AN EXPONENTIAL LEARNING RATE SCHEDULE FOR DEEP LEARNING

Zhiyuan Li
Princeton University
zhiyuanli@cs.princeton.edu

Sanjeev Arora
Princeton University and Institute for Advanced Study
arora@cs.princeton.edu

ABSTRACT

Intriguing empirical evidence exists that deep learning can work well with exotic schedules for varying the learning rate. This paper suggests that the phenomenon may be due to Batch Normalization or BN(Ioffe & Szegedy, 2015), which is ubiquitous and provides benefits in optimization and generalization across all standard architectures. The following new results are shown about BN with weight decay and momentum (in other words, the typical use case which was not considered in earlier theoretical analyses of stand-alone BN (Ioffe & Szegedy, 2015; Santurkar et al., 2018; Arora et al., 2018)

- Training can be done using SGD with momentum and an exponentially *increasing* learning rate schedule, i.e., learning rate increases by some $(1+\alpha)$ factor in every epoch for some $\alpha > 0$. (Precise statement in the paper.) To the best of our knowledge this is the first time such a rate schedule has been successfully used, let alone for highly successful architectures. As expected, such training rapidly blows up network weights, but the network stays well-behaved due to normalization.
- Mathematical explanation of the success of the above rate schedule: a rigorous proof that it is *equivalent* to the standard setting of BN + SGD + Standard Rate Tuning + Weight Decay + Momentum. This equivalence holds for other normalization layers as well, Group Normalization(Wu & He, 2018), Layer Normalization(Ba et al., 2016), Instance Norm(Ulyanov et al., 2016), etc.
- A worked-out toy example illustrating the above linkage of hyper-parameters. Using either weight decay or BN alone reaches global minimum, but convergence fails when both are used.

1 INTRODUCTION

Batch Normalization (BN) offers significant benefits in optimization and generalization across architectures, and has become ubiquitous. Usually best performance is attained by adding weight decay and momentum in addition to BN.

Usually weight decay is thought to improve generalization by controlling the norm of the parameters. However, it is fallacious to try to separately think of optimization and generalization because we are dealing with a nonconvex objective with multiple optima. Even slight changes to the training surely lead to a different trajectory in the loss landscape, potentially ending up at a different solution! One needs trajectory analysis to have a hope of reasoning about the effects of such changes.

In the presence of BN and other normalization schemes, including GroupNorm, LayerNorm, and InstanceNorm, the optimization objective is *scale invariant* to the parameters, which means rescaling parameters would not change the prediction, except the parameters that compute the output which do not have BN. However, Hoffer et al. (2018b) shows that fixing the output layer randomly doesn't harm the performance of the network. So the trainable parameters satisfy scale invariance.(See more in Appendix E) The current paper introduces new modes of analysis for such settings. This rigorous analysis yields the surprising conclusion that the original learning rate (LR) schedule and weight decay(WD) can be folded into a new exponential schedule for learning rate: in each iteration multiplying it by $(1+\alpha)$ for some $\alpha > 0$ that depends upon the momentum and weight decay rate.

Theorem 1.1 (Main, Informal)**.** SGD on a scale-invariant objective with initial learning rate η, weight decay factor λ, and momentum factor γ is equivalent to SGD with momentum factor γ where at iteration t, the learning rate $\tilde{\eta}_t$ in the new exponential learning rate schedule is defined as $\tilde{\eta}_t = \alpha^{-2t-1}\eta$ without weight decay($\tilde{\lambda} = 0$) where α is a non-zero root of equation

$$x^2 - (1 + \gamma - \lambda\eta)x + \gamma = 0, \tag{1}$$

Specifically, when momentum $\gamma = 0$, the above schedule can be simplified as $\tilde{\eta}_t = (1 - \lambda\eta)^{-2t-1}\eta$.

The above theorem requires that the product of learning rate and weight decay factor, $\lambda\eta$, is small than $(1 - \sqrt{\gamma})^2$, which is almost always satisfied in practice. The rigorous and most general version of above theorem is Theorem 2.12, which deals with multi-phase LR schedule, momentum and weight decay.

There are other recently discovered exotic LR schedules, e.g. Triangular LR schedule(Smith, 2017) and Cosine LR schedule(Loshchilov & Hutter, 2016), and our exponential LR schedule is an extreme example of LR schedules that become possible in presence of BN. Such an exponential increase in learning rate seems absurd at first sight and to the best of our knowledge, no deep learning success has been reported using such an idea before. It does highlight the above-mentioned viewpoint that in deep learning, optimization and regularization are not easily separated. Of course, the exponent trumps the effect of initial lr very fast (See Figure 4), which explains why training with BN and WD is not sensitive to the scale of initialization, since with BN, tuning the scale of initialization is equivalent to tuning the initial LR η while fixing the product of LR and WD, $\eta\lambda$ (See Lemma 2.7).

Note that it is customary in BN to switch to a lower LR upon reaching a plateau in the validation loss. According to the analysis in the above theorem, this corresponds to an exponential growth with a smaller exponent, except for a transient effect when a correction term is needed for the two processes to be equivalent (see discussion around Theorem 2.12).

Thus the final training algorithm is roughly as follows: *Start from a convenient LR like* 0.1*, and grow it at an exponential rate with a suitable exponent. When validation loss plateaus, switch to an exponential growth of LR with a lower exponent. Repeat the procedure until the training loss saturates.*

In Section 3, we demonstrate on a toy example how weight decay and normalization are inseparably involved in the optimization process. With either weight decay or normalization alone, SGD will achieve zero training error. But with both turned on, SGD fails to converge to global minimum.

In Section 4, we experimentally verify our theoretical findings on CNNs and ResNets. We also construct better exponential LR schedules by incorporating the Cosine LR schedule on CIFAR10, which opens the possibility of even more general theory of rate schedule tuning towards better performance.

1.1 Related Work

There have been other theoretical analyses of training models with scale-invariance. (Cho & Lee, 2017) proposed to run Riemanian gradient descent on Grassmann manifold $\mathcal{G}(1, n)$ since the weight matrix is scaling invariant to the loss function. observed that the effective stepsize is proportional to $\frac{\eta_w}{\|\boldsymbol{w}_t\|^2}$. (Arora et al., 2019) show the gradient is always perpendicular to the current parameter vector which has the effect that norm of each scale invariant parameter group increases monotonically, which has an auto-tuning effect. (Wu et al., 2018) proposes a new adaptive learning rate schedule motivated by scale-invariance property of Weight Normalization.

Previous work for understanding Batch Normalization. (Santurkar et al., 2018) suggested that the success of BNhas does not derive from reduction in Internal Covariate Shift, but by making landscape smoother. (Kohler et al., 2018) essentially shows linear model with BN could achieve exponential convergence rate assuming gaussian inputs, but their analysis is for a variant of GD with an inner optimization loop rather than GD itself. (Bjorck et al., 2018) observe that the higher learning rates enabled by BN empirically improves generalization. (Arora et al., 2019) prove that with certain mild assumption, (S)GD with BN finds approximate first order stationary point with any fixed learning rate. None of the above analyses incorporated weight decay, but (Zhang et al., 2019; Hoffer et al., 2018a; van Laarhoven, 2017) argued qualitatively that weight decay makes parameters

have smaller norms, and thus the effective learning rate, $\frac{\eta_w}{\|\boldsymbol{w}_t\|^2}$ is larger. They described experiments showing this effect but didn't have a closed form theoretical analysis like ours. None of the above analyses deals with momentum rigorously.

