
Analyzing & Reducing the Need for Learning Rate Warmup in GPT Training

Atli Kosson

Bettina Messmer

Martin Jaggi

EPFL, Switzerland

firstname.lastname@epfl.ch

Abstract

Learning Rate Warmup is a popular heuristic for training neural networks, especially at larger batch sizes, despite limited understanding of its benefits. Warmup decreases the update size $\Delta\mathbf{w}_t = \eta_t \mathbf{u}_t$ early in training by using lower values for the learning rate η_t . In this work we argue that warmup benefits training by keeping the overall size of $\Delta\mathbf{w}_t$ limited, counteracting large initial values of \mathbf{u}_t . Focusing on small-scale GPT training with AdamW/Lion, we explore the following question: *Why and by which criteria are early updates \mathbf{u}_t too large?* We analyze different metrics for the update size including the ℓ_2 -norm, resulting directional change, and impact on the representations of the network, providing a new perspective on warmup. In particular, we find that warmup helps counteract large angular updates as well as a limited critical batch size early in training. Finally, we show that the need for warmup can be significantly reduced or eliminated by modifying the optimizer to explicitly normalize \mathbf{u}_t based on the aforementioned metrics.

1 Introduction

Neural networks are typically trained using variations of stochastic gradient descent. The weight updates $\Delta\mathbf{w}$ have the form $\Delta\mathbf{w} = \eta\mathbf{u}$, where η denotes the *learning rate* and \mathbf{u} an unscaled update vector derived from the history of weight gradients. Throughout training, the learning rate is often adjusted over time t according to a *learning rate schedule*, $\eta = \eta_t$. This schedule frequently includes an initial phase known as a learning rate warmup, where the learning rate starts low and is increased to a target value before being reduced according to a decay schedule. Both the choice of warmup and decay strategy can significantly affect the final model performance. In this work, we focus on the linear warmup introduced by Goyal et al. [10] for large batch size ResNet [11] training, which is also commonly used for transformers [37].

The length of the warmup is a hyperparameter that requires tuning, which is complicated by the fact that the reasons for its effectiveness are somewhat unclear. Warmup empirically helps stabilize training and allows the use of larger learning rates throughout the rest of training, which can speed up the process and provide beneficial regularization [10]. Since the learning rate simply scales the size of the updates $\Delta\mathbf{w} = \eta\mathbf{u}$, warmup must achieve these effects by decreasing the size of early updates. However, it is not fully clear why this helps. *Are the initial updates too large for some reason?* For example, we might need small η_t values to counteract large \mathbf{u} values early in training. *How should we quantify what makes an update $\Delta\mathbf{w}$ large?* *Why do large updates adversely affect training?*

In this work, we explore warmup from this perspective, focusing on GPT2 [29] training with adaptive optimizers like AdamW [24] and Lion [3]. We identify three key issues that necessitate warmup:

1. The way Adam handles momentum can lead to artificially large initial updates $\Delta\mathbf{w}$.
2. Early updates $\Delta\mathbf{w}$ are large compared to the initial weight magnitude of \mathbf{w} for matrices.
3. The gradients of early samples are highly correlated, limiting effective mini-batch sizes.

We demonstrate that simple modifications to the optimizer can mitigate the first two issues: eliminating the momentum bias correction in AdamW and scaling matrix updates to match their magnitude, akin to the Rotational Optimizers by Kosson et al. [19]. We analyze the third issue in terms of the rate at which the internal neural representations of the network are changing (sometimes called feature learning). When the gradients of different samples are highly correlated these internal representations change too fast, which we conjecture can lead to issues with the non-linearities (e.g. activation functions) of the network. This can also be seen as the *critical batch size* [28] being too low early in training to enable the use of the peak learning rate. We derive a scaling factor based on the signal-to-noise ratio of the gradient that can help mitigate this, functioning like an automatic learning rate warmup. Alternatively, we show that using high momentum values in conjunction to the first two methods may suffice to enable efficient training without warmup.

2 Related Work

Learning rate warmups have been used since at least ResNet [11], where a lower constant learning rate was applied at the start of training. Earlier works may have employed similar concepts; for example, Sutskever et al. [35] utilized a momentum schedule that could induce a similar effect in the “effective learning rate” as defined by Fu et al. [6]. The practice of linear warmup in its current form was popularized by Goyal et al. [10] and Vaswani et al. [37].

Warmup has been studied indirectly in various neural network optimizer works. A notable example is RAdam [23], a modification of Adam [18] aimed at reducing the need for warmup. However, Ma and Yarats [26] demonstrated that RAdam essentially incorporates a fixed warmup schedule within the optimizer. In appx. C.1 we show that this warmup effect is insufficient in our setup. Relative optimizers like LARS [47] and LAMB [48] are also believed to reduce the necessity for warmup [21]. Bernstein et al. [2] propose a relative optimizer called Fromage and analyze how relative weight changes relate to representation changes, but differ from our approach in that they do not describe the effects of the gradient signal-to-noise ratio on this relationship. We build upon Kosson et al. [19] which showed that weight decay can make standard optimizers function as approximate relative optimizers and proposed variants that reduce the benefit of warmup without fully eliminating it.

The effect of warmup in transformers was empirically studied by Wortsman et al. [41]. Xiong et al. [43] proposed the pre-LN normalization placement for transformers, showing it reduces the need for warmup. Huang et al. [13] studied initialization in transformers showing a link to warmup.

Finally, warmup has been studied directly on its own. Gotmare et al. [9] studied the effect of warmup, finding it helps avoid overly large updates to the weights of later layers which could be frozen to achieve a similar benefit. Gilmer et al. [7] study the need for warmup from a curvature perspective, showing it may help “push” the optimization trajectory towards flatter regions where higher learning rates are stable. Smith et al. [32] arrive at a similar conclusion, there is a stable learning rate that varies throughout training based on the curvature which limits the learning rate early on, necessitating warmup. These works focus on SGD with momentum, but it is less clear how curvature affects Adam-like or relative optimizers as we discuss later.

The relation between stochastic gradient noise and learning rate has been studied in several works [46, 28, 49, 34, 22, 27]. They find that the update size can be increased roughly linearly with the batch size up to a certain *critical batch size* that depends on ratio of the mean and variance of the mini-batch gradient. We show how the signal-to-noise ratio (SNR) of the mini-batch gradient amplifies changes to the neural representations of a network given a normalized update in weight space. We observe that the SNR starts out high but decreases over time, which translates to large early changes in the internal representations without warmup.

3 Baseline Experimental Setup & Results

Our main experiments focus on the training of a 124M parameter GPT2 [29] model on the OpenWebText corpus [8]. The model has 12 transformer blocks with an embedding dimension of 768. Our base training is performed at batch size 480 with a sequence length of 1024. We train for 5000 iterations which translates into roughly 20 tokens per parameter, as suggested by Chinchila [12]. The baselines use AdamW [24] (see algo. 1) with weight decay $\lambda = 0.1$, momentum coefficient $\beta_1 = 0.9$, smoothing coefficient $\beta_2 = 0.95$, and $\epsilon = 10^{-8}$. The learning rate schedule consists of a

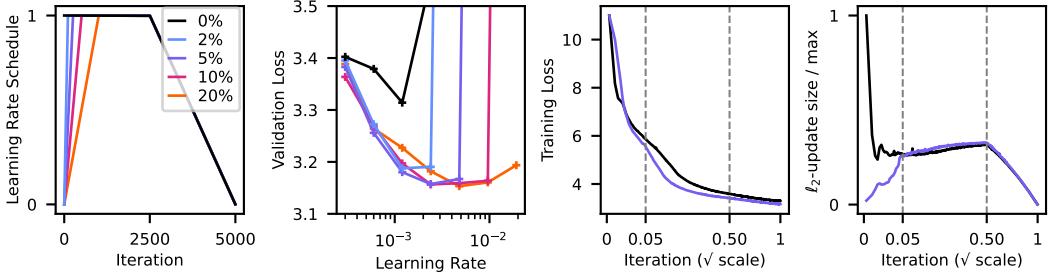


Figure 1: Warmup significantly benefits GPT2 training with AdamW. **Panel 1:** Trapezoidal learning rate schedules with different warmup lengths and 50% linear cooldown. **Panel 2:** Final validation loss for various learning rate and warmup configurations. Note the performance gap between no-warmup (black) and other configurations. **Panel 3:** Training curves comparing the best no-warmup run to a 5% warmup with the same learning rate. The warmup run quickly surpasses the no-warmup run. **Panel 4:** Comparison of ℓ_2 update norms for these runs shows large initial updates without warmup.

