

# On the Adequacy of Untuned Warmup for Adaptive Optimization

Jerry Ma,<sup>1 2</sup> Denis Yarats<sup>3 4</sup>

<sup>1</sup> Booth School of Business, University of Chicago

<sup>2</sup> U.S. Patent and Trademark Office, Department of Commerce

<sup>3</sup> Courant Institute of Mathematical Sciences, New York University

<sup>4</sup> Facebook AI Research

research@jerry.ma.net, denisyarats@cs.nyu.edu

## Abstract

Adaptive optimization algorithms such as Adam are widely used in deep learning. The stability of such algorithms is often improved with a warmup schedule for the learning rate. Motivated by the difficulty of choosing and tuning warmup schedules, recent work proposes automatic variance rectification of Adam’s adaptive learning rate, claiming that this rectified approach (“RAdam”) surpasses the vanilla Adam algorithm and reduces the need for expensive tuning of Adam with warmup. In this work, we refute this analysis and provide an alternative explanation for the necessity of warmup based on the magnitude of the *update term*, which is of greater relevance to training stability. We then provide some “rule-of-thumb” warmup schedules, and we demonstrate that simple untuned warmup of Adam performs more-or-less identically to RAdam in typical practical settings. We conclude by suggesting that practitioners stick to linear warmup with Adam, with a sensible default being linear warmup over  $2/(1 - \beta_2)$  training iterations.

## 1 Introduction

Stochastic gradient-based optimization serves as the workhorse training approach for many classes of parametric models, including neural networks. Stochastic gradient descent and its various first-order cousins (Polyak 1964; Nesterov 1983) have enabled numerous advances in deep learning across domains (Krizhevsky, Sutskever, and Hinton 2012; He et al. 2016; Gehring et al. 2017). More recently, adaptive optimization algorithms have become prevalent in training the largest deep learning models. These adaptive methods, which include Adagrad (Duchi, Hazan, and Singer 2010), RMSProp (Hinton, Srivastava, and Swersky 2012), and Adam (Kingma and Ba 2014), scale the step size for each individual parameter based on various gradient moments.

Many practitioners have adopted the Adam algorithm for general-purpose use; notably, the preponderance of recent state-of-the-art results in natural language processing (Devlin et al. 2018; Radford et al. 2019; Liu et al. 2019b; Brown et al. 2020) have employed Adam, demonstrating the algorithm’s

ability to effectively train neural networks with parameter counts from 100 million to several billion.

In these large-scale settings, Adam’s global learning rate is usually annealed with a “warmup schedule” which promotes early-stage training stability by regulating the size of the parameter updates. The prevalent warmup schedule is a simple linear warmup, in which the global learning rate starts at zero and increases by a constant at each iteration until reaching its intended value.<sup>1</sup> The parameters of these warmup schedules are typically tuned for each problem setting and model.

Liu et al. (2020) performed an analysis of Adam with warmup, concluding that Adam requires a warmup schedule to mitigate the large or divergent variance of the per-parameter scale term. They then propose the rectified Adam (“RAdam”) algorithm, which automatically corrects for this high variance. Liu et al. highlight the robustness of RAdam, noting in particular that RAdam reduces or eliminates the need for tuning warmup schedules when using Adam. RAdam has been applied to domains including generative modeling (Yamamoto, Song, and Kim 2020), natural language processing (Nguyen and Salazar 2019), and video retrieval (Liu et al. 2019a).

## Contributions

Our contributions in this work are as follows:

**Reexamining RAdam and the variance-based motivation for warmup** We dive into the inner operation of RAdam and find that it is precisely Adam with a fixed warmup schedule, with the only deviation being to perform four iterations of heavy-ball momentum (Polyak 1964) at the outset. We then argue that the variance-based motivation for warmup is impaired as it overlooks the correlation between the first and second moment estimators, which is crucial for understanding the actual parameter updates applied by Adam.

**Analyzing Adam’s early-stage update magnitudes** Shifting focus from gradients to parameter updates, we then perform a simulation-based analysis on the magnitudes of Adam’s parameter updates. We find that even at a simulated local minimum of the objective, Adam exhibits considerable non-regularity in its early-stage parameter updates, shedding

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Appendix available at <https://arxiv.org/abs/1910.04209>.

<sup>1</sup>Linear warmup has also been deployed for first-order optimization – see, for example, Goyal et al. (2017).

light on why Adam may require learning rate warmup to a greater extent than first-order optimization methods.

**Demonstrating the sufficiency of untuned warmup** We provide some simple and intuitive “rule-of-thumb” warmup schedules for Adam, all of which require no tuning. As our main empirical result, we demonstrate that these schedules result in substantively *identical* performance and training dynamics to those of RAdam across a wide range of models, problem settings, and hyperparameters, indicating that any claimed benefits can be achieved with lower complexity using off-the-shelf optimization tools. As a sensible untuned default, we recommend linear warmup over  $2 \cdot (1 - \beta_2)^{-1}$  iterations.

## 2 Preliminaries

We begin with notation and a brief review of stochastic gradient descent and Adam.

**Primitives**  $\theta \in \mathbb{R}^p$  denotes a vector of model parameters.  $\mathcal{L}(\theta) : \mathbb{R}^p \rightarrow \mathbb{R}$  denotes a loss function to be minimized over the model parameters.  $\hat{\mathcal{L}}(\theta) : \mathbb{R}^p \rightarrow \mathbb{R}$  denotes an unbiased approximator of the loss function (*e.g.* over a minibatch).  $\nabla \mathcal{L}(\theta)$  and  $\nabla \hat{\mathcal{L}}(\theta)$  denote the gradient of  $\mathcal{L}(\theta)$  and  $\hat{\mathcal{L}}(\theta)$ , respectively. The terms  $\theta$ ,  $\hat{\mathcal{L}}(\theta)$ , and  $\nabla \hat{\mathcal{L}}(\theta)$  are subscriptable by  $t \geq 0$ , the optimization time step (“training iteration”).  $\theta_0$  represents the initial model parameters.