1.2 Preliminaries and Notations

For batch $\mathcal{B} = \{x_i\}_{i=1}^B$, network parameter $\boldsymbol{\theta}$, we denote the network by $f_{\boldsymbol{\theta}}$ and the loss function at iteration t by $L_t(f_{\boldsymbol{\theta}}) = L(f_{\boldsymbol{\theta}}, \mathcal{B}_t)$. When there's no ambiguity, we also use $L_t(\boldsymbol{\theta})$ for convenience.

We say a loss function $L(\boldsymbol{\theta})$ is *scale invariant* to its parameter $\boldsymbol{\theta}$ is for any $c \in \mathbb{R}^+$, $L(\boldsymbol{\theta}) = L(c\boldsymbol{\theta})$. In practice, the source of scale invariance is usually different types of normalization layers, including Batch Normalization (Ioffe & Szegedy, 2015), Group Normalization (Wu & He, 2018), Layer Normalization (Ba et al., 2016), Instance Norm (Ulyanov et al., 2016), etc.

Implementations of SGD with Momentum/Nesterov comes with subtle variations in literature. We adopt the variant from Sutskever et al. (2013), also the default in PyTorch (Paszke et al., 2017). $L2$ *regularization* (a.k.a. *Weight Decay*) is another common trick used in deep learning. Combining them together, we get the one of the mostly used optimization algorithms below.

Definition 1.2. [SGD with Momentum and Weight Decay] At iteration t, with randomly sampled batch $\mathcal{B}_t$, update the parameters $\boldsymbol{\theta}_t$ and momentum $\boldsymbol{v}_t$ as following:

$$\boldsymbol{\theta}_t = \boldsymbol{\theta}_{t-1} - \eta_{t-1}\boldsymbol{v}_t \tag{2}$$

$$\boldsymbol{v}_t = \gamma \boldsymbol{v}_{t-1} + \nabla_{\boldsymbol{\theta}} \left(L_t(\boldsymbol{\theta}_{t-1}) + \frac{\lambda_{t-1}}{2}\|\boldsymbol{\theta}_{t-1}\|^2 \right), \tag{3}$$

where η_t is the learning rate at epoch t, γ is the momentum coefficient, and λ is the factor of weight decay. Usually, $\boldsymbol{v}_0$ is initialized to be $\mathbf{0}$.

For ease of analysis, we will use the following equivalent of Definition 1.2.

$$\frac{\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-1}}{\eta_{t-1}} = \gamma \frac{\boldsymbol{\theta}_{t-1} - \boldsymbol{\theta}_{t-2}}{\eta_{t-2}} - \nabla_{\boldsymbol{\theta}} \left((L(\boldsymbol{\theta}_{t-1}) + \frac{\lambda_{t-1}}{2}\|\boldsymbol{\theta}_{t-1}\|_2^2 \right), \tag{4}$$

where η_{-1} and $\boldsymbol{\theta}_{-1}$ must be chosen in a way such that $\boldsymbol{v}_0 = \frac{\boldsymbol{\theta}_0 - \boldsymbol{\theta}_{-1}}{\eta_{-1}}$ is satisfied, e.g. when $\boldsymbol{v}_0 = \mathbf{0}$, $\boldsymbol{\theta}_{-1} = \boldsymbol{\theta}_0$ and η_{-1} could be arbitrary.

A key source of intuition is the following simple lemma about scale-invariant networks Arora et al. (2019). The first property ensures GD (with momentum) always increases the norm of the weight.(See Lemma C.1 in Appendix C) and the second property says that the gradients are smaller for parameteres with larger norm, thus stabilizing the trajectory from diverging to infinity.

Lemma 1.3 (Scale Invariance)**.** *If for any $c \in \mathbb{R}^+$, $L(\boldsymbol{\theta}) = L(c\boldsymbol{\theta})$, then*
(1). $\langle \nabla_{\boldsymbol{\theta}} L, \boldsymbol{\theta} \rangle = 0$*;*
(2). $\nabla_{\boldsymbol{\theta}} L\big|_{\boldsymbol{\theta}=\boldsymbol{\theta}_0} = c\nabla_{\boldsymbol{\theta}} L\big|_{\boldsymbol{\theta}=c\boldsymbol{\theta}_0}$*, for any* $c > 0$

2 Deriving Exponential Learning Rate Schedule

As a warm-up in Section 2.1 we show that if momentum is turned off then *Fixed LR + Fixed WD* can be translated to an equivalent *Exponential LR*. In Section 2.2 we give a more general analysis on the equivalence between *Fixed LR + Fixed WD + Fixed Momentum Factor* and *Exponential LR + Fixed Momentum Factor*. While interesting, this still does completely apply to real-life deep learning where reaching full accuracy usually requires multiple phases in training where LR is fixed within a phase and reduced by some factor from one phase to the next. Section 2.3 shows how to interpret such a multi-phase LR schedule + WD + Momentum as a certain multi-phase exponential LR schedule with Momentum.

2.1 Replacing WD by Exponential LR in Momentum-Free SGD

We use notation of Section 1.2 and assume LR is fixed over iterations, i.e. $\eta_t = \eta_0$, and γ (momentum factor) is set as 0. We also use λ to denote WD factor and $\boldsymbol{\theta}_0$ to denote the initial parameters.

The intuition should be clear from Lemma 1.3, which says that shrinking parameter weights by factor ρ (where $\rho < 1$) amounts to making the gradient ρ^{-1} times larger without changing its direction. Thus in order to restore the ratio between original parameter and its update (LR$\times$Gradient), the easiest way would be scaling LR by ρ^2. This suggests that scaling the parameter $\boldsymbol{\theta}$ by ρ at each step is equivalent to scaling the LR η by ρ^{-2}.

To prove this formally we use the following formalism. We'll refer to the vector $(\boldsymbol{\theta}, \eta)$ the *state* of a training algorithm and study how this evolves under various combinations of parameter changes. We will think of each step in training as a *mapping* from one state to another. Since mappings can be composed, any finite number of steps also correspond to a mapping. The following are some basic mappings used in the proof.

1. Run GD with WD for a step: $\quad \mathrm{GD}_t^\rho(\boldsymbol{\theta}, \eta) = (\rho\boldsymbol{\theta} - \eta\nabla L_t(\boldsymbol{\theta}), \eta)$;
2. Scale the parameter $\boldsymbol{\theta}$: $\quad \Pi_1^c(\boldsymbol{\theta}, \eta) = (c\boldsymbol{\theta}, \eta)$;
3. Scale the LR η: $\quad \Pi_2^c(\boldsymbol{\theta}, \eta) = (\boldsymbol{\theta}, c\eta)$.