linear warmup followed by a constant phase and eventually linear cooldown spanning half of training (see examples in fig. 1). This schedule keeps the peak learning rate and decay phase identical for different warmup lengths. This differs from other schedules, e.g. cosine, where the warmup length typically affects the whole shape. The learning rate value and the warmup length are swept for various configurations. Our code is based on NanoGPT [16] with additional utilities by Kosson et al. [19]. The hyperparameter values and base training configuration are adopted from NanoGPT. See appx. C.2 for experiments on additional architectures and datasets.

Figure 1 shows the baseline performance for our setup. We observe that even short warmup can significantly improve performance. Not using warmup results in faster initial progress for a given learning rate, but eventually falls behind leaving a permanent gap. Warmup not only stabilizes higher learning rates, but also prevents a lasting degradation of the model that can not be mitigated by simply training for slightly longer. We notice that although Adam normalizes the update size, its ℓ_2 -magnitude varies significantly throughout training with a large spike at the start of training.

4 The Interaction of Momentum and the ℓ_2 -Update Norm in AdamW

In this section we analyze the reason behind the large ℓ_2 -norm of early updates in our AdamW baseline seen in panel 4 of fig. 1. We find that this primarily stems from the β_1 bias correction. We then explore to what extent these large initial updates contribute to the need for warmup by modifying the optimizer to directly control the ℓ_2 -norm of the update. Although we find this is to be insufficient to replace warmup on its own, these changes are an important component of our later methods.

Adam-like optimizers such as AdamW (algo. 1) differ from simpler methods like SGD with momentum in that they normalize the update size with the gradient magnitude. This makes them invariant to a rescaling of the loss function and helps counteract potential differences in the gradient magnitude between layers. An important consequence of this is that the unscaled updates u are not large simply due to large initial gradients, unlike in plain SGD and other optimizers that don't normalize their

Algorithm 1 AdamW (PyTorch variant, differing from the original by Loshchilov and Hutter [24])

Require: Learning rate η_t , weight decay λ , momentum β_1 , magnitude smoothing β_2, ε for numerical stability
1: **Initialize:** Time step $t \leftarrow 0$, parameter vector θ_0 , momentum vector $m_0 \leftarrow 0$, magnitude vector $v_0 \leftarrow 0$
2: **while** stopping criteria not met :
3: $t \leftarrow t + 1$
4: $g_t \leftarrow$ Mini-batch gradient w.r.t. θ_{t-1}
5: $m_t \leftarrow \beta_1 m_{t-1} + (1 - \beta_1) g_t$
6: $v_t \leftarrow \beta_2 v_{t-1} + (1 - \beta_2) g_t^2$
7: $\hat{m}_t \leftarrow m_t / (1 - \beta_1^t)$
8: $\hat{v}_t \leftarrow v_t / (1 - \beta_2^t)$
9: $\theta_t \leftarrow (1 - \eta_t \lambda) \theta_{t-1} - \eta_t \hat{m}_t / (\sqrt{\hat{v}_t} + \varepsilon)$

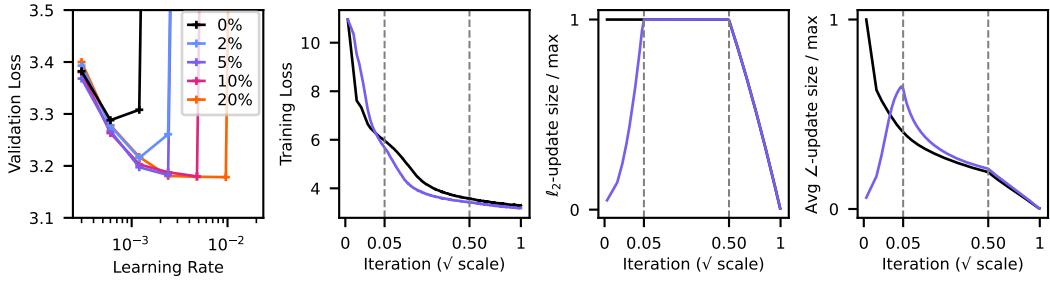


Figure 2: LionA (algo. 2) fails to significantly reduce the warmup advantage. **Panel 1:** Final validation loss across various learning rates and warmup percentages shows a reduced but still significant no-warmup penalty compared to AdamW (fig. 1). **Panel 2:** Training curves for 0% vs. 5% warmup at the highest stable learning rate for 0%, with warmup quickly overtaking no-warmup as before. **Panel 3:** LionA successfully controls the ℓ_2 -update norm. **Panel 4:** Early angular updates (see §5) are large without warmup and do not follow the learning rate schedule throughout training.

updates. Such un-normalized optimizers might diverge to infinity if a high learning rate is combined with large initial gradients or large curvature, as the update size is unbounded. Preventing this could be an additional benefit of warmup for SGD on top of the effects discussed in this work.

Although AdamW normalizes the update size based on the gradient, its magnitude can still vary throughout training as seen in fig. 1. This can be caused by changes in the gradient magnitude over time, especially when using different values of β_1 and β_2 . However, it can also be caused by momentum and especially the bias correction (algo. 1, line 7). The magnitude of \mathbf{m}_t depends on the alignment of subsequent gradients $\mathbf{g}_1, \dots, \mathbf{g}_t$ whereas the normalization factor \mathbf{v}_t does not. For example, when each \mathbf{g}_t is an independent zero-mean random vector with a fixed second moment $\mathbb{E}[\mathbf{g}_t^2] = \sigma^2$, we have (see appx. B.1 for details):

$$\mathbb{E}[\mathbf{m}_t^2] = (1 - \beta_1^{2t}) \frac{1 - \beta_1}{1 + \beta_1} \sigma^2, \quad \mathbb{E}[\mathbf{v}_t] = (1 - \beta_2^t) \sigma^2 \quad (1)$$

In this case the bias correction for β_1 is incorrect since it is derived for a constant gradient. With the bias correction the size becomes $\mathbb{E}[\|\hat{\mathbf{m}}\|^2] = \frac{1 + \beta_1^t}{1 - \beta_1^t} \frac{1 - \beta_1}{1 + \beta_1} \sigma^2$, amplifying the norm of early updates by $\sqrt{(1 + \beta_1^t)/(1 - \beta_1^t)}$. This factor is larger if the gradients between successive steps are negatively correlated, which we empirically observe happening in our setup (see §6.2).

The ℓ_2 -norm of AdamW updates can therefore vary significantly due to the initial bias correction, changes in the alignment of the gradients throughout training, and potential variations in the gradient norm over time. Lion [3] is a closely related optimizer that uses an element-wise sign operation to normalize the update, giving +1 for positive values, -1 for negative values and 0 for zeros. Ignoring the possibility of zeros, this gives a constant update norm. Lion is closely related to Adam, and can be obtained by tracking the size of \mathbf{m}_t instead of \mathbf{g}_t in line 6 while setting $\beta_2 = 0$. It also uses a form of scaled Nesterov momentum instead of the traditional heavy-ball variant and a hyperparameter specification that differs significantly from AdamW. We propose a Lion variant, LionA (algo. 2), that keeps the hyperparameters more compatible with those of AdamW. The learning rate is kept comparable by scaling the ℓ_2 update size to match that of AdamW in the random-gradient scenario, see appx. B.1 for the derivation of the scaling factors. Due to its ability to perfectly control the size of each update, we use Lion based methods for the remaining of this paper. This avoids confounding effects in Adam-like optimizers, such as \mathbf{v} being inaccurate from rapidly decreasing gradient magnitudes early in training, which can induce an additional warmup-like effect.