We write optimization algorithms as per-iteration procedures (“update rules”), taking the basic form:

$$\theta_t \leftarrow \theta_{t-1} - \underbrace{\left\{ \dots \right\}}_{\text{“update step”}}$$

**Stochastic gradient descent** The SGD algorithm, parameterized by learning rate  $\alpha > 0$ , performs the following procedure at each iteration  $t$ :

$$\theta_t \leftarrow \theta_{t-1} - \alpha \cdot \nabla \hat{\mathcal{L}}_{t-1}(\theta_{t-1}) \quad (1)$$

**Adam** The Adam algorithm (Kingma and Ba 2014), parameterized by global learning rate  $\alpha > 0$ , discount factors  $\beta_1, \beta_2 \in (0, 1)$ , and stability constant  $\epsilon > 0$ , performs the following procedure at each iteration  $t$ :

$$m_t \leftarrow \beta_1 \cdot m_{t-1} + (1 - \beta_1) \cdot \nabla \hat{\mathcal{L}}_{t-1}(\theta_{t-1}) \quad (2)$$

$$v_t \leftarrow \beta_2 \cdot v_{t-1} + (1 - \beta_2) \cdot \left[ \nabla \hat{\mathcal{L}}_{t-1}(\theta_{t-1}) \right]^2 \quad (3)$$

$$\theta_t \leftarrow \theta_{t-1} - \alpha \left[ \frac{(1 - \beta_1^t)^{-1} \cdot m_t}{\sqrt{(1 - \beta_2^t)^{-1} \cdot v_t + \epsilon}} \right] \quad (4)$$

where  $m, v \in \mathbb{R}^p$  denote auxiliary memory (interpretable as first moment and second moment estimators of  $\nabla \hat{\mathcal{L}}_t$ , respectively). By convention,  $m_0 = v_0 = 0$ .

**Warmup schedules** For any optimization algorithm parameterized with a learning rate  $\alpha$ , a *warmup schedule*  $\omega$  can be applied.  $\omega$  is a sequence of “warmup factors”  $\omega_t \in [0, 1]$ , which serve to dampen the step size of each iteration  $t$ . Specifically, a warmup schedule is imposed by replacing  $\alpha$  with  $\alpha_t = \alpha \cdot \omega_t$  in the algorithm’s update rule.

Perhaps the most common functional form for the schedule is *linear warmup*, parameterized by a “warmup period”  $\tau$ :

$$\omega_t^{\text{linear}, \tau} = \min \left\{ 1, \frac{1}{\tau} \cdot t \right\} \quad (5)$$

**Rectified Adam** The RAdam algorithm (Liu et al. 2020), parameterized identically to Adam, performs the following procedure at each iteration  $t$ :

$$\rho_\infty \leftarrow 2/(1 - \beta_2) - 1 \quad (6)$$

$$\rho_t \leftarrow \rho_\infty - 2t \cdot \beta_2^t / (1 - \beta_2^t) \quad (7)$$

$$\omega_t \leftarrow \sqrt{\frac{(\rho_t - 4)(\rho_t - 2)\rho_\infty}{(\rho_\infty - 4)(\rho_\infty - 2)\rho_t}} \quad (8)$$

$$m_t \leftarrow \beta_1 \cdot m_{t-1} + (1 - \beta_1) \cdot \nabla \hat{\mathcal{L}}_{t-1}(\theta_{t-1}) \quad (9)$$

$$v_t \leftarrow \beta_2 \cdot v_{t-1} + (1 - \beta_2) \cdot \left[ \nabla \hat{\mathcal{L}}_{t-1}(\theta_{t-1}) \right]^2 \quad (10)$$

$$\theta_t \leftarrow \theta_t - \begin{cases} \alpha \cdot (1 - \beta_1^t)^{-1} \cdot m_t & \rho_t \leq 4 \\ \alpha \cdot \omega_t \cdot \left[ \frac{(1 - \beta_1^t)^{-1} \cdot m_t}{\sqrt{(1 - \beta_2^t)^{-1} \cdot v_t + \epsilon}} \right] & \rho_t > 4 \end{cases} \quad (11)$$

## 3 Rectified Adam, Adaptive Variance, and Update Steps

We begin by uncovering the precise behavior of RAdam, before delving into its underlying variance-based motivation.

### 3.1 RAdam: Perform 4 Iterations of Momentum SGD, Then Use Adam with Fixed Warmup

Liu et al. describe RAdam as having two modes of operation: “divergent variance” and “convergent variance”, corresponding respectively to the cases  $\rho_t \leq 4$  and  $\rho_t > 4$  in Equation 11. In the “divergent” phase, RAdam performs a variant of heavy-ball momentum SGD (Polyak 1964).<sup>2</sup> Then, in the “convergent” phase, RAdam performs Adam, with the learning rate scaled down by  $\omega_t$ .

However, this is *not* dynamic scaling based on the training-time behavior of the optimizer *or* the distribution of the gradients. Rather,  $\omega_t$  is a deterministic function of solely  $t$  and  $\beta_2$ . Thus, the “convergent” phase is simply Adam with a fixed warmup schedule. We find that for all practically relevant values of  $\beta_2$ , the condition  $\rho_t \leq 4$  is simply  $t \leq 4$ :

**Fact 3.1.** Assume that  $0.8 \leq \beta_2 < 1$  and  $t$  is a positive integer. Then, for  $\rho_t$  as defined in Equation 7:

$$\rho_t \leq 4 \iff t \leq 4$$

*Proof.* See Appendix B.1. □

Thus follows a layman’s description of RAdam:

1. Perform four iterations of heavy-ball momentum.
2. At iteration five and beyond, use Adam with a fixed warmup schedule.

<sup>2</sup>The departure from standard heavy-ball momentum is in the bias correction by  $(1 - \beta_1^t)$ .

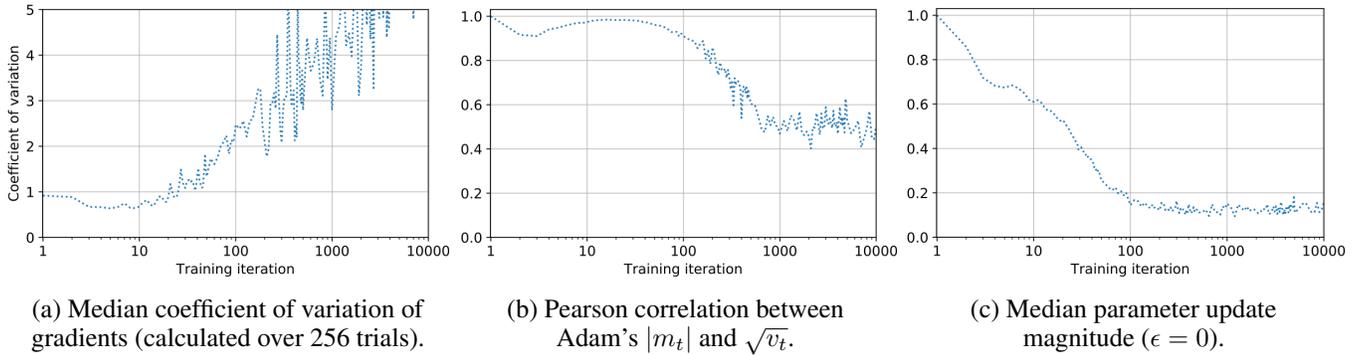


Figure 1: Analysis of gradients and updates during the training of a simple feed-forward network on the EMNIST digit recognition task with the Adam optimizer – see Appendix A.5 for comprehensive details.