For example, when $\rho = 1$, GD_t^1 is vanilla GD update without WD, also abbreviated as GD_t. When $\rho = 1 - \lambda\eta_0$, $\mathrm{GD}_t^{1-\lambda\eta_0}$ is GD update with WD λ and LR η_0. Here L_t is the loss function at iteration t, which is decided by the batch of the training samples $\mathcal{B}_t$ in tth iteration. Below is the main result of this subsection, showing our claim that GD + WD $\Leftrightarrow$ GD+ Exp LR (when Momentum is zero). It will be proved after a series of lemmas.

Theorem 2.1 (WD $\Leftrightarrow$ Exp LR)**.** For every $\rho < 1$ and positive integer t following holds:

$$\mathrm{GD}_{t-1}^\rho \circ \cdots \circ \mathrm{GD}_0^\rho \;=\; \left[\Pi_1^{\rho^t} \circ \Pi_2^{\rho^{2t}}\right] \circ \Pi_2^{\rho^{-1}} \circ \mathrm{GD}_{t-1} \circ \Pi_2^{\rho^{-2}} \circ \cdots \circ \mathrm{GD}_1 \circ \Pi_2^{\rho^{-2}} \circ \mathrm{GD}_0 \circ \Pi_2^{\rho^{-1}}.$$

With WD being λ, ρ is set as $1 - \lambda\eta_0$ and thus the scaling factor of LR per iteration is $\rho^{-2} = (1 - \lambda\eta_0)^{-2}$, except for the first iteration it's $\rho^{-1} = (1 - \lambda\eta_0)^{-1}$.

We first show how to write GD update with WD as a composition of above defined basic maps.

Lemma 2.2. *$GD_t^\rho = \Pi_2^\rho \circ \Pi_1^\rho \circ GD_t \circ \Pi_2^{\rho^{-1}}$.*

Below we will define the proper notion of equivalence such that (1). $\Pi_1^\rho \sim \Pi_2^{\rho^{-2}}$, which implies $\mathrm{GD}_t^\rho \sim \Pi_2^{\rho^{-1}} \circ \mathrm{GD}_t \circ \Pi_2^{\rho^{-1}}$; (2) the equivalence is preserved under future GD updates.

We first extend the equivalence between weights (same direction) to that between states, with additional requirement that the ratio between the size of GD update and that of parameter are the same among all equivalent states, which yields the notion of *Equivalent Scaling*.

Definition 2.3 (Equivalent States)**.** $(\boldsymbol{\theta}, \eta)$ is *equivalent* to $(\boldsymbol{\theta}', \eta')$ iff $\exists c > 0$, $(\widetilde{\boldsymbol{\theta}}, \widetilde{\eta}) = [\Pi_1^c \circ \Pi_2^{c^2}](\boldsymbol{\theta}, \eta) = (c\boldsymbol{\theta}, c^2\eta)$, which is also denoted by $(\widetilde{\boldsymbol{\theta}}, \widetilde{\eta}) \overset{c}{\sim} (\boldsymbol{\theta}, \eta)$. $\Pi_1^c \circ \Pi_2^{c^2}$ is called *Equivalent Scaling* for all $c > 0$.

The following lemma shows that equivalent scaling commutes with GD update with WD, implying that equivalence is preserved under GD update (Lemma 2.4). This anchors the notion of equivalence — we could insert equivalent scaling anywhere in a sequence of basic maps(GD update, LR/parameter scaling), without changing the final network.

Lemma 2.4. *For any constant $c, \rho > 0$ and $t \geq 0$, $GD_t^\rho \circ [\Pi_1^c \circ \Pi_2^{c^2}] = [\Pi_1^c \circ \Pi_2^{c^2}] \circ GD_t^\rho$. In other words, $(\boldsymbol{\theta}, \eta) \overset{c}{\sim} (\boldsymbol{\theta}', \eta') \Longrightarrow GD_t^\rho(\boldsymbol{\theta}, \eta) \overset{c}{\sim} GD_t^\rho(\boldsymbol{\theta}', \eta')$.*

Now we formally define equivalence relationship between maps using equivalent scalings.

Definition 2.5 (Equivalent Maps)**.** Two maps F, G are *equivalent* iff $\exists c > 0$, $F = \Pi_1^c \circ \Pi_2^{c^2} \circ G$, which is also denoted by $F \overset{c}{\sim} G$.

Proof of Theorem 2.1. By Lemma 2.2,, $\mathrm{GD}_t^\rho \overset{\rho}{\sim} \Pi_2^{\rho^{-1}} \circ \mathrm{GD}_t \circ \Pi_2^{\rho^{-1}}$. By Lemma 2.4, GD update preserves map equivalence, i.e. $F \overset{c}{\sim} G \Rightarrow \mathrm{GD}_t^\rho \circ F \overset{c}{\sim} \mathrm{GD}_t^\rho \circ G, \forall c, \rho > 0$. Thus,

$$\mathrm{GD}_{t-1}^\rho \circ \cdots \circ \mathrm{GD}_0^\rho \overset{\rho^t}{\sim} \Pi_2^{\rho^{-1}} \circ \mathrm{GD}_{t-1} \circ \Pi_2^{\rho^{-2}} \circ \cdots \circ \mathrm{GD}_1 \circ \Pi_2^{\rho^{-2}} \circ \mathrm{GD}_0 \circ \Pi_2^{\rho^{-1}}. \qquad \square$$

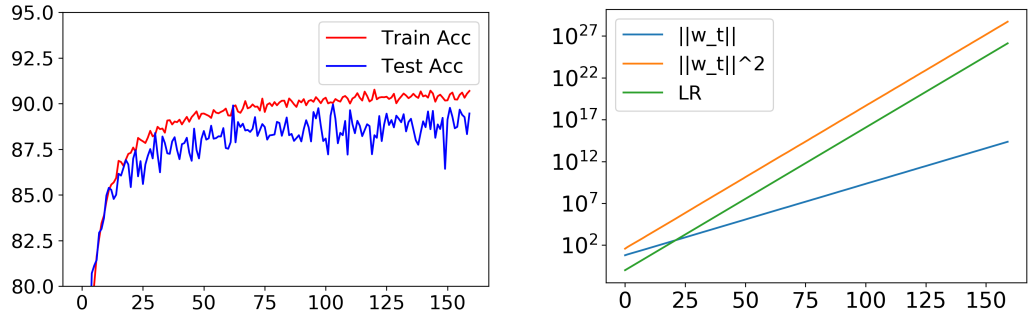

Figure 1: Taking PreResNet32 with standard hyperparameters and replacing WD during first phase (Fixed LR) by exponential LR according to Theorem 2.9 to the schedule $\widetilde{\eta}_t = 0.1 \times 1.481^t$, momentum 0.9. Plot on right shows weight norm $\boldsymbol{w}$ of the first convolutional layer in the second residual block grows exponentially, satisfying $\frac{\|\boldsymbol{w}_t\|^2}{\widetilde{\eta}_t}$ = constant. Reason being that according to the proof it is essentially the norm square of the weights when trained with Fixed LR + WD + Momentum, and published hyperparameters kept this norm roughly constant during training.