In fig. 2 we repeat the baseline sweep using LionA. **Despite perfect control of the ℓ_2 update norm (as seen in panel 3), the benefit of warmup remains. This leads us to conclude that the ℓ_2 update size is not sufficient to quantify the “effectively” large updates that we conjecture warmup mitigates.** The final panel shows that the angular update size (see definition in the following section), proposed to be a better measure of an effective step size by Wan et al. [38], still varies throughout training with a spike at the start of training. In the next section we explore the reasons for the large initial angular updates and how they significantly contribute to the need for warmup.

Algorithm 2 LionA: A modified version of the Lion [3] optimizer for greater compatibility with AdamW (algo. 1). The `sign` operation replaces the magnitude smoothing, explicitly controlling the ℓ_2 -norm of each update. `Additional scaling` keeps the hyperparameters η, λ comparable to AdamW.

Require: Learning rate η_t , weight decay λ , momentum β , Nesterov flag ν

```

1: Initialize: Time step  $t \leftarrow 0$ , parameter vector  $\theta_0$ , momentum vector  $m_0 \leftarrow 0$ 
2: while stopping criteria not met :
3:    $t \leftarrow t + 1$ 
4:    $g_t \leftarrow$  Mini-batch gradient w.r.t.  $\theta_{t-1}$ 
5:    $m_t \leftarrow \beta m_{t-1} + (1 - \beta)g_t$ 
6:   if Nesterov flag  $\nu$  is set :
7:      $\theta_t \leftarrow (1 - \eta_t \lambda) \theta_{t-1} - \eta_t \cdot \sqrt{(1 - \beta^2)^2 + \beta^4 \frac{1-\beta}{1+\beta}} \cdot \text{sign}(\beta m_t + (1 - \beta)g_t)$ 
8:   else:
9:      $\theta_t \leftarrow (1 - \eta_t \lambda) \theta_{t-1} - \eta_t \cdot \sqrt{\frac{1-\beta}{1+\beta}} \cdot \text{sign}(m_t)$ 

```

5 The Importance and Irregularity of the Angular Update Size

The effect of a weight vector $\mathbf{w}_t \in \mathbb{R}^C$ used in a dot product with some vector \mathbf{x} (e.g., in a neuron):

$$\langle \mathbf{w}_t, \mathbf{x} \rangle = \|\mathbf{w}_t\| \|\mathbf{x}\| \cos(\angle(\mathbf{w}_t, \mathbf{x})) \quad (2)$$

can be understood in terms of its magnitude $\|\mathbf{w}_t\|$ and direction $\mathbf{w}_t / \|\mathbf{w}_t\|$. The magnitude acts like a gain, scaling the outputs, whereas the direction determines which input representations \mathbf{x} the system responds to. The angular update size [38] of an update $\mathbf{w}_t \mapsto \mathbf{w}_{t+1}$ is defined as:

$$\angle(\mathbf{w}_{t+1}, \mathbf{w}_t) = \arccos \left(\frac{\langle \mathbf{w}_{t-1}, \mathbf{w}_{t+1} \rangle}{\|\mathbf{w}_t\| \|\mathbf{w}_t\|} \right) \quad (3)$$

and measures how fast the direction of \mathbf{w}_t changes during training, and thus its “preference” for \mathbf{x} .

With BatchNorm [15] and similar operations [1, 14, 42, 4], a network can become invariant to the magnitude of weight vectors like \mathbf{w}_t , such that only the direction matters and the vector is said to be *scale-invariant*. Weight Normalization [30] provides a good example of this, changing the system to:

$$\langle \mathbf{w}_t / \|\mathbf{w}_t\|, \mathbf{x} \rangle = \|\mathbf{x}\| \cos(\angle(\mathbf{w}_t, \mathbf{x})) \quad (4)$$

Note that although the system output is invariant to the magnitude $\|\mathbf{w}_t\|$, traditional optimizers are not. Scaling the value of a scale-invariant weight vector by a factor of $c > 0$, results in a gradient that is scaled by c^{-1} and curvature that is scaled by c^{-2} (see appx. B.2). For SGD this scales the angular update by c^{-2} and for Adam-like optimizers it is scaled by c^{-1} . With weight decay the magnitude of scale-invariant vectors trends towards a certain stable equilibrium value over time which also results in a specific expected angular update size as described by Wan et al. [38], Kosson et al. [19].

This has several important implications. Changing the initialization magnitude of scale-invariant weights will scale the angular updates over time for standard optimizers, resulting in effects similar to modifying the learning rate schedule. For initial weight magnitudes that are small compared to the equilibrium magnitude, the early angular updates will be large and these optimizers may benefit from learning rate warmup to counteract this. These effects also make the notion of “curvature” somewhat arbitrary as it can be scaled without changing the encoded function. Optimizers that specifically account for the weight magnitude would be invariant to these effects which may reduce the need for warmup from the traditional curvature perspective. Although standard transformers are not fully scale-invariant, many of the angular update insights still approximately hold for un-normalized weights [19].

In light of this, we modify LionA to better control the angular update size by making the updates to weight matrices proportional to their weight magnitude, resulting in algo. 3. We normalize the angular update size to match the equilibrium value, replacing weight decay with projections similar to Kosson et al. [19]. However, unlike their RVs, we make the angular updates proportional to the learning rate schedule which we found was necessary for good performance in our case. We also do not rely on additional exponential moving averages to control the angular update size, instead utilizing the fixed update size from the LionA optimizer. This is similar to the Adam scheme used by Karras et al. [17] with good results for diffusion models. No additional normalization operations or scaling factors are introduced, which we still find to result in decent performance.

Algorithm 3 LionAR: A rotational version of algo. 2 inspired by Kosson et al. [19]. The parameter vector is divided into sub-vectors $\theta = [\theta^{(1)}, \dots, \theta^{(P)}]$, each corresponding to either the weight vector of a neuron (e.g. a matrix row / a convolutional filter), or other parameters such as gains and biases. The **updates of neuronal weight vectors are scaled to be proportional to their magnitude** which is kept constant through **projections that replace weight decay**. Additional **hyperparameter adjustments** are made for compatibility with AdamW. The weight decay hyperparameter remains, fulfilling its primary role as a scaling factor for the relative updates of neurons [19].

Require: Learning rate η_t , weight decay λ , momentum β , Nesterov flag ν

- 1: **Initialize:** Time step $t \leftarrow 0$, parameter vector θ_0 , momentum vector $m_0 \leftarrow 0$
- 2: **while** stopping criteria not met :
- 3: $t \leftarrow t + 1$
- 4: $[g_t^{(1)}, \dots, g_t^{(P)}] \leftarrow$ Mini-batch gradient w.r.t. θ_{t-1} , divided into sub-vectors like θ
- 5: **for** $p \in \{1, \dots, P\}$:
- 6: $m_t^{(p)} \leftarrow \beta m_{t-1}^{(p)} + (1 - \beta) g_t^{(p)}$
- 7: **if** Nesterov flag ν is set :
- 8: $u_t^{(p)} \leftarrow \beta m_t^{(p)} + (1 - \beta) g_t^{(p)}$
- 9: $\gamma \leftarrow \sqrt{(1 - \beta^2)^2 + \beta^4 \frac{1-\beta}{1+\beta}}$ *# Nesterov momentum scaling factor*
- 10: **else:**
- 11: $u_t^{(p)} \leftarrow m_t^{(p)}$
- 12: $\gamma \leftarrow \sqrt{\frac{1-\beta}{1+\beta}}$ *# Heavy-ball momentum scaling factor*
- 13: **if** $\theta^{(p)} \in \mathbb{R}^C$ is a neuronal weight vector :
- 14: $\hat{\theta}_t^{(p)} \leftarrow \theta_{t-1}^{(p)} - \frac{\eta_t}{\max_{\tau}(\eta_{\tau}) \lambda} \cdot \sqrt{2 \max_{\tau}(\eta_{\tau}) \lambda} \cdot \gamma \cdot (\|\theta_0^{(p)}\| / \sqrt{C}) \cdot \text{sign}(u_t^{(p)})$
- 15: $\theta_t^{(p)} \leftarrow \hat{\theta}_t^{(p)} \cdot \|\theta_0^{(p)}\| / \|\hat{\theta}_t^{(p)}\|$ *# Reset the magnitude to the initial value*
- 16: **else:**
- 17: $\theta_t^{(p)} \leftarrow \theta_{t-1}^{(p)} - \eta_t \cdot \gamma \cdot \text{sign}(u_t^{(p)})$