On its face, using four iterations of momentum at the beginning of training seems arbitrary. In preliminary experimentation (including the experimental settings described in Section 5), we performed ablations over the following options for these four initial iterations:

- Do absolutely nothing.
- Use Adam with learning rate  $\alpha \cdot \omega_5$  (i.e. do exactly what RAdam does at the fifth iteration).
- Use Adam with linear warmup to  $\alpha \cdot \omega_5$  (i.e. gradually warm up the learning rate to RAdam’s fifth iteration).

As expected for a decision affecting only four training iterations, the practical difference between these choices is uniformly negligible. Thus, the only possible benefit of RAdam stems from its custom warmup schedule  $\omega_t$  for the fifth iteration and beyond. We revisit this topic in Sections 4 and 5.

### 3.2 Variance-Based Motivation for RAdam and Warmup

Given the arbitrary nature of RAdam’s operation, we proceed to investigate the motivation for RAdam, which Liu et al. also identify as the underlying motivation for warmup’s crucial role in Adam.

Liu et al. focus their principal analysis on the term  $\sqrt{\frac{1-\beta_2^t}{v_t}}$ . Fixing  $\epsilon = 0$ , this term can be interpreted as Adam’s “adaptive learning rate”, which scales the global learning rate for each parameter before computing Adam’s final update for that parameter. They identify that the quantity  $\text{Var} \left[ \sqrt{\frac{1-\beta_2^t}{v_t}} \right]$

does not exist during the first few training iterations,<sup>3</sup> and even after converging to a finite value, continues to remain elevated for some time.

Perhaps the most immediate observation is that early-stage gradients are *not* zero-mean. In fact, at the beginning of optimization, the expected magnitude of the gradients  $\nabla \hat{\mathcal{L}}_t(\theta_t)$

<sup>3</sup>The authors approximate  $\frac{1-\beta_2^t}{v_t}$  as having a scaled inverse  $\chi^2$  distribution, under the assumption that (1) all gradients are i.i.d. and zero-mean, and (2) a simple average approximates an exponential moving average.

(i.e. absolute value of the deterministic gradients  $\nabla \mathcal{L}(\theta_t)$ ) should dominate the gradient variance, since a randomly-initialized model is exceedingly unlikely to be near a local minimum of  $\mathcal{L}(\theta_t)$ . Indeed, on a demonstration training run of a feed-forward network on the EMNIST digit recognition task, we observe that the median coefficient of variation of the gradients (Figure 1a) starts at below 1, indicating that for most parameters, the expected value of the gradient exceeds the standard deviation during early-stage training. Only beyond training iteration 50 does the coefficient of variation consistently exceed 1. Relaxing the zero-mean assumption decreases  $\text{Var} \left[ \frac{1-\beta_2^t}{v_t} \right]$  considerably.<sup>4</sup>

More important, however, is that  $m_t$  and  $v_t$  are not at all independent. Figure 1b reveals that in the EMNIST setting, the absolute value of the first moment estimator ( $|m_t|$ ) is extremely correlated with the square root of the second moment estimator ( $\sqrt{v_t}$ ). Since Adam’s parameter updates are proportional to  $m_t/\sqrt{v_t}$ , high correlation between these two quantities implies that the magnitude of the updates are quite regular, despite the high variance of  $\sqrt{\frac{1}{v_t}}$ .

Indeed, during the first training iteration ( $t = 1$ ), it is guaranteed that  $|m_t| = \sqrt{v_t}$  for all parameters, making all Adam parameter updates either  $-\alpha$  or  $\alpha$  (assuming  $\epsilon = 0$ ). Thus, even though  $\text{Var} \left[ \frac{1-\beta_2^t}{v_t} \right]$  is divergent, the magnitude of the parameter updates themselves are constant. Ironically, it is *precisely* when the adaptive learning rate’s variance is “divergent” that the actual parameter update magnitudes have zero variance. This suggests that the adaptive learning rate may not be the best medium of analysis for understanding the role of warmup in Adam.

### 3.3 High Initial Update Step Magnitudes Necessitate Warmup in Adam

We provide an alternative view of the frequent necessity of learning rate warmup when using Adam. We do so by directly

<sup>4</sup>Although Liu et al. do not comment on the relative magnitudes of  $\text{Var} \left[ \frac{1-\beta_2^t}{v_t} \right]$ , their Fig. 9 reveals that coefficients of variation below 1 dampen that quantity by an order of magnitude or more.

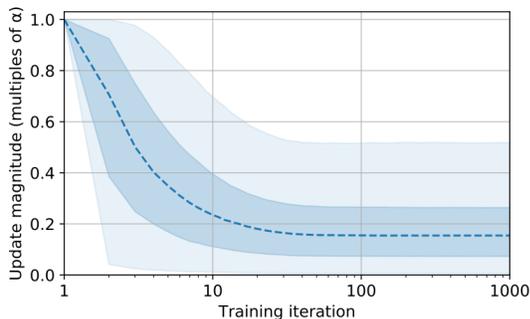


Figure 2: Distribution of Adam’s update step magnitudes at a simulated local minimum of  $\mathcal{L}(\theta)$  (quantiles:  $\{2.5\%, 25\%, 50\%, 75\%, 97.5\%\}$ ).

investigating the magnitudes of the update steps, perhaps the most proximate determinant of training stability.

In stochastic gradient descent, parameter updates are simply the gradients multiplied by the learning rate. Warmup for SGD can thus be motivated as mitigating the large expected magnitudes of the gradients (directly proportional to update magnitudes) and rapid change in gradients at the beginning of training (Goyal et al. 2017; Gotmare et al. 2019). Similar logic can be employed for adaptive methods.