2.2 Replacing WD by Exponential LR: Case of constant LR with momentum

In this subsection the setting is the same to that in Subsection 2.1 except that the momentum factor is γ instead of 0. Suppose the initial momentum is $\boldsymbol{v}_0$, we set $\boldsymbol{\theta}_{-1} = \boldsymbol{\theta}_0 - \boldsymbol{v}_0\eta$. Presence of momentum requires representing the state of the algorithm with four coordinates, $(\boldsymbol{\theta}, \eta, \boldsymbol{\theta}', \eta')$, which stand respectively for the current parameters/LR and the buffered parameters/LR (from last iteration) respectively. Similarly, we define the following basic maps and equivalence relationships.

1. Run GD with WD for a step: $\mathrm{GD}_t^\rho(\boldsymbol{\theta}, \eta, \boldsymbol{\theta}', \eta') = \left(\rho\boldsymbol{\theta} + \eta\left(\gamma\frac{\boldsymbol{\theta}-\boldsymbol{\theta}'}{\eta'} - \nabla L_t(\boldsymbol{\theta})\right), \eta, \boldsymbol{\theta}, \eta\right)$;
2. Scale Current parameter $\boldsymbol{\theta}$ $\Pi_1^c(\boldsymbol{\theta}, \eta, \boldsymbol{\theta}', \eta') = (c\boldsymbol{\theta}, \eta, \boldsymbol{\theta}', \eta')$;
3. Scale Current LR η: $\Pi_2^c(\boldsymbol{\theta}, \eta, \boldsymbol{\theta}', \eta') = (\boldsymbol{\theta}, c\eta, \boldsymbol{\theta}', \eta')$;
4. Scale Buffered parameter $\boldsymbol{\theta}'$: $\Pi_3^c(\boldsymbol{\theta}, \eta, \boldsymbol{\theta}', \eta') = (\boldsymbol{\theta}, \eta, c\boldsymbol{\theta}', \eta')$;
5. Scale Buffered parameter η': $\Pi_4^c(\boldsymbol{\theta}, \eta, \boldsymbol{\theta}', \eta') = (\boldsymbol{\theta}, \eta, \boldsymbol{\theta}', c\eta')$.

Definition 2.6 (Equivalent States)**.** $(\boldsymbol{\theta}, \eta, \boldsymbol{\theta}', \eta')$ is *equivalent* to $(\widetilde{\boldsymbol{\theta}}, \widetilde{\eta}, \widetilde{\boldsymbol{\theta}}', \widetilde{\eta}')$ iff $\exists c > 0$, $(\boldsymbol{\theta}, \eta, \boldsymbol{\theta}', \eta') = \left[\Pi_1^c \circ \Pi_2^{c^2} \circ \Pi_3^c \circ \Pi_4^{c^2}\right](\widetilde{\boldsymbol{\theta}}, \widetilde{\eta}, \widetilde{\boldsymbol{\theta}}', \widetilde{\eta}') = (c\widetilde{\boldsymbol{\theta}}, c^2\widetilde{\eta}, c\widetilde{\boldsymbol{\theta}}', c^2\widetilde{\eta}')$, which is also denoted by $(\boldsymbol{\theta}, \eta, \boldsymbol{\theta}', \eta') \overset{c}{\sim} (\widetilde{\boldsymbol{\theta}}, \widetilde{\eta}, \widetilde{\boldsymbol{\theta}}', \widetilde{\eta}')$. We call $\Pi_1^c \circ \Pi_2^{c^2} \circ \Pi_3^c \circ \Pi_4^{c^2}$ *Equivalent Scalings* for all $c > 0$.

Again by expanding the definition, we show equivalent scalings commute with GD update.

Lemma 2.7. *$\forall c, \rho > 0$ and $t \geq 0$, $GD_t^\rho \circ \left[\Pi_1^c \circ \Pi_2^{c^2} \circ \Pi_3^c \circ \Pi_4^{c^2}\right] = \left[\Pi_1^c \circ \Pi_2^{c^2} \circ \Pi_3^c \circ \Pi_4^{c^2}\right] \circ GD_t^\rho$.*

Similarly, we can rewrite GD_t^ρ as a composition of vanilla GD update and other scalings by expanding the definition, when the current and buffered LR are the same in the input of GD_t^ρ.

Lemma 2.8. *For any input $(\boldsymbol{\theta}, \eta, \boldsymbol{\theta}', \eta)$, if $\alpha > 0$ is a root of $\alpha + \gamma\alpha^{-1} = \rho + \gamma$, then $GD_t^\rho(\boldsymbol{\theta}, \eta, \boldsymbol{\theta}', \eta) = \left[\Pi_4^\alpha \circ \Pi_2^\alpha \circ \Pi_1^\alpha \circ GD_t \circ \Pi_2^{\alpha^{-1}} \circ \Pi_3^\alpha \circ \Pi_4^\alpha\right](\boldsymbol{\theta}, \eta, \boldsymbol{\theta}', \eta)$. In other words,*

$$GD_t^\rho(\boldsymbol{\theta}, \eta, \boldsymbol{\theta}', \eta) \overset{\alpha}{\sim} \left[\Pi_3^{\alpha^{-1}} \circ \Pi_4^{\alpha^{-1}} \circ \Pi_2^{\alpha^{-1}} \circ GD_t \circ \Pi_2^{\alpha^{-1}} \circ \Pi_3^\alpha \circ \Pi_4^\alpha\right](\boldsymbol{\theta}, \eta, \boldsymbol{\theta}', \eta). \tag{5}$$

Though looking complicated, the RHS of Equation 5 is actually the desired $\Pi_2^{\alpha^{-1}} \circ \mathrm{GD}_t \circ \Pi_2^{\alpha^{-1}}$ conjugated with some scaling on momentum part $\Pi_3^\alpha \circ \Pi_4^\alpha$, and $\Pi_3^{\alpha^{-1}} \circ \Pi_4^{\alpha^{-1}}$ in the current update cancels with the $\Pi_3^\alpha \circ \Pi_4^\alpha$ in the next update. Now we are ready to show the equivalence between WD and Exp LR schedule when momentum is turned on for both.

Theorem 2.9 (GD + WD $\Leftrightarrow$ GD+ Exp LR; With Momentum)**.** The following defined two sequences of parameters ,$\{\boldsymbol{\theta}_t\}_{t=0}^\infty$ and $\{\widetilde{\boldsymbol{\theta}}_t\}_{t=0}^\infty$, satisfy $\widetilde{\boldsymbol{\theta}}_t = \alpha^t\boldsymbol{\theta}_t$, thus they correspond to the same networks in function space, i.e. $f_{\boldsymbol{\theta}_t} = f_{\widetilde{\boldsymbol{\theta}}_t}$, $\forall t \in \mathbb{N}$, given $\widetilde{\boldsymbol{\theta}}_0 = \boldsymbol{\theta}_0$, $\widetilde{\boldsymbol{\theta}}_{-1} = \boldsymbol{\theta}_{-1}\alpha$, and $\widetilde{\eta}_t = \eta_0\alpha^{-2t-1}$.