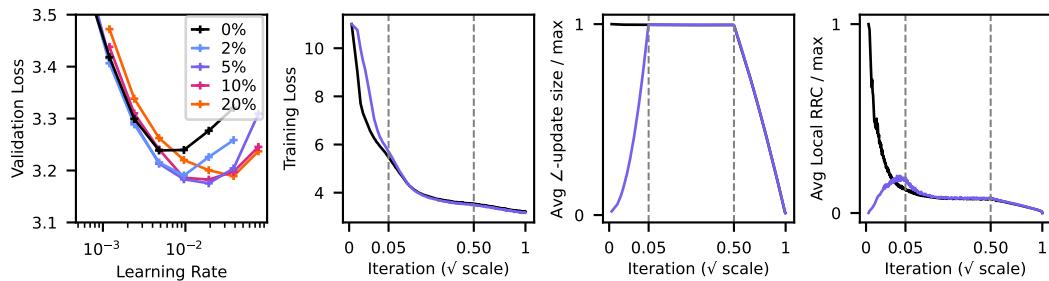


Figure 3: LionAR (algo. 3) reduces but does not fully eliminate the benefit of warmup. **Panel 1:** LionAR is more stable across learning rates and shows a reduced but still significant performance gap without warmup. **Panel 2:** Comparing the 0% and 5% warmup for learning rate $\approx 10^{-2}$ shows the warmup run overtaking early in training. **Panel 3:** LionAR precisely controls the angular update size throughout training. **Panel 4:** Despite fixed angular (and thus relative) updates in weight space, the relative change of the internal representations (see §6) is large initially without warmup.

Figure 3 repeats the GPT2 training sweep with LionAR. Consistent with the findings of Kosson et al. [19] we find that controlling the angular updates stabilizes training and decreases the benefit from warmup, but does not entirely eliminate it in this setting. Both the angular change and the ℓ_2 -norm are simple measures of the update magnitude in parameter space that do not account for the direction or other aspects of the update. In the next section we show how a fixed update size in parameter space can result in large changes to the internal representations of the network (a.k.a. features, activations etc), as shown in the final panel of fig. 3.

6 Early Gradient Alignment Results in Large Representation Changes

Our two approaches to measuring and controlling the update size in weight space failed to fully explain the need for warmup. As an alternative to the parameters, we can analyze changes in

the internal representations or activations of the neural network (feature learning). Although this is harder to analyze and control, it may ultimately be a better measure of the true impact of an update. A parameter update can only affect the network output, and hence the loss, by changing the representation of the network inputs at some layer. Large, sudden, changes in the representations could significantly affect the non-linearities, potentially causing lasting issues such as dead ReLUs or vanishing gradients from saturated sigmoids. This could in turn explain the lasting performance degradation observed without warmup.

A given parameter update will affect the representations of each distinct input sample differently. The gradients computed on these samples also generally differ, but can align to some extent. For a higher gradient alignment, the impact of a parameter update of a given magnitude on the representations will be larger than otherwise. We will analyze this for a dot product making up a single neuron:

$$\mathbf{y} = \mathbf{w}^\top \mathbf{X} = [y_1, \dots, y_B]^\top = [\langle \mathbf{w}, \mathbf{x}_1 \rangle, \dots, \langle \mathbf{w}, \mathbf{x}_B \rangle]^\top \quad (5)$$

where $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_B] \in \mathbb{R}^{C \times B}$ are the C -dimensional representations of a random mini-batch of B inputs that is fed into the neuron, $\mathbf{w} \in \mathbb{R}^C$ is the weight vector, and $\mathbf{y} \in \mathbb{R}^B$ is a batch of outputs. For a weight update $\mathbf{w} \mapsto \mathbf{w} + \Delta\mathbf{w}$, we aim to quantify the size of the output change $\Delta\mathbf{y} = \Delta\mathbf{w}^\top \mathbf{X}$ computed on the same inputs. We focus on the *Relative Representation Change* (RRC):

$$\frac{\|\Delta\mathbf{y}\|}{\|\mathbf{y}\|} = \frac{\|\Delta\mathbf{w}^\top \mathbf{X}\|}{\|\mathbf{w}^\top \mathbf{X}\|} \quad (6)$$

similar to the angular weight updates, as the sensitivity to the absolute change $\|\Delta\mathbf{y}\|$ can be unclear due to normalization or other scaling operations. Note that this is a measure of a *local change*, not accounting for changes in the inputs \mathbf{X} from updates to preceding layers (*global change*).

6.1 Normalized Gradient Descent

We will focus our analysis on the relatively tractable case of normalized gradient descent with updates:

$$\Delta\mathbf{w} = -\eta \frac{\mathbf{g}}{\sqrt{\mathbb{E}[\|\mathbf{g}\|^2]}}, \quad \mathbf{g} = \frac{1}{B} \sum_{b=1}^B \mathbf{g}_b \quad (7)$$

where \mathbf{g}_b is the gradient of some loss w.r.t. \mathbf{w} for the b -th element of the mini-batch. We will use the following definitions, properties, lemmas and assumptions for this system (see appx. B.4 for details):

- D1: We define $\mathbf{g}_b =: \bar{\mathbf{g}} + \tilde{\mathbf{g}}_b$ where $\bar{\mathbf{g}} = \mathbb{E}[\mathbf{g}]$ and $\tilde{\mathbf{g}}_b$ is the difference with $\mathbb{E}[\tilde{\mathbf{g}}_b] = \mathbf{0}$.
- D2: We define $\varphi := \mathbb{E}[\|\bar{\mathbf{g}}\|^2]/\mathbb{E}[\|\tilde{\mathbf{g}}_b\|^2]$ as the Signal-to-Noise Ratio (SNR) of the gradient.
- P1: For a neuron, $\mathbf{g}_b \parallel \mathbf{x}_b$, and hence $\mathbf{x}_b = \text{sign}(\langle \mathbf{x}_b, \mathbf{g}_b \rangle) \cdot (\|\mathbf{x}_b\|/\|\mathbf{g}_b\|) \cdot (\bar{\mathbf{g}} + \tilde{\mathbf{g}}_b)$.
- L1: Consider two independent random vectors $\mathbf{a} \in \mathbb{R}^C$ and $\mathbf{b} \in \mathbb{R}^C$, whose elements are independent and identically distributed (IID). If at least one of the vectors has a zero-mean distribution, then the expected value of the squared inner product of \mathbf{a} and \mathbf{b} is given by $\mathbb{E}[\langle \mathbf{a}, \mathbf{b} \rangle^2] = \mathbb{E}[\|\mathbf{a}\|^2]\mathbb{E}[\|\mathbf{b}\|^2]/C$.
- A1: We assume the following vector pairs satisfy L1: $(\mathbf{x}_i, \tilde{\mathbf{g}}_b)$ when $i \neq b$, $(\bar{\mathbf{g}}, \tilde{\mathbf{g}}_b)$ and $(\mathbf{w}, \mathbf{x}_b)$.