On the other hand, if a model’s gradients have near-zero means and low gradient variances, the update steps are similarly well-regulated and optimization via SGD can be stable without any learning rate warmup. For example, a nearly-converged model (thus having near-zero expected gradients and low gradient magnitudes) trained via SGD can have its optimization be stably restarted without learning rate warmup.

This is not the case with Adam. We proceed to computationally analyze the magnitude of Adam’s update step over the course of training. Specifically, we demonstrate via simulation that even when the model parameters  $\theta_t$  are initialized at an idealized local minimum of  $\mathcal{L}(\theta)$  (i.e.  $\nabla \hat{\mathcal{L}}_t(\theta_t)$  has zero mean and is i.i.d. across time), the magnitude of Adam’s update steps will still be quite high at the start of training, only gradually decaying toward a stationary distribution.

**Simulation configuration** All gradients are simulated as i.i.d. normal variables with zero mean and constant isotropic variance  $10^{-9}$ , thus approximating the optimization dynamics at an exact local minimum of  $\mathcal{L}(\theta)$ .<sup>5</sup> We sample independent gradient trajectories (each 1000 iterations long) for 25000 parameters. We then run the Adam optimizer with these sampled gradients and evaluate the distribution of the update step magnitudes (before multiplication by the global learning rate  $\alpha$ ) at each iteration. The Adam optimizer configuration is  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ , and  $\epsilon = 0$ .

**Simulation results** Figure 2 depicts the outcome of this computational simulation. As alluded to in Section 3.2, the update magnitudes for all parameters start at  $1 \cdot \alpha$ . The update magnitudes gradually decay but continue to remain high for quite some time, only beginning to settle into a stationary

<sup>5</sup>Note that the behavior of Adam in this setting is invariant to the choice of variance constant.

distribution after 40 or so training iterations (with median update magnitude  $\approx 0.16 \cdot \alpha$ ). We extend the trajectory length to 10000 and find that the median update step of the stationary distribution is approximately  $0.153 \cdot \alpha$ .

These results imply that unlike SGD, Adam will always encounter early-stage training instability by way of large update magnitudes, even when the model is *already initialized at a local minimum*. This stands as a contributing factor to Adam’s need for learning rate warmup above and beyond that of first-order methods.

**Comparison to real-world, random initialization settings** Finally, we examine the update step distribution of a model initialized away from a local minimum of  $\mathcal{L}(\theta)$ . Figure 1c depicts the median parameter update magnitudes of Adam in the EMNIST setting from Section 3.2. We observe a qualitative similarity to the local minimum simulation results – the update magnitudes start at  $1 \cdot \alpha$ , only gradually settling into a stationary distribution around  $0.15 \cdot \alpha$ .

Note that the EMNIST optimization decreases more slowly in update magnitude and takes longer ( $\approx 100$  training iterations) to settle into the stationary distribution. This suggests that the update step non-regularity observed in the idealized local minimum initialization setting is only exacerbated in the more realistic setting of random initialization.

## 4 Rules of Thumb

Turning to the practical application of learning rate warmup, we first define a simple heuristic function, the *effective warmup period*, to characterize the dampening effect of warmup schedules. We then present and intuitively motivate two Adam warmup schedules that require no tuning and are thus usable as rules of thumb.

### 4.1 Effective Warmup Period

We define the *effective warmup period*  $\mathcal{T}(\omega)$  of a warmup schedule  $\omega$  as follows:

$$\mathcal{T}(\omega) = \sum_{t=1}^{\infty} (1 - \omega_t)$$

Intuitively, this is the sum of the warmup’s dampening effect across all of training.

### 4.2 Exponential Warmup

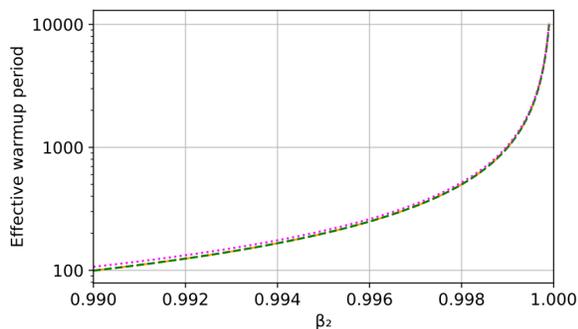
We propose a simple “exponential warmup” schedule based on a decaying exponential and a constant  $\tau$ :

$$\omega_t^{\text{expo}, \tau} = 1 - \exp\left(-\frac{1}{\tau} \cdot t\right) \quad (12)$$

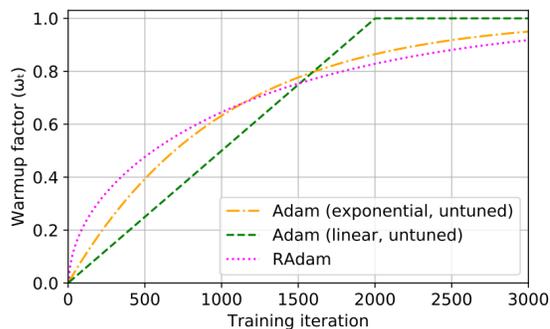
The constant  $\tau$  is analogous to a linear warmup period, and we recommend  $\tau = (1 - \beta_2)^{-1}$  as a rule of thumb:

$$\omega_t^{\text{expo}, \text{untuned}} = 1 - \exp(-(1 - \beta_2) \cdot t) \quad (13)$$

In choosing  $\tau$ , our guiding (albeit extremely speculative) intuition is to have the warmup factor  $\omega_t^{\text{expo}, \tau}$  be roughly



(a) Effective warmup periods of RAdam and rule-of-thumb warmup schedules, as a function of  $\beta_2$ .



(b) RAdam and rule-of-thumb warmup schedules over time for  $\beta_2 = 0.999$ .