1. $\frac{\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-1}}{\eta_0} = \frac{\gamma(\boldsymbol{\theta}_{t-1} - \boldsymbol{\theta}_{t-2})}{\eta_0} - \nabla_{\boldsymbol{\theta}}(L(\boldsymbol{\theta}_{t-1}) + \frac{\lambda}{2}\|\boldsymbol{\theta}_{t-1}\|_2^2)$
2. $\frac{\widetilde{\boldsymbol{\theta}}_t - \widetilde{\boldsymbol{\theta}}_{t-1}}{\widetilde{\eta}_t} = \frac{\gamma(\widetilde{\boldsymbol{\theta}}_{t-1} - \widetilde{\boldsymbol{\theta}}_{t-2})}{\widetilde{\eta}_{t-1}} - \nabla_{\boldsymbol{\theta}} L(\widetilde{\boldsymbol{\theta}}_{t-1})$

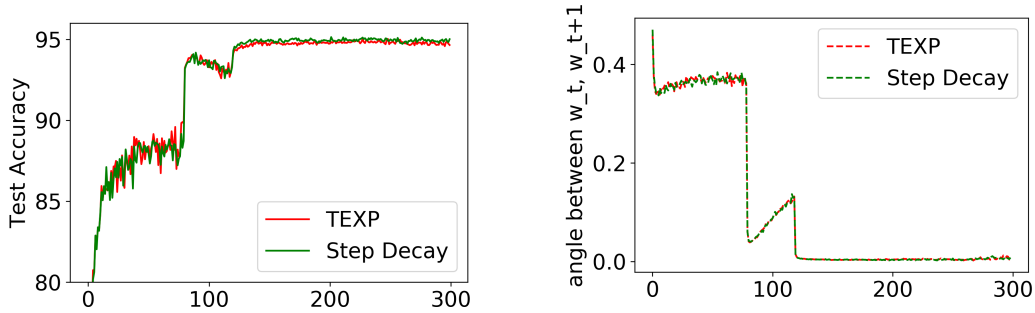

Figure 2: PreResNet32 trained with standard Step Decay and its corresponding Tapered-Exponential LR schedule. As predicted by Theorem 2.12, they have similar trajectories and performances.

where α is a positive root of equation $x^2 - (1+\gamma-\lambda\eta_0)x + \gamma = 0$, which is always smaller than 1(See Appendix B.1). When $\gamma = 0$, $\alpha = 1 - \lambda\eta_0$ is the unique non-zero solution.

Remark 2.10. Above we implicitly assume that $\lambda\eta_0 \leq (1-\sqrt{\gamma})^2$ such that the roots are real and this is always true in practice. For instance of standard hyper-parameters where $\gamma = 0.9, \eta_0 = 0.1, \lambda = 0.0005$, $\frac{\lambda\eta_0}{(1-\sqrt{\gamma})^2} \approx 0.019 \ll 1$.

Proof. Note that $(\widetilde{\boldsymbol{\theta}}_0, \widetilde{\eta}_0, \widetilde{\boldsymbol{\theta}}_{-1}, \widetilde{\eta}_{-1}) = \left[\Pi_2^{\alpha^{-1}} \circ \Pi_3^{\alpha} \circ \Pi_4^{\alpha}\right](\boldsymbol{\theta}_0, \eta_0, \boldsymbol{\theta}_0, \eta_0)$, it suffices to show that

$$\left[\Pi_3^{\alpha^{-1}} \circ \Pi_4^{\alpha^{-1}} \circ \Pi_2^{\alpha^{-1}} \circ \mathrm{GD}_{t-1} \circ \Pi_2^{\alpha^{-2}} \circ \cdots \circ \mathrm{GD}_1 \circ \Pi_2^{\alpha^{-2}} \circ \mathrm{GD}_0 \circ \Pi_2^{\alpha^{-1}} \circ \Pi_3^{\alpha} \circ \Pi_4^{\alpha}\right](\boldsymbol{\theta}_0, \eta_0, \boldsymbol{\theta}_0, \eta_0)$$
$$\overset{\alpha^t}{\sim} \mathrm{GD}_{t-1}^{1-\lambda\eta_0} \circ \cdots \circ \mathrm{GD}_0^{1-\lambda\eta_0}(\boldsymbol{\theta}_0, \eta_0, \boldsymbol{\theta}_0, \eta_0), \quad \forall t \geq 0.$$

which follows immediately from Lemma 2.7 and Lemma 2.8 by induction. □

2.3 Replacing WD by Exponential LR: Case of multiple LR phases

Usual practice in deep learning shows that reaching full training accuracy requires reducing the learning rate a few times.

Definition 2.11. *Step Decay* is the (standard) learning rate schedule, where training has K phases $I = 0, 1, \ldots, K-1$, where phase I starts at iteration T_I ($T_0 = 0$), and all iterations in phase I use a fixed learning rate of η_I^*.

The algorithm state in Section 2.2, consists of 4 components including buffered and current LR. When LR changes, the buffered and current LR are not equal, and thus Lemma 2.8 cannot be applied any more. In this section we show how to fix this issue by adding extra momentum correction. In detail, we show the below defined Exp LR schedule leads the same trajectory of networks in function space, with one-time momentum correction at the start of each phase. We empirically find on CIFAR10 that ignoring the correction term does not change performance much.

Theorem 2.12 (Tapered-Exponential LR Schedule)**.** There exists a way to correct the momentum only at the first iteration of each phase, such that the following Tapered-Exponential LR schedule (TEXP) $\{\widetilde{\eta}_t\}$ with momentum factor γ and no WD, leads the same sequence networks in function space as that of Step Decay LR schedule(Definition 2.11) with momentum factor γ and WD λ.

$$\widetilde{\eta}_t = \begin{cases} \widetilde{\eta}_{t-1} \times (\alpha_{I-1}^*)^{-2} & \text{if } T_{I-1}+1 \leq t \leq T_I - 1, I \geq 1; \\ \widetilde{\eta}_{t-1} \times \frac{\eta_I^*}{\eta_{I-1}^*} \times (\alpha_I^*)^{-1}(\alpha_{I-1}^*)^{-1} & \text{if } t = T_I, I \geq 1, \end{cases} \tag{6}$$

where $\alpha_I^* = \frac{1+\gamma-\lambda\eta_I^*+\sqrt{(1+\gamma-\lambda\eta_I^*)^2-4\gamma}}{2}$, $\widetilde{\eta}_0 = \eta_0 \cdot (\alpha_0^*)^{-1} = \eta_0^* \cdot (\alpha_0^*)^{-1}$.

The analysis in previous subsection give the equivalence within each phase, where the same LR is used throughout the phase. To deal with the difference between buffered LR and current LR when entering new phases, the idea is to pretend $\eta_{t-1} = \eta_t$ and $\boldsymbol{\theta}_{t-1}$ becomes whatever it needs to maintain $\frac{\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-1}}{\eta_{t-1}}$ such that we can again apply Lemma 2.8, which requires the current LR of the input state is equal to its buffered LR. Because scaling α in RHS of Equation 5 is different in different phases, so unlike what happens within each phase, they don't cancel with each other at phase transitions, thus remaining as a correction of the momentum. The proofs are delayed to Appendix B, where we proves a more general statement allowing phase-dependent WD, $\{\lambda_I\}_{I=0}^{K-1}$.