This allows us to compute the expected square relative representation change (detailed in appx. B.4):

$$\begin{aligned} \frac{\mathbb{E}[(\Delta y_b)^2]}{\mathbb{E}[y_b^2]} &= \frac{\eta^2 C}{B^2 \|\mathbf{w}\|^2} \frac{1}{\mathbb{E}[\|\mathbf{g}\|^2]} \left(\mathbb{E}[\|\mathbf{g}_b\|^2] + \frac{B-1}{C} \mathbb{E}[\|\tilde{\mathbf{g}}_b\|^2] \right. \\ &\quad \left. + \frac{(B-1)^2}{\mathbb{E}[\|\mathbf{g}_b\|^2]} \left(\|\bar{\mathbf{g}}\|^4 + \frac{\|\bar{\mathbf{g}}\|^2 \mathbb{E}[\|\tilde{\mathbf{g}}_b\|^2]}{C} \right) + 2(B-1)\|\bar{\mathbf{g}}\|^2 \right) \quad (8) \end{aligned}$$

$$= \frac{\eta^2 C}{B^2 \|\mathbf{w}\|^2} \frac{1}{\varphi + \frac{1}{B}} \left((\varphi + 1) + \frac{B-1}{C} + \left(\frac{(B-1)^2 \varphi}{\varphi + 1} \left(\varphi + \frac{1}{C} \right) + 2(B-1)\varphi \right) \right) \quad (9)$$

The expected relative change in the output for a given sample can be broken down into three sources, the contribution of the sample itself (first term), random interference from the “noise” $\tilde{\mathbf{g}}_i$ of other samples (second term), and finally amplification of the common mean component $\bar{\mathbf{g}}$ (third term).

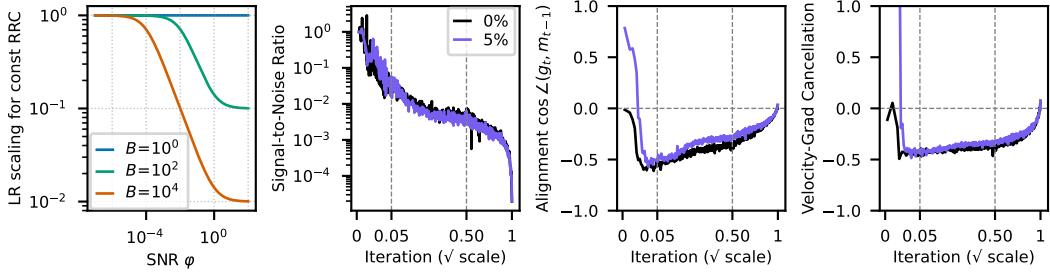


Figure 4: Equation (9) predicts that the learning rate needs to be downscaled for higher signal to noise ratios (φ) to keep the relative representation change constant. Larger batch sizes are affected more, with scaling becoming significant when $\varphi > B^{-1}$. **Panel 2:** Measurements of the SNR for the two highlighted runs in fig. 3. Note the SNR starts very high but is also remains large in comparison to our $B = 480$ for almost all of training. **Panel 3:** The gradient is strongly oppositely aligned with the momentum vector for most of training (shown for an example layer). **Panel 4:** Projecting the momentum component of the updates onto the gradient component shows that this results in the momentum vector ‘‘cancelling’’ roughly half the gradient on average.

The RRC expression provides many interesting insights. In the case of large input dimension $C \rightarrow \infty$ and small SNR $\varphi \approx 0$, keeping the RRC constant for different batch sizes involves scaling the learning rate $\eta \propto \sqrt{B}$, as suggested by Malladi et al. [27] for Adam. When the SNR φ is some finite and value and C is still large, this scaling rule instead starts to break down around $B = 1/\varphi$, matching the predicted critical batch size of e.g. McCandlish et al. [28]. The role of the dimension C in the expression is curious, suggesting that narrower layers experience larger changes due to random inference from other samples in a given batch. The C in the leading factor also suggests that the angular updates can be smaller for a larger input dimension, similar to what is proposed in μ -parameterization [44, 45]. Most importantly, **this expression shows that if the SNR changes throughout training the learning rate needs to be adjusted to keep the RRC constant. In particular, with large batch sizes, a high initial SNR results in large representation changes which warmup can help prevent.** The first panel of fig. 4 shows how eq. (9) predicts we should downscale the learning rate for different batch sizes and SNRs, assuming we originally scaled the learning rate $\eta \propto \sqrt{B}$ and that C is large. The second panel confirms that the SNR indeed starts out large, suggesting lower initial learning rates are needed, i.e. warmup.

In the first panel of fig. 5, we show the results of adding a term that scales the update size as predicted by eq. (9). **This acts similar to an automatic warmup based on online measurements of the SNR which we obtain from the gradient accumulation of micro-batches.** Although this helps close the gap between warmup and no-warmup, the overall performance is slightly worse. One potential issue is that our batch size of 480 is quite large compared to the measured SNR, exceeding the critical batch size estimation throughout most of training. This results in a scaling of the step size throughout training, which distorts the decay phase. It also requires large learning rate values to counteract the scaling, which may destabilize the training of non-matrix weights like gains. We increase the weight decay by a factor of 32 \times to try to increase the angular updates relative to gains in order to compensate, but this value was not tuned and is unlikely to be optimal. We believe approaches that aim to directly control the RRC are a promising direction but require further work to be practical.

6.2 The Role of Momentum

Momentum is believed to be a key enabler of optimization with larger batch sizes [32, 39, 31]. However, it is unclear how it should change predictions for the critical batch size or relative representation change. Momentum spreads the application of a gradient out over multiple steps which tends to make each update smaller, especially for a random walk, which is reflected in our scaling coefficients in algo. 2 and 3. The smaller updates are counteracted by an increased correlation in their direction, which can result in similar ‘‘long term’’ changes from each gradient sample, especially for simpler methods like SGD that don’t normalize the step size. In the last two panels of fig. 4 we observe that in our setup the gradient and momentum are negatively correlated, counteracting each other. We find momentum crucial for performant training, panel 2 of fig. 5 shows significant degradation without it.

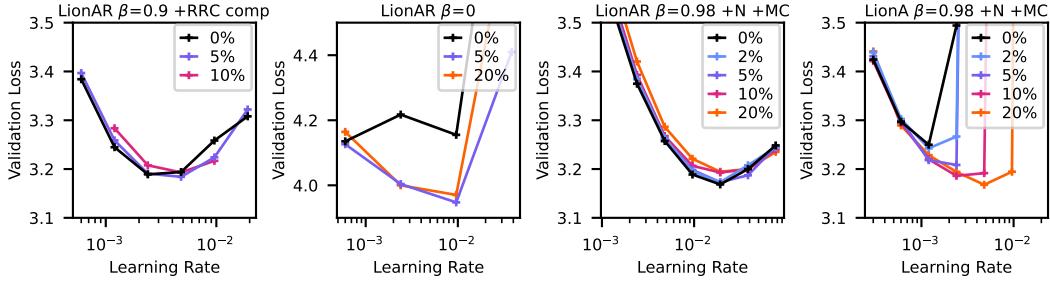


Figure 5: Panel 1: LionAR with a correction factor for the RRC based on eq. (9) does not benefit from a warmup. **Panel 2:** LionAR training without momentum results in drastically lower performance. **Panel 3:** In LionAR with increased momentum $\beta = 0.98$, Nesterov momentum and an inverse bias correction for early momentum, no warmup performs best. **Panel 4:** The same does not apply to LionA, suggesting that these changes are not sufficient without controlling the angular updates.

