. To this aim, we design an LSTM-type meta-learner (MLR-SNet) to parameterize LR schedules, which gives more flexibility to adaptively learn a proper LR schedule to comply with the significantly complex training dynamics of DNN. Meanwhile, the meta-learned LR schedules are plug-and-play and transferrable, which can be transferred how to schedule LR for SGD to new heterogeneous tasks. Comprehensive experiments substantiate the superiority of our method on various image and text benchmarks in its adaptability, transferrability and robustness, as compared with current LR schedules policies. The MLR-SNet is highly practical as it requires negligible increase in the parameter size and computation time (see supplementary material), and no transferrable cost for new tasks. We believe our proposed method has a potential to become a new tool to study how to design LR schedules to help improve current DNN training, as well as more practical applications.

The preliminary experimental evaluations show that our method gives good convergence performance on various tasks. We find that the meta-learned LR schedules in our experiments follow a consistent trajectory as shown in Fig.1, sharing a similar tendency as the pre-set LR schedules. The convergence guarantee in (Li et al., 2020) can roughly explain our good convergence performance for such DNN training. The detailed theoretical analysis for convergence of our method is left for further work. Furthermore, Keskar et al. (2017); Dinh et al. (2017) suggested that the width of a local optimum is related to generalization. Wider optima leads to better generalization. We use the visualization technique in Izmailov et al. (2018) to visualize the "width" of the solutions for different LR schedules on CIFAR-100 with ResNet-18. As shown in Fig.3, our method lies a wide flat region of the train loss. This could explain the better generalization of our method compared with pre-set LR schedules. Deeper understandings on this point will be further investigated.

[5]The training code can be found on https://github.com/pytorch/examples/tree/master/imagenet.

REFERENCES

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 *NeurIPS*, 2016.

Atilim Gunes Baydin, Robert Cornish, David Martinez Rubio, Mark Schmidt, and Frank Wood. Online learning rate adaptation with hypergradient descent. In *ICLR*, 2018.

Y Bengio, S Bengio, and J Cloutier. Learning a synaptic learning rule. In *IJCNN*, volume 2, pp. 969–vol. IEEE, 1991.

Yoshua Bengio. Practical recommendations for gradient-based training of deep architectures. In *Neural networks: Tricks of the trade*, pp. 437–478. Springer, 2012.

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

Leonard Berrada, Andrew Zisserman, and M Pawan Kumar. Deep frank-wolfe for neural network optimization. In *ICLR*, 2019.

Yutian Chen, Matthew W Hoffman, Sergio Gómez Colmenarejo, Misha Denil, Timothy P Lillicrap, Matt Botvinick, and Nando De Freitas. Learning to learn without gradient descent by gradient descent. In *ICML*, 2017.

Damek Davis, Dmitriy Drusvyatskiy, and Vasileios Charisopoulos. Stochastic algorithms with geometric step decay converge linearly on sharp functions. *arXiv:1907.09547*, 2019.

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

Laurent Dinh, Razvan Pascanu, Samy Bengio, and Yoshua Bengio. Sharp minima can generalize for deep nets. In *ICML*, 2017.

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.

Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*, 2017.

Luca Franceschi, Michele Donini, Paolo Frasconi, and Massimiliano Pontil. Forward and reverse gradient-based hyperparameter optimization. In *ICML*, 2017.

Rong Ge, Sham M Kakade, Rahul Kidambi, and Praneeth Netrapalli. The step decay schedule: A near optimal, geometrically decaying learning rate procedure for least squares. In *NeurIPS*, 2019.

Robert Mansel Gower, Nicolas Loizou, Xun Qian, Alibek Sailanbayev, Egor Shulgin, and Peter Richtárik. Sgd: General analysis and improved rates. In *ICML*, 2019.

Fengxiang He, Tongliang Liu, and Dacheng Tao. Control batch size and learning rate to generalize well: Theoretical and empirical evidence. In *NeurIPS*, 2019.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.

Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. In *ICLR*, 2019.

Sepp Hochreiter and Jürgen Schmidhuber. Flat minima. *Neural Computation*, 9(1):1–42, 1997a.

Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8): 1735–1780, 1997b.

Rein Houthooft, Yuhua Chen, Phillip Isola, Bradly Stadie, Filip Wolski, OpenAI Jonathan Ho, and Pieter Abbeel. Evolved policy gradients. In *NeurIPS*, 2018.

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

Pavel Izmailov, Dmitrii Podoprikhin, Timur Garipov, Dmitry Vetrov, and Andrew Gordon Wilson. Averaging weights leads to wider optima and better generalization. In *UAI*, 2018.

Stanisław Jastrzebski, Zachary Kenton, Devansh Arpit, Nicolas Ballas, Asja Fischer, Yoshua Bengio, and Amos Storkey. Three factors influencing minima in sgd. *arXiv:1711.04623*, 2017.

Nitish Shirish Keskar, Dheevatsa Mudigere, Jorge Nocedal, Mikhail Smelyanskiy, and Ping Tak Peter Tang. On large-batch training for deep learning: Generalization gap and sharp minima. In *ICLR*, 2017.

Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.

Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, 2009.

Brenden M Lake, Tomer D Ullman, Joshua B Tenenbaum, and Samuel J Gershman. Building machines that learn and think like people. *Behavioral and brain sciences*, 40, 2017.