Alternative interpretation of Step Decay to exponential LR schedule: Below we present a new LR schedule, *TEXP++*, which is exactly equivalent to *Step Decay* without the need of one-time correction of momentum when entering each phase. We further show in Appendix B.1 that when translating from *Step Decay*, the *TEXP++* we get is very close to the original *TEXP*(Equation 9), i.e. the ratio between the LR growth per round, $\frac{\widetilde{\eta}_{t+1}}{\widetilde{\eta}_t} / \frac{\widetilde{\eta}'_{t+1}}{\widetilde{\eta}'_t}$ converges to 1 exponentially each phase. For example, with WD 0.0005, max LR 0.1, momentum factor 0.9, the ratio is within $1 \pm 0.0015 * 0.9^{t-T_I}$, meaning TEXP and TEXP++ are very close for Step Decay with standard hyperparameters.

Theorem 2.13. The following two sequences of parameters ,$\{\boldsymbol{\theta}_t\}_{t=0}^{\infty}$ and $\{\widetilde{\boldsymbol{\theta}}_t\}_{t=0}^{\infty}$, define the same sequence of network functions, i.e. $f_{\boldsymbol{\theta}_t} = f_{\widetilde{\boldsymbol{\theta}}_t}$, $\forall t \in \mathbb{N}$, given the initial conditions, $\widetilde{\boldsymbol{\theta}}_0 = P_0 \boldsymbol{\theta}_0$, $\widetilde{\boldsymbol{\theta}}_{-1} = P_{-1} \boldsymbol{\theta}_{-1}$.

1. $\frac{\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-1}}{\eta_{t-1}} = \gamma \frac{\boldsymbol{\theta}_{t-1} - \boldsymbol{\theta}_{t-2}}{\eta_{t-2}} - \nabla_{\boldsymbol{\theta}} \left((L(\boldsymbol{\theta}_{t-1}) + \frac{\lambda_{t-1}}{2} \|\boldsymbol{\theta}_{t-1}\|_2^2 \right)$, for $t = 1, 2, \ldots$;
2. $\frac{\widetilde{\boldsymbol{\theta}}_t - \widetilde{\boldsymbol{\theta}}_{t-1}}{\widetilde{\eta}_{t-1}} = \gamma \frac{\widetilde{\boldsymbol{\theta}}_{t-1} - \widetilde{\boldsymbol{\theta}}_{t-2}}{\widetilde{\eta}_{t-2}} - \nabla_{\boldsymbol{\theta}} L(\widetilde{\boldsymbol{\theta}}_{t-1})$, for $t = 1, 2, \ldots$,

where $\widetilde{\eta}_t = P_t P_{t+1} \eta_t$, $P_t = \prod_{i=-1}^{t} \alpha_i^{-1}$, $\forall t \geq -1$ and α_t recursively defined as

$$\alpha_t = -\eta_{t-1} \lambda_{t-1} + 1 + \frac{\eta_{t-1}}{\eta_{t-2}} \gamma (1 - \alpha_{t-1}^{-1}), \forall t \geq 1. \tag{7}$$

The LR schedule $\{\widetilde{\eta}_t\}_{t=0}^{\infty}$ is called *Tapered Exponential* ++, or *TEXP*++.

3 Example illustrating interplay of WD and BN

The paper so far has shown that effects of different hyperparameters in training are not easily separated, since their combined effect on the trajectory is complicated. We give a simple example to illustrate this, where convergence is guaranteed if we use either BatchNorm or weight decay in isolation, but convergence fails if both are used. (Momentum is turned off for clarity of presentation)

Setting: Suppose we are fine-tuning the last linear layer of the network, where the input of the last layer is assumed to follow a standard Gaussian distribution $\mathcal{N}(0, I_m)$, and m is the input dimension of last layer. We also assume this is a binary classification task with logistic loss, $l(u, y) = \ln(1 + \exp(-uy))$, where label $y \in \{-1, 1\}$ and $u \in \mathbb{R}$ is the output of the neural network. The training algorithm is SGD with constant LR and WD, and without momentum. For simplicity we assume the batch size B is very large so we could assume the covariance of each batch $\mathcal{B}_t$ concentrates and is approximately equal to identity, namely $\frac{1}{B} \sum_{i=1}^{B} \boldsymbol{x}_{t,b} \boldsymbol{x}_{t,b}^{\top} \approx I_m$. We also assume the the input of the last layer are already separable, and w.l.o.g. we assume the label is equal to the sign of the first coordinate of $\boldsymbol{x} \in \mathbb{R}^m$, namely sign (x_1) . Thus the training loss and training error are simply

$$L(\boldsymbol{w}) = \mathop{\mathbb{E}}_{\boldsymbol{x} \sim \mathcal{N}(0, I_m), y = \text{sign}(x_1)} \left[\ln(1 + \exp(-\boldsymbol{x}^{\top} \boldsymbol{w} y)) \right], \quad \mathop{\Pr}_{\boldsymbol{x} \sim \mathcal{N}(0, I_m), y = \text{sign}(x_1)} \left[\boldsymbol{x}^{\top} \boldsymbol{w} y \leq 0 \right] = \frac{1}{\pi} \arccos \frac{w_1}{\|\boldsymbol{w}\|}$$

Case 1: WD alone: Since both the above objective with L2 regularization is strongly convex and smooth in $\boldsymbol{w}$, vanilla GD *with suitably small learning rate* could get arbitrarily close to the global minimum for this regularized objective. In our case, large batch SGD behaves similarly to GD and can achieve $O(\sqrt{\frac{\eta \lambda}{B}})$ test error following the standard analysis of convex optimization.

Case 2: BN alone: Add a BN layer after the linear layer, and fix scalar and bias term to 1 and 0. The objective becomes

$$L_{BN}(\boldsymbol{w}) = \mathop{\mathbb{E}}_{\boldsymbol{x} \sim \mathcal{N}(0, I_m), y = \text{sign}(x_1)} \left[L_{BN}(\boldsymbol{w}, \boldsymbol{x}) \right] = \mathop{\mathbb{E}}_{\boldsymbol{x} \sim \mathcal{N}(0, I_m), y = \text{sign}(x_1)} \left[\ln(1 + \exp(-\boldsymbol{x}^{\top} \frac{\boldsymbol{w}}{\|\boldsymbol{w}\|} y)) \right].$$

From Appendix B.6, there's some constant C, such that $\forall \boldsymbol{w} \in \mathbb{R}^m$ with constant probability, $\|\nabla_{\boldsymbol{w}} L_{BN}(\boldsymbol{w}, \boldsymbol{x})\| \geq \frac{C}{\|\boldsymbol{w}\|}$. By Pythagorean Theorem, $\|\boldsymbol{w}_{t+1}\|^4 = (\|\boldsymbol{w}_t\|^2 + \eta^2 \|\nabla_{\boldsymbol{w}} L_{BN}(\boldsymbol{w}_t, \boldsymbol{x})\|^2)^2 \geq \|\boldsymbol{w}_t\|^4 + 2\eta^2 \|\boldsymbol{w}_t\|^2 \|\nabla_{\boldsymbol{w}} L_{BN}(\boldsymbol{w}_t, \boldsymbol{x})\|^2$. As a result, for any *fixed learning rate*, $\|\boldsymbol{w}_{t+1}\|^4 \geq 2 \sum_{i=1}^{t} \eta^2 \|\boldsymbol{w}\|^2 \|\nabla_{\boldsymbol{w}} L_{BN}(\boldsymbol{w}_i, \boldsymbol{x})\|^2$ grows at least linearly with high probability. Following the analysis of Arora et al. (2019), this is like reducing the effective learning

rate, and when $\|\boldsymbol{w}_t\|$ is large enough, the effective learning rate is small enough, and thus SGD can find the local minimum, which is the unique global minimum.