We believe the smaller update sizes for momentum combined with the potential for later gradients to counteract earlier gradients during their application over time, can help stabilize training. An otherwise large relative representation change is spread out over multiple steps and counteracted by later gradients. Higher values of momentum should amplify these effects. Looking at the total contribution of each gradient also implies that **with momentum early updates should be smaller when measured in parameter space, otherwise the relative representation change for those samples is too large**. This is equivalent to entirely removing the β_1 bias correction in AdamW (algo. 1, line 7), or introducing *an inverse bias correction* in Lion like algorithms (see appx. B.1 for details). Higher β values should help amplify the stabilization effects of momentum. **In fig. 5 we find that at higher momentum values LionAR no longer benefits from warmup unlike LionA which still needs it.** These experiments use Nesterov momentum and the additional inverse bias correction, though these adjustments offer only minor improvements compared to higher momentum.

7 The Detrimental Effects of Large Updates

In appx. A we empirically investigate potential causes for the lasting performance degradation from large initial updates for a small ResNet model. We find that the performance impact best correlates with the number of dead ReLUs and is improved by the use of leaky-ReLUs, which fits well with our perspective of large changes in the internal representations. We also investigated whether overfitting to initial training samples or the correlation between weight vectors of different neurons could explain the necessity for warmup, but these factors did not show a significant impact.

8 The Role of Larger Batch Sizes

Warmup is often used with larger batch sizes in particular, for example in the setting where Goyal et al. [10] first proposed using linear warmup. Although this was for SGD, we expect the need for warmup to be amplified at larger batch sizes for two of the reasons we identified. The first is that larger batch sizes are more likely to exceed the critical batch size early in training. The second is the size of early angular updates. As shown by Kosson et al. [19], the equilibrium weight magnitude depends on the learning rate and weight decay value. Common hyperparameter scaling rules for a modified batch size only change the learning rate but not the weight decay, which shifts the equilibrium magnitude. The smaller the initialization magnitude is compared to the equilibrium magnitude, the larger the early angular updates will be relative to their steady state value, potentially necessitating warmup.

9 Limitations

Our main experiments focus on a single network which may not be broad enough to generalize to a wide range of networks. In appx. C.2 we experiments with an additional dataset and architecture but the scale of the experiments is still limited and they cover a limited range of hyperparameters. We believe we identify real factors that contribute to the need for warmup, but these may not be the only

ones across a broader range of settings. Similarly, the promising results for reducing or eliminating the warmup with higher momentum values or the relative representation correction would benefit from broader validation.

10 Conclusion

In this work we explored how the size of the updates $\Delta \mathbf{w} = \eta \mathbf{u}$ impacts the need for learning rate warmup. We showed that \mathbf{u} can be large initially when measured in terms of its ℓ_2 -norm (§4), the resulting directional change in \mathbf{w} (angular update, §5), as well as the resulting change in the internal representations of the network (relative representation change, §6). We argued that small initial values of the learning rate η are beneficial to counteract large values of \mathbf{u} , i.e. that a learning rate warmup simply keeps some notion of the overall “effective” update size reasonable. We showed this experimentally rather than theoretically by modifying the optimizers to normalize the size of \mathbf{u} based on each metric and measuring how these changes affected the benefit of using learning rate warmup.

The two weight-based measures of the update size, the ℓ_2 -norm and angular update did not fully account for the need for warmup. However, quantifying the update size in terms of the relative change in neural representations shows potential. This measure is closely linked to the angular update size but accounts for changes in the signal characteristics of the gradient, which can vary significantly throughout training. Effectively controlling neural representation changes is a challenging task we leave for future work, but our initial attempts show encouraging results in reducing the need for a manually configured warmup. We also highlighted the importance of high momentum for warmup; when combined with angular update control and an inverse bias correction, it may enable efficient warmup-free training. Overall, our work provides new insights into the benefit of learning rate warmup with modern optimizers beyond SGD and suggests potential directions for eliminating it.

Although we present new methods we consider promising, we still recommend the use of a short warmup in practice. Fully eliminating it seems to require significant modifications that also need further validation across additional settings. However, we hope to have provided the reader with a new perspective and simple intuition for why warmup is beneficial for training. We also hope our work inspires further exploration of how learning should be controlled and scheduled in neural network training. In particular, it seems that the learning rate in current optimizers does not really control the “rate of learning”, making learning rate schedules and the use of warmup highly arbitrary.

References

- [1] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016. 5
- [2] Jeremy Bernstein, Arash Vahdat, Yisong Yue, and Ming-Yu Liu. On the distance between two neural networks and the stability of learning. *Advances in Neural Information Processing Systems*, 33:21370–21381, 2020. arXiv:2002.03432. 2
- [3] Xiangning Chen, Chen Liang, Da Huang, Esteban Real, Kaiyuan Wang, Hieu Pham, Xuanyi Dong, Thang Luong, Cho-Jui Hsieh, Yifeng Lu, and Quoc V Le. Symbolic discovery of optimization algorithms. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=ne6zeqLFCZ>. arXiv:2302.06675. 1, 4, 5
- [4] Vitaliy Chiley, Ilya Sharapov, Atli Kosson, Urs Koster, Ryan Reece, Sofia Samaniego de la Fuente, Vishal Subbiah, and Michael James. Online normalization for training neural networks. *Advances in Neural Information Processing Systems*, 32, 2019. arXiv:1905.05894. 5
- [5] Katie E Everett, Lechao Xiao, Mitchell Wortsman, Alexander A Alemi, Roman Novak, Peter J Liu, Izzeddin Gur, Jascha Sohl-Dickstein, Leslie Pack Kaelbling, Jaehoon Lee, and Jeffrey Pennington. Scaling exponents across parameterizations and optimizers. In *Forty-first International Conference on Machine Learning*, 2024. URL <https://openreview.net/forum?id=0ksNeD1SJT>. arXiv:2407.05872. 18
- [6] Jingwen Fu, Bohan Wang, Huishuai Zhang, Zhizheng Zhang, Wei Chen, and Nanning Zheng. When and why momentum accelerates sgd: An empirical study. *arXiv preprint arXiv:2306.09000*, 2023. 2

[7] Justin Gilmer, Behrooz Ghorbani, Ankush Garg, Sneha Kudugunta, Behnam Neyshabur, David Cardoze, George Edward Dahl, Zachary Nado, and Orhan Firat. A loss curvature perspective on training instabilities of deep learning models. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=0cKMT-36vUs>. arXiv:2110.04369. 2

[8] Aaron Gokaslan and Vanya Cohen. Openwebtext corpus. <http://Skylion007.github.io/OpenWebTextCorpus>, 2019. 2, 22

[9] Akhilesh Gotmare, Nitish Shirish Keskar, Caiming Xiong, and Richard Socher. A closer look at deep learning heuristics: Learning rate restarts, warmup and distillation. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=r14E0sCqKX>. arXiv:1810.13243. 2

[10] Priya Goyal, Piotr Dollár, Ross Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He. Accurate, large minibatch sgd: Training imagenet in 1 hour. *arXiv preprint arXiv:1706.02677*, 2017. 1, 2, 9

[11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. arXiv:1512.03385. 1, 2, 14

[12] Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, et al. Training compute-optimal large language models. *arXiv preprint arXiv:2203.15556*, 2022. URL <https://arxiv.org/abs/2203.15556>. 2

[13] Xiao Shi Huang, Felipe Perez, Jimmy Ba, and Maksims Volkovs. Improving transformer optimization through better initialization. In *International Conference on Machine Learning*, pages 4475–4483. PMLR, 2020. URL <https://proceedings.mlr.press/v119/huang20f.html>. 2

[14] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *Proceedings of the IEEE international conference on computer vision*, pages 1501–1510, 2017. arXiv:1703.06868. 5

[15] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pages 448–456. pmlr, 2015. arXiv:1502.03167. 5

[16] Andrej Karpathy. nanogpt. <https://github.com/karpathy/nanoGPT/>, 2023. 3

[17] Tero Karras, Miika Aittala, Jaakko Lehtinen, Janne Hellsten, Timo Aila, and Samuli Laine. Analyzing and improving the training dynamics of diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 24174–24184, 2024. arXiv:2312.02696. 5