Figure 3: Comparison of various characteristics of RAdam and rule-of-thumb warmup schedules.

equivalent to Adam’s second moment bias correction term in Adam. This term,  $1 - \beta_2^t$ , is the sum of the coefficients in the moving average estimation of the second moment, and can thus be interpreted as how “complete” the second moment estimator is at any given point in time. We briefly show the approximate correspondence between the bias correction term and the warmup factor:<sup>6</sup>

$$\begin{aligned} 1 - \beta_2^t &= 1 - \exp(\log(\beta_2) \cdot t) \\ &\approx 1 - \exp((\beta_2 - 1) \cdot t) \\ &= 1 - \exp(-(1 - \beta_2) \cdot t) \end{aligned}$$

### 4.3 Linear Warmup

Recall the formulation of linear warmup:

$$\omega_t^{\text{linear}, \tau} = \min \left\{ 1, \frac{1}{\tau} \cdot t \right\}$$

As a similar rule of thumb to the exponential warmup schedule, we suggest performing linear warmup over  $\tau = 2 \cdot (1 - \beta_2)^{-1}$  iterations:

$$\omega_t^{\text{linear}, \text{untuned}} = \min \left\{ 1, \frac{1 - \beta_2}{2} \cdot t \right\} \quad (14)$$

Our choice of  $\tau$  is carried over from exponential warmup as a starting point. To preserve the same effective warmup period, the  $\tau$  from the exponential rule-of-thumb is multiplied by 2 to account for the fact that exponential warmup decelerates over time, whereas linear warmup does not. We elaborate in Appendix B.2.

### 4.4 Comparison with RAdam

We first compare RAdam with the rule-of-thumb schedules (Equations 13 and 14) by computing their effective warmup

<sup>6</sup>The second step follows from a first-order Taylor expansion of  $\log(\beta_2)$  around  $\beta_2 = 1$ . In practice, this approximation is extremely accurate for typical values of  $\beta_2$ .

periods across a range of  $\beta_2$  values.<sup>7</sup> Figure 3a reveals that the effective warmup periods of RAdam and the rules of thumb are nearly identical across all practical values of  $\beta_2$ , indicating that they have similar dampening effects over early-stage training.

We then proceed to examine the trajectory of the warmup schedule for the commonly used setting of  $\beta_2 = 0.999$ . Figure 3b reveals that the functional forms of the warmup factors are qualitatively similar in magnitudes. The warmup schedules for RAdam and the rule-of-thumb exponential warmup closely correspond in shape as well.

We thus posit that RAdam and the untuned rule-of-thumb warmup schedules are more or less interchangeable. An empirical verification follows.

## 5 Experiments

We evaluate untuned exponential warmup (Equation 13), untuned linear warmup (Equation 14), and RAdam across a variety of supervised machine learning tasks. For brevity, all experimental settings are summarized in the main text and comprehensively detailed in Appendix A.

### 5.1 Image Classification

Using each of the three warmup methods, we train a ResNet-50 model (He et al. 2016) on the ILSVRC (“ImageNet”) image classification dataset with various configurations of Adam. Specifically, we sweep over:

$$\begin{aligned} \alpha \text{ (learning rate)} &\in \{10^{-4}, 10^{-3}, 10^{-2}\} \\ \beta_2 &\in \{0.99, 0.997, 0.999\} \end{aligned}$$

Table 1 presents the top-1 error rates at the end of training for the three warmup methods. Across all configurations of Adam, the top-1 error rates are indistinguishable between the warmup methods.<sup>8</sup>

<sup>7</sup>For the purpose of this analysis,  $w_{\{1,2,3,4\}}$  are all defined to be zero for RAdam.

<sup>8</sup>The best error rates fall roughly 3% behind those from SGD, as is typical with Adam on computer vision tasks.

LR	$\beta_2$	Exponential	Linear	RAdam
$10^{-4}$	0.99	$34.2\% \pm 0.1$	$34.2\% \pm 0.1$	$34.2\% \pm 0.1$
$10^{-4}$	0.997	$34.3\% \pm 0.2$	$34.2\% \pm 0.2$	$34.1\% \pm 0.1$
$10^{-4}$	0.999	$34.5\% \pm 0.1$	$34.4\% \pm 0.1$	$34.2\% \pm 0.3$
$10^{-3}$	0.99	$27.9\% \pm 0.1$	$28.0\% \pm 0.1$	$28.4\% \pm 0.1$
$10^{-3}$	0.997	$27.9\% \pm 0.1$	$27.9\% \pm 0.1$	$28.3\% \pm 0.1$
$10^{-3}$	0.999	$28.2\% \pm 0.1$	$28.3\% \pm 0.1$	$28.4\% \pm 0.1$
$10^{-2}$	0.99	$29.3\% \pm 0.1$	$29.3\% \pm 0.3$	$29.4\% \pm 0.2$
$10^{-2}$	0.997	$29.2\% \pm 0.2$	$29.3\% \pm 0.1$	$29.4\% \pm 0.5$
$10^{-2}$	0.999	$28.9\% \pm 0.2$	$28.7\% \pm 0.1$	$29.8\% \pm 0.4$

Table 1: Top-1 error rates of ResNet-50 on ImageNet (means and standard deviations over 5 random seeds).

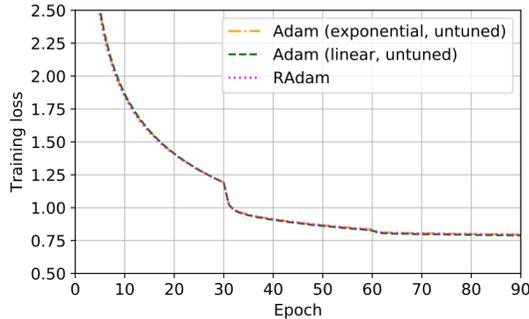


Figure 4: Mean training loss (5 seeds) of ResNet-50 on ImageNet, using Adam with  $\alpha = 10^{-3}$  and  $\beta_2 = 0.999$ .

We next examine the course of optimization for individual configurations of Adam’s  $\alpha$  and  $\beta_2$ . Figure 4 depicts the training loss using the popular “default” Adam configuration of learning rate  $\alpha = 10^{-3}$  and  $\beta_2 = 0.999$ , revealing that the behavior of these warmup methods is indeed nearly indistinguishable.

Appendix C.1 provides both training and validation metrics (Figures 7 and 8 respectively) for all tested configurations, reinforcing this trend.

## 5.2 Language Modeling

Using each of the three warmup methods, we train a state-of-the-art Transformer-based language model from Baevski and Auli (2018) on WIKITEXT-103. We sweep over the following grid of Adam hyperparameters:

$$\alpha \text{ (learning rate)} \in \{1 \cdot 10^{-4}, 3 \cdot 10^{-4}, 5 \cdot 10^{-4}\}$$

$$\beta_2 \in \{0.99, 0.998, 0.999\}$$

with  $\beta_1 = 0.9$  and  $\epsilon = 10^{-7}$  fixed. As with image classification, we observe in Table 2 that the choice of warmup method has a minimal impact on training across different hyperparameters.