Case 3: Both BN and WD: Suppose WD factor is λ, LR is η, we have the following theorem:

Theorem 3.1. [Nonconvergence] Starting from iteration any T_0, with probability $1-\delta$ over the randomness of samples, the training error will be larger than $\frac{\varepsilon}{\pi}$ at least once for the following consecutive $\frac{1}{2(\eta\lambda-2\varepsilon^2)}\ln\frac{64\|w_{T_0}\|^2\varepsilon\sqrt{B}}{\eta\sqrt{m-2}}+9\ln\frac{1}{\delta}$ iterations.

Sketch. (See full proof in Appendix B.) The high level idea of this proof is that if the test error is low, the weight is restricted in a small cone around the global minimum, and thus the amount of the gradient update is bounded by the size of the cone. In this case, the growth of the norm of the weight by Pythagorean Theorem is not large enough to cancel the shrinkage brought by weight decay. As a result, the norm of the weight converges to 0 geometrically. Again we need to use the lower bound for size of the gradient, that $\|\nabla_{\boldsymbol{w}}L_t\|=\Theta(\frac{\eta\sqrt{m}}{\|\boldsymbol{w}_t\|})$ holds with constant probability. Thus the size of the gradient will grow along with the shrinkage of $\|\boldsymbol{w}_t\|$ until they're comparable, forcing the weight to leave the cone in next iteration. □

4 Experiments

The translation to exponential LR schedule is exact except for one-time momentum correction term entering new phases. The experiments explore the effect of this correction term. The Tapered Exponential(TEXP) LR schedule contains two parts when entering a new phase I: an instant LR decay ($\frac{\eta_I}{\eta_{I-1}}$) and an adjustment of the growth factor ($\alpha^*_{I-1}\to\alpha^*_I$). The first part is relative small compared to the huge exponential growing. Thus a natural question arises: *Can we simplify TEXP LR schedule by dropping the part of instant LR decay?*

Also, previously we have only verified our equivalence theorem in Step Decay LR schedules. But it's not sure how would the Exponential LR schedule behave on more rapid time-varying LR schedules such as Cosine LR schedule.

Settings: We train PreResNet32 on CIFAR10. The initial learning rate is 0.1 and the momentum is 0.9 in all settings. We fix all the scalar and bias of BN, because otherwise they together with the following conv layer grow exponentially, sometimes exceeding the range of Float32 when trained with large growth rate for a long time. We fix the parameters in the last fully connected layer for scale invariance of the objective.

4.1 The benefit of instant LR decay

We tried the following LR schedule (we call it *TEXP--*). Interestingly, up to correction of momentum when entering a new phase, this schedule is equivalent to a constant LR schedule, but with the weight decay coefficient reduced correspondingly at the start of each phase. (See Theorem B.2 and Figure 6)

$$\text{TEXP--:}\quad \widetilde{\eta}_{t+1}=\begin{cases}\widetilde{\eta}_t\times(\alpha^*_{I-1})^{-2} & \text{if } T_{I-1}+1\le t\le T_I-1, I\ge 1;\\ \widetilde{\eta}_t\times(\alpha^*_I)^{-1}(\alpha^*_{I-1})^{-1} & \text{if } t=T_I, I\ge 1,\end{cases}\tag{8}$$

where $\alpha^*_I=\frac{1+\gamma-\lambda\eta^*_I+\sqrt{(1+\gamma-\lambda\eta^*_I)^2-4\gamma}}{2}$, $\widetilde{\eta}_0=\eta_0\cdot(\alpha^*_0)^{-1}=\eta^*_0\cdot(\alpha^*_0)^{-1}$.

4.2 Better Exponential LR Schedule with Cosine LR

We applied the TEXP LR schedule (Theorem 2.12) on the Cosine LR schedule (Loshchilov & Hutter, 2016), where the learning rate changes every epoch, and thus correction terms cannot be ignored. The LR at epoch $t\le T$ is defined as: $\eta_t=\eta_0\frac{1+\cos(\frac{t}{T}\pi)}{2}$. Our experiments show this hybrid schedule with Cosine LR performs better on CIFAR10 than Step Decay, but this finding needs to be verified on other datasets.

5 Conclusions

The paper shows rigorously how BN allows a host of very exotic learning rate schedules in deep learning, and verifies these effects in experiments. The lr increases exponentially in almost every

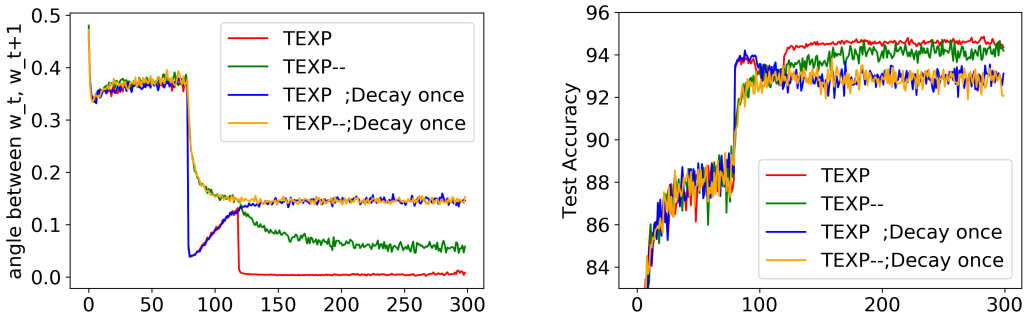

Figure 3: Instant LR decay is crucial when LR growth $\widetilde{\eta}_t/\widetilde{\eta}_{t-1}-1$ is very small. The original LR of Step Decay is decayed by 10 at epoch $80, 120$ respectively. In the third phase, LR growth $\widetilde{\eta}_t/\widetilde{\eta}_{t-1}-1$ is approximately 100 times smaller than that in the third phase, it would take TEXP-- hundreds of epochs to reach its equilibrium. As a result, TEXP achieves better test accuracy than TEXP--. As a comparison, in the second phase, $\widetilde{\eta}_t/\widetilde{\eta}_{t-1}-1$ is only 10 times smaller than that in the first phase and it only takes 70 epochs to return to equilibrium.