[18] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, San Diego, CA, USA, 2015. arXiv:1412.6980. 2

[19] Atli Kosson, Bettina Messmer, and Martin Jaggi. Rotational equilibrium: How weight decay balances learning across neural networks. In *Forty-first International Conference on Machine Learning*, 2024. URL <https://openreview.net/forum?id=MQirNNU2pC>. arXiv:2305.17212. 2, 3, 5, 6, 9, 14, 16

[20] Alex Krizhevsky. Learning multiple layers of features from tiny images. *self-published*, 2009. URL <https://www.cs.toronto.edu/~kriz/learning-features-2009-TR.pdf>. 14

[21] Zhiyuan Li, Kaifeng Lyu, and Sanjeev Arora. Reconciling modern deep learning with traditional optimization analyses: The intrinsic learning rate. *Advances in Neural Information Processing Systems*, 33:14544–14555, 2020. arXiv:2010.02916. 2

[22] Zhiyuan Li, Sadhika Malladi, and Sanjeev Arora. On the validity of modeling SGD with stochastic differential equations (SDEs). In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021. URL https://openreview.net/forum?id=goEdyJ_nVQI. arXiv:2102.12470. 2

[23] Liyuan Liu, Haoming Jiang, Pengcheng He, Weizhu Chen, Xiaodong Liu, Jianfeng Gao, and Jiawei Han. On the variance of the adaptive learning rate and beyond. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=rkgz2aEKDr>. arXiv:1908.03265. 2, 21

[24] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=Bkg6RiCqY7>. arXiv:1711.05101. 1, 2, 3, 21

[25] Kaifeng Lyu, Zhiyuan Li, and Sanjeev Arora. Understanding the generalization benefit of normalization layers: Sharpness reduction. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022. URL <https://openreview.net/forum?id=xp5V0BxTxZ>. arXiv:2206.07085. 16

[26] Jerry Ma and Denis Yarats. On the adequacy of untuned warmup for adaptive optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 8828–8836, 2021. arXiv:1910.04209. 2, 21

[27] Sadhika Malladi, Kaifeng Lyu, Abhishek Panigrahi, and Sanjeev Arora. On the SDEs and scaling rules for adaptive gradient algorithms. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022. URL <https://openreview.net/forum?id=F2mhzjHkQP>. arXiv:2205.10287. 2, 8

[28] Sam McCandlish, Jared Kaplan, Dario Amodei, and OpenAI Dota Team. An empirical model of large-batch training. *arXiv preprint arXiv:1812.06162*, 2018. 2, 8

[29] Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. *self-published*, 2019. URL https://d4mucfpksyv.w.cdncloudfront.net/better-language-models/language_models_are_unsupervised_multitask_learners.pdf. 1, 2

[30] Tim Salimans and Durk P Kingma. Weight normalization: A simple reparameterization to accelerate training of deep neural networks. *Advances in neural information processing systems*, 29, 2016. arXiv:1602.07868. 5

[31] Christopher J Shallue, Jaehoon Lee, Joseph Antognini, Jascha Sohl-Dickstein, Roy Frostig, and George E Dahl. Measuring the effects of data parallelism on neural network training. *Journal of Machine Learning Research*, 20(112):1–49, 2019. arXiv:1811.03600. 8

[32] Samuel Smith, Erich Elsen, and Soham De. On the generalization benefit of noise in stochastic gradient descent. In *International Conference on Machine Learning*, pages 9058–9067. PMLR, 2020. arXiv:2006.15081. 2, 8

[33] Daria Soboleva, Faisal Al-Khateeb, Robert Myers, Jacob R Steeves, Joel Hestness, and Nolan Dey. Slimpajama: A 627b token cleaned and deduplicated version of redpajama. <https://cerebras.ai/blog/slimpjama-a-627b-token-cleaned-and-deduplicated-version-of-redpajama>, 2023. 21, 22

[34] Sebastian Stich, Amirkeivan Mohtashami, and Martin Jaggi. Critical parameters for scalable distributed learning with large batches and asynchronous updates. In *International Conference on Artificial Intelligence and Statistics*, pages 4042–4050. PMLR, 2021. arXiv:2103.02351. 2

[35] Ilya Sutskever, James Martens, George Dahl, and Geoffrey Hinton. On the importance of initialization and momentum in deep learning. In *International conference on machine learning*, pages 1139–1147. PMLR, 2013. URL <https://proceedings.mlr.press/v28/sutskever13.html>. 2

[36] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwäl Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023. 22

[37] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. arXiv:1706.03762. 1, 2

[38] Ruosi Wan, Zhanxing Zhu, Xiangyu Zhang, and Jian Sun. Spherical motion dynamics: Learning dynamics of normalized neural network using sgd and weight decay. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 6380–6391. Curran Associates, Inc., 2021. URL <https://proceedings.neurips.cc/paper/2021/file/326a8c055c0d04f5b06544665d8bb3ea-Paper.pdf>. arXiv:2006.08419. 4, 5

[39] Runzhe Wang, Sadhika Malladi, Tianhao Wang, Kaifeng Lyu, and Zhiyuan Li. The marginal value of momentum for small learning rate SGD. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=3JjJezzVkt>. arXiv:2307.15196. 8

[40] Ross Wightman. Pytorch image models. <https://github.com/rwightman/pytorch-image-models>, 2019. 14

[41] Mitchell Wortsman, Peter J Liu, Lechao Xiao, Katie Everett, Alex Alemi, Ben Adlam, John D Co-Reyes, Izzeddin Gur, Abhishek Kumar, Roman Novak, et al. Small-scale proxies for large-scale transformer training instabilities. *arXiv preprint arXiv:2309.14322*, 2023. URL <https://arxiv.org/abs/2309.14322>. 2

[42] Yuxin Wu and Kaiming He. Group normalization. In *Proceedings of the European conference on computer vision (ECCV)*, pages 3–19, 2018. arXiv:1803.08494. 5

[43] Ruibin Xiong, Yunchang Yang, Di He, Kai Zheng, Shuxin Zheng, Chen Xing, Huishuai Zhang, Yanyan Lan, Liwei Wang, and Tieyan Liu. On layer normalization in the transformer architecture. In *International Conference on Machine Learning*, pages 10524–10533. PMLR, 2020. arXiv:2002.04745. 2

[44] Greg Yang, Edward J Hu, Igor Babuschkin, Szymon Sidor, Xiaodong Liu, David Farhi, Nick Ryder, Jakub Pachocki, Weizhu Chen, and Jianfeng Gao. Tuning large neural networks via zero-shot hyperparameter transfer. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021. URL <https://openreview.net/forum?id=Bx6qKuBM2AD>. arXiv:2203.03466. 8

[45] Greg Yang, James B Simon, and Jeremy Bernstein. A spectral condition for feature learning. *arXiv preprint arXiv:2310.17813*, 2023. 8

[46] Dong Yin, Ashwin Pananjady, Max Lam, Dimitris Papailiopoulos, Kannan Ramchandran, and Peter Bartlett. Gradient diversity: a key ingredient for scalable distributed learning. In *International Conference on Artificial Intelligence and Statistics*, pages 1998–2007. PMLR, 2018. arXiv:1706.05699. 2

[47] Yang You, Igor Gitman, and Boris Ginsburg. Large batch training of convolutional networks. *arXiv preprint arXiv:1708.03888*, 2017. 2

[48] Yang You, Jing Li, Sashank Reddi, Jonathan Hseu, Sanjiv Kumar, Srinadh Bhojanapalli, Xiaodan Song, James Demmel, Kurt Keutzer, and Cho-Jui Hsieh. Large batch optimization for deep learning: Training bert in 76 minutes. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=Syx4wnEtvH>. arXiv:1904.00962. 2

[49] Guodong Zhang, Lala Li, Zachary Nado, James Martens, Sushant Sachdeva, George Dahl, Chris Shallue, and Roger B Grosse. Which algorithmic choices matter at which batch sizes? Insights from a noisy quadratic model. *Advances in Neural Information Processing Systems*, 32, 2019. arXiv:1907.04164. 2