Figure 5 depicts the validation perplexity throughout training for the best Adam parametrization ( $\alpha = 10^{-4}$  and  $\beta_2 = 0.999$ ), which similarly supports the indistinguishability of the warmup methods.

LR	$\beta_2$	Exponential	Linear	RAdam
$1 \cdot 10^{-4}$	0.99	$21.0 \pm 0.1$	$21.0 \pm 0.1$	$21.1 \pm 0.1$
$1 \cdot 10^{-4}$	0.998	$19.9 \pm 0.0$	$19.9 \pm 0.0$	$20.0 \pm 0.0$
$1 \cdot 10^{-4}$	0.999	$20.0 \pm 0.0$	$20.0 \pm 0.0$	$20.1 \pm 0.1$
$3 \cdot 10^{-4}$	0.99	$21.3 \pm 0.3$	$20.8 \pm 0.1$	$22.4 \pm 0.0$
$3 \cdot 10^{-4}$	0.998	$19.6 \pm 0.0$	$19.6 \pm 0.0$	$19.6 \pm 0.1$
$3 \cdot 10^{-4}$	0.999	$19.5 \pm 0.0$	$19.5 \pm 0.0$	$19.5 \pm 0.0$
$5 \cdot 10^{-4}$	0.99	$24.4 \pm 2.4$	$24.1 \pm 1.4$	$26.0 \pm 1.8$
$5 \cdot 10^{-4}$	0.998	$20.1 \pm 0.0$	$20.0 \pm 0.0$	$20.1 \pm 0.0$
$5 \cdot 10^{-4}$	0.999	$19.8 \pm 0.0$	$19.7 \pm 0.1$	$19.7 \pm 0.0$

Table 2: Validation perplexity of a Transformer LM on the WIKITEXT-103 dataset (means and standard deviations over 3 random seeds).

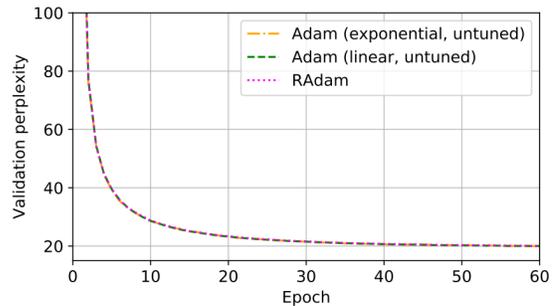


Figure 5: Mean validation perplexity (3 seeds) of Transformer LM on WIKITEXT-103, using Adam with  $\alpha = 10^{-4}$  and  $\beta_2 = 0.999$ .

## 5.3 Machine Translation

Finally, we evaluate the warmup methods on a large scale machine translation task. Using each of the three warmup methods, we train a Transformer model (Vaswani et al. 2017) on the WMT16 English-German (“EN-DE”) dataset. We fix Adam’s  $\beta_1 = 0.9$  and  $\epsilon = 10^{-7}$  and sweep over the following grid of Adam hyperparameters:

$$\alpha \text{ (learning rate)} \in \{5 \cdot 10^{-5}, 8 \cdot 10^{-5}, 1 \cdot 10^{-4}\}$$

$$\beta_2 \in \{0.98, 0.99, 0.998, 0.999\}$$

We observe no perceptible differences between the warmup methods in either final performance (Table 3), or in the training-time metrics of a single canonical configuration ( $\alpha = 10^{-4}$  and  $\beta_2 = 0.999$ , shown in Figure 6).

## 6 Discussion

We discuss various consequences of our findings, along with directions for future work.

### 6.1 Extended Warmup Periods

The analysis of the update step magnitudes in Section 3.3 suggests shorter warmup periods than those typically used in practice. For example, using the setting of  $\beta_2 = 0.999$ ,

LR	$\beta_2$	Exponential	Linear	RAdam
$5 \cdot 10^{-5}$	0.98	$24.5 \pm 0.1$	$24.4 \pm 0.1$	$24.4 \pm 0.1$
$5 \cdot 10^{-5}$	0.99	$24.5 \pm 0.0$	$24.5 \pm 0.0$	$24.5 \pm 0.1$
$5 \cdot 10^{-5}$	0.998	$24.3 \pm 0.2$	$24.4 \pm 0.2$	$24.4 \pm 0.1$
$5 \cdot 10^{-5}$	0.999	$24.2 \pm 0.1$	$24.2 \pm 0.1$	$24.1 \pm 0.1$
$8 \cdot 10^{-5}$	0.98	$25.9 \pm 0.1$	$25.9 \pm 0.1$	$25.9 \pm 0.1$
$8 \cdot 10^{-5}$	0.99	$25.9 \pm 0.2$	$25.9 \pm 0.1$	$25.9 \pm 0.0$
$8 \cdot 10^{-5}$	0.998	$26.0 \pm 0.1$	$25.2 \pm 1.0$	$25.9 \pm 0.1$
$8 \cdot 10^{-5}$	0.999	$25.7 \pm 0.1$	$25.8 \pm 0.1$	$25.7 \pm 0.0$
$1 \cdot 10^{-4}$	0.98	$26.5 \pm 0.1$	$26.6 \pm 0.1$	$26.6 \pm 0.1$
$1 \cdot 10^{-4}$	0.99	$26.7 \pm 0.1$	$26.6 \pm 0.1$	$26.6 \pm 0.0$
$1 \cdot 10^{-4}$	0.998	$25.9 \pm 0.9$	$26.5 \pm 0.1$	$26.6 \pm 0.0$
$1 \cdot 10^{-4}$	0.999	$26.2 \pm 0.2$	$26.4 \pm 0.0$	$26.4 \pm 0.0$

Table 3: BLEU score of Transformer on WMT16-EN-DE (means and standard deviations over 3 random seeds).