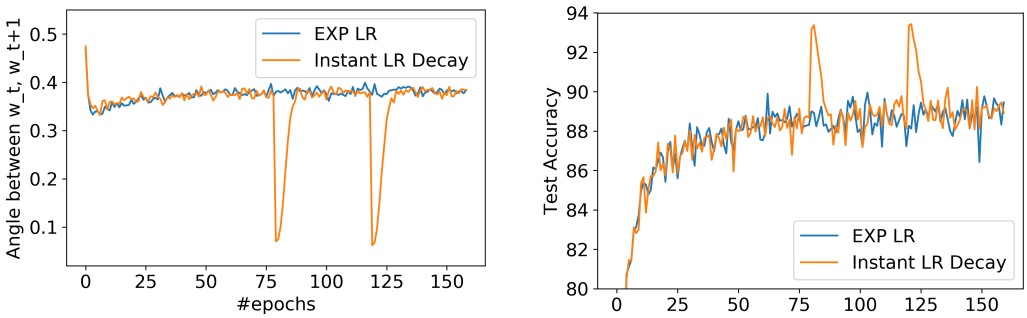

Figure 4: Instant LR decay has only temporary effect when LR growth $\widetilde{\eta}_t/\widetilde{\eta}_{t-1}-1$ is large. The blue line uses an exponential LR schedule with constant exponent. The orange line multiplies its LR by the same constant each iteration, but also divide LR by 10 at the start of epoch 80 and 120. The instant LR decay only allows the parameter to stay at good local minimum for 1 epoch and then diverges, behaving similarly to the trajectories without no instant LR decay.

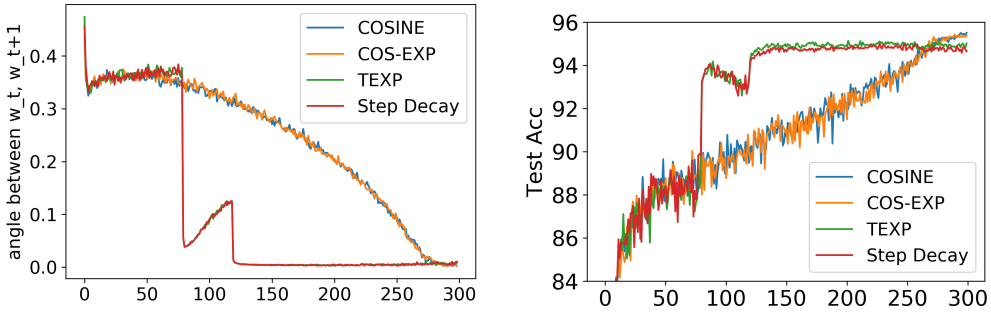

Figure 5: Both Cosine and Step Decay schedule behaves almost the same as their exponential counterpart, as predicted by our equivalence theorem. The (exponential) Cosine LR schedule achieves better test accuracy, with a entirely different trajectory.

iteration during training. The exponential increase derives from use of weight decay, but the precise expression involves momentum as well. We suggest that the efficacy of this rule may be hard to explain with canonical frameworks in optimization.

Our analyses of BN is a substantial improvement over earlier theoretical analyses, since it accounts for weight decay and momentum, which are always combined in practice.

Our tantalising experiments with a hybrid of exponential and cosine rates suggest that more surprises may lie out there. Our theoretical analysis of interrelatedness of hyperparameters could also lead to faster hyperparameter search.

REFERENCES

Sanjeev Arora, Nadav Cohen, and Elad Hazan. On the optimization of deep networks: Implicit acceleration by overparameterization. In *International Conference on Machine Learning*, pp. 244–253, 2018.

Sanjeev Arora, Zhiyuan Li, and Kaifeng Lyu. Theoretical analysis of auto rate-tuning by batch normalization. In *International Conference on Learning Representations*, 2019. URL https://openreview.net/forum?id=rkxQ-nA9FX.

Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.

Nils Bjorck, Carla P Gomes, Bart Selman, and Kilian Q Weinberger. Understanding batch normalization. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (eds.), *Advances in Neural Information Processing Systems 31*, pp. 7705–7716. Curran Associates, Inc., 2018.

Minhyung Cho and Jaehyung Lee. Riemannian approach to batch normalization. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Advances in Neural Information Processing Systems 30*, pp. 5225–5235. Curran Associates, Inc., 2017.

Sanjoy Dasgupta and Anupam Gupta. An elementary proof of a theorem of johnson and lindenstrauss. *Random Structures & Algorithms*, 22(1):60–65, 2003.

Robert Mansel Gower, Nicolas Loizou, Xun Qian, Alibek Sailanbayev, Egor Shulgin, and Peter Richtárik. Sgd: General analysis and improved rates. *arXiv preprint arXiv:1901.09401*, 2019.

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*, pp. 770–778, 2016a.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European conference on computer vision*, pp. 630–645. Springer, 2016b.

Elad Hoffer, Ron Banner, Itay Golan, and Daniel Soudry. Norm matters: efficient and accurate normalization schemes in deep networks. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (eds.), *Advances in Neural Information Processing Systems 31*, pp. 2164–2174. Curran Associates, Inc., 2018a.

Elad Hoffer, Itay Hubara, and Daniel Soudry. Fix your classifier: the marginal value of training the last weight layer. In *International Conference on Learning Representations*, 2018b. URL https://openreview.net/forum?id=S1Dh8Tg0-.

Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International Conference on Machine Learning*, pp. 448–456, 2015.

Jonas Kohler, Hadi Daneshmand, Aurelien Lucchi, Ming Zhou, Klaus Neymeyr, and Thomas Hofmann. Exponential convergence rates for batch normalization: The power of length-direction decoupling in non-convex optimization. *arXiv preprint arXiv:1805.10694*, 2018.

Ilya Loshchilov and Frank Hutter. SGDR: Stochastic Gradient Descent with Warm Restarts. *arXiv e-prints*, art. arXiv:1608.03983, Aug 2016.

David Page. How to train your resnet 6: Weight decay? URL https://myrtle.ai/how-to-train-your-resnet-6-weight-decay/.

Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017.

Shibani Santurkar, Dimitris Tsipras, Andrew Ilyas, and Aleksander Madry. How does batch normalization help optimization? In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (eds.), *Advances in Neural Information Processing Systems 31*, pp. 2488–2498. Curran Associates, Inc., 2018.

Leslie N Smith. Cyclical learning rates for training neural networks. In *2017 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pp. 464–472. IEEE, 2017.

Ilya Sutskever, James Martens, George Dahl, and Geoffrey Hinton. On the importance of initialization and momentum in deep learning. In *Proceedings of the 30th International Conference on International Conference on Machine Learning - Volume 28*, ICML'13, pp. III–1139–III–1147. JMLR.org, 2013. URL http://dl.acm.org/citation.cfm?id=3042817.3043064.

Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Instance normalization: The missing ingredient for fast stylization. *arXiv preprint arXiv:1607.08022*, 2016.

Twan van Laarhoven. L2 regularization versus batch and weight normalization. *arXiv preprint arXiv:1706.05350*, 2017.

Xiaoxia Wu, Rachel Ward, and Léon Bottou. WNGrad: Learn the Learning Rate in Gradient Descent. *arXiv preprint arXiv:1803.02865*, 2018.

Yuxin Wu and Kaiming He. Group normalization. In *The European Conference on Computer Vision (ECCV)*, September 2018.

Yang You, Igor Gitman, and Boris Ginsburg. Large Batch Training of Convolutional Networks. *arXiv e-prints*, art. arXiv:1708.03888, Aug 2017.

Guodong Zhang, Chaoqi Wang, Bowen Xu, and Roger Grosse. Three mechanisms of weight decay regularization. In *International Conference on Learning Representations*, 2019. URL https://openreview.net/forum?id=B1lz-3Rct7.