Jason D Lee, Ioannis Panageas, Georgios Piliouras, Max Simchowitz, Michael I Jordan, and Benjamin Recht. First-order methods almost always avoid saddle points. *Mathematical Programming*, 2019.

Ke Li and Jitendra Malik. Learning to optimize neural nets. In *ICLR*, 2017.

Xiaoyu Li, Zhenxun Zhuang, and Francesco Orabona. Exponential step sizes for non-convex optimization. *arXiv preprint arXiv:2002.05273*, 2020.

Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. In *ICLR*, 2017.

Kaifeng Lv, Shunhua Jiang, and Jian Li. Learning gradient descent: Better generalization and longer horizons. In *ICML*, 2017.

Ningning Ma, Xiangyu Zhang, Hai-Tao Zheng, and Jian Sun. Shufflenet v2: Practical guidelines for efficient cnn architecture design. In *ECCV*, 2018.

Mitchell P Marcus and Mary Ann Marcinkiewicz. Building a large annotated corpus of english: The penn treebank. *Computational Linguistics*, 19(2).

Kamil Nar and Shankar Sastry. Step size matters in deep learning. In *NeurIPS*, 2018.

Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. In *NeurIPS Workshop on Deep Learning and Unsupervised Feature Learning*, 2011.

Jorge Nocedal and Stephen Wright. *Numerical optimization*. Springer Science & Business Media, 2006.

Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10):1345–1359, 2009.

Ioannis Panageas, Georgios Piliouras, and Xiao Wang. First-order methods almost always avoid saddle points: The case of vanishing step-sizes. In *NeurIPS*, 2019.

Eunbyung Park and Junier B Oliva. Meta-curvature. In *NeurIPS*, 2019.

Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *NeurIPS*, 2019.

Boris T Polyak. Some methods of speeding up the convergence of iteration methods. *Computational Mathematics and Mathematical Physics*, 4(5):1–17, 1964.

Sachin Ravi and Hugo Larochelle. Optimization as a model for few-shot learning. In *ICLR*, 2017.

Herbert Robbins and Sutton Monro. A stochastic approximation method. *The annals of mathematical statistics*, pp. 400–407, 1951.

Michal Rolinek and Georg Martius. L4: Practical loss-based stepsize adaptation for deep learning. In *NeurIPS*, 2018.

Nir Rosenfeld, Eric Balkanski, Amir Globerson, and Yaron Singer. Learning to optimize combinatorial functions. In *ICML*, 2018.

Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *CVPR*, 2018.

Tom Schaul, Sixin Zhang, and Yann LeCun. No more pesky learning rates. In *ICML*, 2013.

Jürgen Schmidhuber. Learning to control fast-weight memories: An alternative to dynamic recurrent networks. *Neural Computation*, 4(1):131–139, 1992.

Jun Shu, Zongben Xu, and Deyu Meng. Small sample learning in big data era. *arXiv:1808.04572*, 2018.

Jun Shu, Qi Xie, Lixuan Yi, Qian Zhao, Sanping Zhou, Zongben Xu, and Deyu Meng. Meta-weight-net: Learning an explicit mapping for sample weighting. In *NeurIPS*, 2019.

Jun Shu, Qian Zhao, Keyu Chen, Zongben Xu, and Deyu Meng. Learning adaptive loss for robust learning with noisy labels. *arXiv:2002.06482*, 2020.

Jasper Snoek, Hugo Larochelle, and Ryan P Adams. Practical bayesian optimization of machine learning algorithms. In *NeurIPS*, 2012.

Tijmen Tieleman and Geoffrey Hinton. Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude. *Neural networks for machine learning*, 2012.

Paul Tseng. An incremental gradient (-projection) method with momentum term and adaptive stepsize rule. *SIAM Journal on Optimization*, 8(2):506–531, 1998.

Sharan Vaswani, Aaron Mishkin, Issam Laradji, Mark Schmidt, Gauthier Gidel, and Simon Lacoste-Julien. Painless stochastic gradient: Interpolation, line-search, and convergence rates. In *NeurIPS*, 2019.

Lewis B Ward. Reminiscence and rote learning. *Psychological Monographs*, 49(4), 1937.

Olga Wichrowska, Niru Maheswaranathan, Matthew W Hoffman, Sergio Gomez Colmenarejo, Misha Denil, Nando de Freitas, and Jascha Sohl-Dickstein. Learned optimizers that scale and generalize. In *ICML*, 2017.

Ashia C Wilson, Rebecca Roelofs, Mitchell Stern, Nati Srebro, and Benjamin Recht. The marginal value of adaptive gradient methods in machine learning. In *NeurIPS*, 2017.

Lei Wu, Chao Ma, and E Weinan. How sgd selects the global minima in over-parameterized learning: A dynamical stability perspective. In *NeurIPS*, 2018a.

Yuhuai Wu, Mengye Ren, Renjie Liao, and Roger Grosse. Understanding short-horizon bias in stochastic meta-optimization. In *ICLR*, 2018b.

Qiang Yang, Yu Zhang, Wenyuan Dai, and Sinno Jialin Pan. *Transfer learning*. Cambridge University Press, 2020.

Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016.

Matthew D Zeiler. Adadelta: an adaptive learning rate method. *arXiv:1212.5701*, 2012.

Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. In *ICLR*, 2017.

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