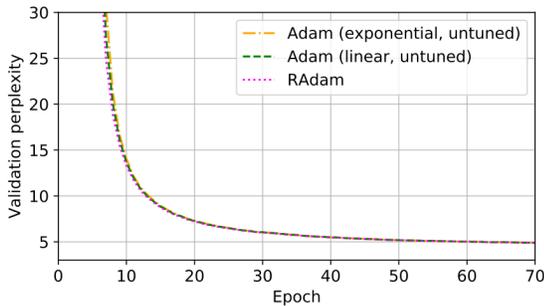


Figure 6: Mean validation perplexity (3 seeds) of Transformer on WMT16-EN-DE, using Adam with  $\alpha = 10^{-4}$  and  $\beta_2 = 0.999$ .

Adam’s update magnitudes in the theoretical model converge to a stationary distribution in roughly 40 iterations. If update magnitudes were the only relevant consideration, then a warmup schedule over a few hundred iterations would suffice to stabilize training. In contrast, the effective warmup periods of both RAdam and our rule-of-thumb schedules are roughly 1000 iterations for  $\beta_2 = 0.999$ . State-of-the-art methods with hand-tuned warmup schedules often go well beyond, using up to 10000 iterations of linear warmup in some cases (Liu et al. 2019b; Baevski and Auli 2018; Ott et al. 2019).

Accordingly, we surmise that the precise channel by which Adam necessitates an extended period of warmup is still an unresolved question, likely related to the properties of the gradients at random initialization. Future work could rigorously investigate the effect of extended warmup periods on the training dynamics of Adam, beyond simple per-iteration statistics.

## 6.2 Consequences of Update Step Invariance to Gradients

One ancillary finding of Section 3.3 is that the magnitudes of Adam’s update steps during later stages of training are largely invariant to the properties or dynamics of the gradient dis-

tribution – both the simulated local optimum and real-world random initialization settings result in convergence to similar stationary distributions of update magnitudes. This suggests that learning rate *decay* at later stages of training could be the only way to improve late-stage convergence, as Adam’s late-stage update magnitudes do not appear to be very sensitive to the variance or stationarity of gradients. In particular, we suspect that variance-based methods of improving the late-stage convergence of SGD, such as increasing the batch size (Smith et al. 2018), will not yield comparable benefits when applied to Adam, as the stationary distribution of the update magnitudes will remain largely the same. Partially adaptable methods (Chen and Gu 2018; Keskar and Socher 2017; Luo et al. 2019), which interpolate between the full adaptivity of Adam and the non-adaptivity of SGD, may hold more promise for improving late-stage convergence.

## 6.3 Dynamic Warmup

All methods considered by this work use fixed warmup schedules, computed only as a function of the training iteration  $t$  and various hyperparameters. Such schedules will inevitably be brittle to some combination of problem setting, model, and optimizer configuration. Another direction for future work could be to devise truly dynamic mechanisms for scheduling warmup in Adam. Such a mechanism could (among other things) track and utilize auxiliary statistics, such as the running moments of the applied updates, in order to determine the stability of training at each iteration.

This direction comes dangerously close to seeking the “holy grail” of an automatic learning rate tuner; existing attempts to devise such a method have achieved limited adoption as of yet (Li, Tai, and E 2017; Zhang, Mitliagkas, and Ré 2017; Baydin et al. 2018). What makes this potentially more tractable is that a maximum learning rate is still tuned and given *a priori* to the optimizer; the task is then restricted to dynamic scheduling of the learning rate from zero to this known constant, instead of an arbitrary range  $(0, \infty)$ .

## 7 Conclusion

We show that the Rectified Adam (RAdam) algorithm can be characterized as four steps of momentum SGD, followed by Adam with a fixed warmup schedule. We also examine the shortcomings of a variance-based approach to analyzing the learning rate warmup heuristic, and we illustrate that Adam’s frequent need for learning rate warmup can be partially explained by inspecting Adam’s early-stage update step magnitudes when applied to an already-converged model.

RAdam’s claimed benefits are its superior performance to Adam and its elimination of costly warmup schedule tuning. We obviate RAdam by providing two simple “rule-of-thumb” warmup schedules for Adam, both of which require no tuning. Linear warmup of Adam’s learning rate over  $2 \cdot (1 - \beta_2)^{-1}$  iterations is functionally equivalent to RAdam across a wide range of settings. Hence, we suggest that practitioners considering the need for untuned warmup of Adam’s learning rate first try linear warmup over  $2 \cdot (1 - \beta_2)^{-1}$  training iterations.

## References

- Baevski, A.; and Auli, M. 2018. Adaptive Input Representations for Neural Language Modeling. *CoRR* abs/1809.10853. URL <http://arxiv.org/abs/1809.10853>.
- Baydin, A. G.; Cornish, R.; Martínez-Rubio, D.; Schmidt, M.; and Wood, F. 2018. Online Learning Rate Adaptation with Hypergradient Descent. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. URL <https://openreview.net/forum?id=BkrsAzWAb>.
- Brown, T. B.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; Agarwal, S.; Herbert-Voss, A.; Krueger, G.; Henighan, T.; Child, R.; Ramesh, A.; Ziegler, D. M.; Wu, J.; Winter, C.; Hesse, C.; Chen, M.; Sigler, E.; Litwin, M.; Gray, S.; Chess, B.; Clark, J.; Berner, C.; McCandlish, S.; Radford, A.; Sutskever, I.; and Amodei, D. 2020. Language Models are Few-Shot Learners. *CoRR* abs/2005.14165. URL <https://arxiv.org/abs/2005.14165>.
- Chen, J.; and Gu, Q. 2018. Closing the Generalization Gap of Adaptive Gradient Methods in Training Deep Neural Networks. *CoRR* abs/1806.06763. URL <http://arxiv.org/abs/1806.06763>.
- Cohen, G.; Afshar, S.; Tapson, J.; and van Schaik, A. 2017. EMNIST: Extending MNIST to handwritten letters. In *2017 International Joint Conference on Neural Networks, IJCNN 2017, Anchorage, AK, USA, May 14-19, 2017*, 2921–2926. doi:10.1109/IJCNN.2017.7966217. URL <https://doi.org/10.1109/IJCNN.2017.7966217>.
- Devlin, J.; Chang, M.; Lee, K.; and Toutanova, K. 2018. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *CoRR* abs/1810.04805. URL <http://arxiv.org/abs/1810.04805>.
- Duchi, J. C.; Hazan, E.; and Singer, Y. 2010. Adaptive Subgradient Methods for Online Learning and Stochastic Optimization. In *COLT 2010 - The 23rd Conference on Learning Theory, Haifa, Israel, June 27-29, 2010*, 257–269. URL <http://colt2010.haifa.il.ibm.com/papers/COLT2010proceedings.pdf#page=265>.
- Gehring, J.; Auli, M.; Grangier, D.; Yarats, D.; and Dauphin, Y. N. 2017. Convolutional Sequence to Sequence Learning. In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, 1243–1252. URL <http://proceedings.mlr.press/v70/gehring17a.html>.
- Gotmare, A.; Keskar, N. S.; Xiong, C.; and Socher, R. 2019. A Closer Look at Deep Learning Heuristics: Learning rate restarts, Warmup and Distillation. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. URL <https://openreview.net/forum?id=r14EOsCqKX>.
- Goyal, P.; Dollár, P.; Girshick, R. B.; Noordhuis, P.; Wesolowski, L.; Kyrola, A.; Tulloch, A.; Jia, Y.; and He, K. 2017. Accurate, Large Minibatch SGD: Training ImageNet in 1 Hour. *CoRR* abs/1706.02677. URL <http://arxiv.org/abs/1706.02677>.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, 770–778. doi:10.1109/CVPR.2016.90. URL <https://doi.org/10.1109/CVPR.2016.90>.
- Hinton, G.; Srivastava, N.; and Swersky, K. 2012. Neural networks for machine learning: Lecture 6a.
- Keskar, N. S.; and Socher, R. 2017. Improving Generalization Performance by Switching from Adam to SGD. *CoRR* abs/1712.07628. URL <http://arxiv.org/abs/1712.07628>.
- Kingma, D. P.; and Ba, J. 2014. Adam: A Method for Stochastic Optimization. *CoRR* abs/1412.6980. URL <http://arxiv.org/abs/1412.6980>.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. ImageNet Classification with Deep Convolutional Neural Networks. In *Advances in Neural Information Processing Systems 25: 26th Annual Conference on Neural Information Processing Systems 2012. Proceedings of a meeting held December 3-6, 2012, Lake Tahoe, Nevada, United States.*, 1106–1114. URL <http://papers.nips.cc/paper/4824-imagenet-classification-with-deep-convolutional-neural-networks>.
- Li, Q.; Tai, C.; and E, W. 2017. Stochastic Modified Equations and Adaptive Stochastic Gradient Algorithms. In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, 2101–2110. URL <http://proceedings.mlr.press/v70/li17f.html>.
- Liu, L.; Jiang, H.; He, P.; Chen, W.; Liu, X.; Gao, J.; and Han, J. 2020. On the Variance of the Adaptive Learning Rate and Beyond. In *International Conference on Learning Representations*. URL <https://openreview.net/forum?id=rkgz2aEKDr>.
- Liu, Y.; Albanie, S.; Nagrani, A.; and Zisserman, A. 2019a. Use What You Have: Video retrieval using representations from collaborative experts. In *30th British Machine Vision Conference 2019, BMVC 2019, Cardiff, UK, September 9-12, 2019*, 279. BMVA Press. URL <https://bmv2019.org/wp-content/uploads/papers/0363-paper.pdf>.
- Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; and Stoyanov, V. 2019b. RoBERTa: A Robustly Optimized BERT Pretraining Approach. *CoRR* abs/1907.11692. URL <http://arxiv.org/abs/1907.11692>.
- Loshchilov, I.; and Hutter, F. 2019. Decoupled Weight Decay Regularization. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. URL <https://openreview.net/forum?id=Bkg6RiCqY7>.
- Luo, L.; Xiong, Y.; Liu, Y.; and Sun, X. 2019. Adaptive Gradient Methods with Dynamic Bound of Learning Rate. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. URL <https://openreview.net/forum?id=Bkg3g2R9FX>.

- Ma, J.; and Yarats, D. 2019. Quasi-hyperbolic momentum and Adam for deep learning. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. URL <https://openreview.net/forum?id=S1fUpoR5FQ>.
- Nesterov, Y. E. 1983. A method for solving the convex programming problem with convergence rate  $\mathcal{O}(1/k^2)$ . In *Dokl. Akad. Nauk SSSR*, volume 269, 543–547.
- Nguyen, T. Q.; and Salazar, J. 2019. Transformers without Tears: Improving the Normalization of Self-Attention. *CoRR* abs/1910.05895. URL <http://arxiv.org/abs/1910.05895>.
- Ott, M.; Edunov, S.; Baevski, A.; Fan, A.; Gross, S.; Ng, N.; Grangier, D.; and Auli, M. 2019. fairseq: A Fast, Extensible Toolkit for Sequence Modeling. In *Proceedings of NAACL-HLT 2019: Demonstrations*.
- Paszke, A.; Gross, S.; Chintala, S.; Chanan, G.; Yang, E.; DeVito, Z.; Lin, Z.; Desmaison, A.; Antiga, L.; and Lerer, A. 2016. PyTorch Examples. <https://github.com/pytorch/examples>.
- Paszke, A.; Gross, S.; Chintala, S.; Chanan, G.; Yang, E.; DeVito, Z.; Lin, Z.; Desmaison, A.; Antiga, L.; and Lerer, A. 2017. Automatic differentiation in PyTorch. In *NIPS-W*.
- Polyak, B. 1964. Some methods of speeding up the convergence of iteration methods. *USSR Computational Mathematics and Mathematical Physics* 4(5): 1–17.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; and Sutskever, I. 2019. Language models are unsupervised multi-task learners. *OpenAI Blog*.
- Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; Berg, A. C.; and Fei-Fei, L. 2015. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)* 115(3): 211–252. doi: 10.1007/s11263-015-0816-y.
- Smith, S. L.; Kindermans, P.; Ying, C.; and Le, Q. V. 2018. Don't Decay the Learning Rate, Increase the Batch Size. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. URL <https://openreview.net/forum?id=B1Yy1BxCZ>.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention Is All You Need. *CoRR* abs/1706.03762. URL <http://arxiv.org/abs/1706.03762>.
- Yamamoto, R.; Song, E.; and Kim, J.-M. 2020. Parallel WaveGAN: A fast waveform generation model based on generative adversarial networks with multi-resolution spectrogram. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 6199–6203. IEEE.
- Zhang, J.; Mitliagkas, I.; and Ré, C. 2017. YellowFin and the Art of Momentum Tuning. *CoRR* abs/1706.03471. URL <http://arxiv.org/abs/1706.03471>.