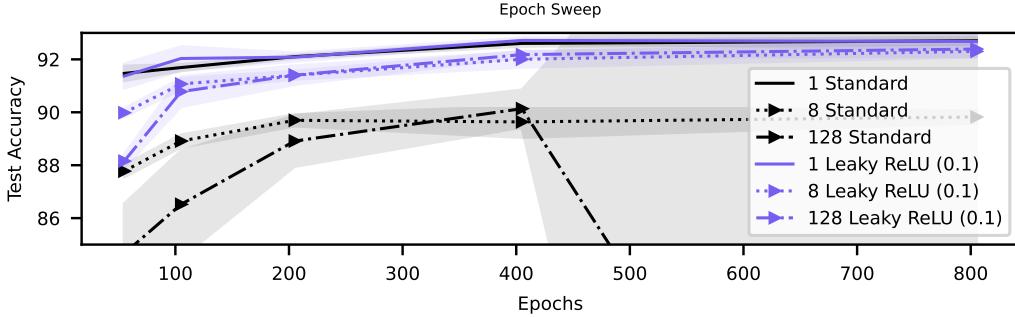


Figure 6: The performance gap caused by large initial updates persists despite extended training (800 epochs) in a standard ResNet-20. We investigate the influence of network non-linearities by comparing two training methods while scaling update sizes during the 5 epoch initial phase by factors of 1, 8, and 128: *Standard* (S), which employs traditional ReLU activations, and *Leaky ReLU*, which replaces ReLUs with Leaky ReLUs using a scaling factor of $\alpha = 0.1$. We observe that training with *Leaky ReLU* results in smaller performance degradation from large initial updates, suggesting that the non-linearities in the network might substantially impact the observed performance degradation.

A The Detrimental Effects of Large Updates

To investigate the effects of large updates at the beginning of training, we conducted controlled experiments on a ResNet-20 model on CIFAR-10 [20] due to resource constraints. We controlled the average angular update throughout training using the rotational optimizer variant of AdamW proposed by Kosson et al. [19]. For the initial phase of training, 5 epochs, we either use a standard learning rate of 0.05 or amplify it by a factor of 8 or 128. This results in very large updates, scaling the rotation by either $\sqrt{8}$ or $\sqrt{128}$. For all experiments, we used a weight decay of 0.01, $\beta_1 = 0.9$, $\beta_2 = 0.999$, 5 epoch initial phase, and trained for 205 epochs in total with a cosine schedule unless otherwise specified. The data was pre-processed by normalizing it with a mean of (0.4914, 0.4822, 0.4465) and a standard deviation of (0.2023, 0.1994, 0.2010) and applying simple data augmentation techniques as described by He et al. [11]. To run the experiment, we used the codebase from Wightman [40] and extended the utilities from Kosson et al. [19].

As shown in fig. 6, the performance of standard training does not recover when large updates are used at the beginning of training, even when the training time is extended to four times the normal duration for ReLU networks. This suggest that large, controlled, initial updates can result in a permanent performance degradation, similar to what we observe without warmup in advanced settings. The impact is much smaller when replacing ReLUs with leaky ReLUs, suggesting that the non-linearities in the network might substantially contribute to the performance degradation.

In fig. 7 we measure the fraction of dead ReLUs directly across different settings and scaling factors. We find that large initial updates do indeed lead to a large number of permanently dead units and that the final accuracy suffers when this is the case. This effect can be mitigated by freezing the biases at the beginning of training, as shown in the table in fig. 7. We also observe that replacing the actual gradients with random gradients has a much smaller impact, suggesting that the direction of the updates also matters for the degradation, not only their size.

Interestingly, we did not find a connection to overfitting to a small number of samples at the beginning of training. The performance of 92.1 can be recovered in this case. Additionally, we explored stable rank measurements as a potential factor but did not find a significant relation, as detailed in fig. 8.

B Additional Mathematical and Technical Details

B.1 The magnitude of the Momentum Vector

Let's assume a scalar gradient g_t (e.g. for some coordinate) that is a random variable that is independent across time and has a zero mean distribution that does not change across time, i.e. $E[g_t] = 0$

Method	Scale 1	Scale 8	Scale 128
Standard	92.13 \pm 0.2	89.48 \pm 0.2	88.22 \pm 1.6
Standard frozen bias	92.30 \pm 0.3	92.08 \pm 0.3	92.30 \pm 0.2
Random	92.05 \pm 0.3	91.74 \pm 0.2	89.54 \pm 0.3
Random frozen bias	92.27 \pm 0.2	92.12 \pm 0.4	92.20 \pm 0.1
Leaky ReLU	92.16 \pm 0.3	91.48 \pm 0.3	91.82 \pm 0.4
Leaky ReLU frozen bias	92.46 \pm 0.2	92.49 \pm 0.1	92.35 \pm 0.2

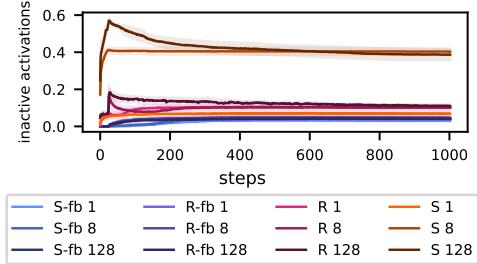


Figure 7: Comparison of the performance (final test accuracy) and fraction of dead ReLUs (inactive activations) across different settings. The learning rate in the initial phase of 5 epochs is scaled by a factor of either 1, 8, or 128. *Standard (S)* denotes normal training, while *Frozen Biases (fb)* involves freezing the biases at the onset of training. The *Random (R)* approach employs random gradient directions at the start of training, and *Leaky ReLU* replaces the ReLUs in standard models with Leaky ReLUs using a scaling factor of $\alpha = 0.1$. We observe a notable correspondence between large initial updates, higher ratios of dead ReLUs in ResNet-20, and performance degradation.

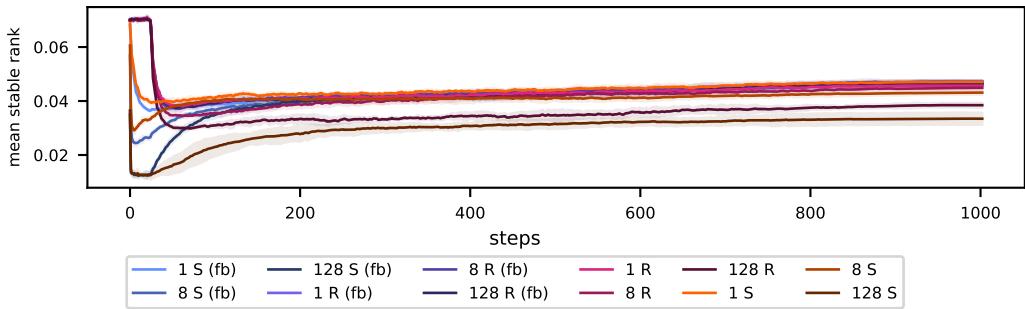


Figure 8: Impact of varying update sizes during the warmup phase on the stable rank of a standard ResNet-20. The learning rate in the initial phase of 5 epochs is scaled by a factor of either 1, 8, or 128. We evaluate the effects of across different training configurations: *Standard (S)* denotes normal training; *Frozen Biases (fb)* involves freezing the biases at the onset of training; and *Random (R)* employs random gradient directions at the start of training. The stable rank remains largely consistent across these methods, except when using extremely large updates—specifically, scaling the learning rate by a factor of 128 without freezing biases—which leads to noticeable variations in the rank.

and $\mathbb{E}[g_t^2] = \sigma^2$. For standard **heavyball-momentum** m_t , with $m_0 = 0$ and coefficient β (equivalent to β_1 for Adam) we have:

$$\mathbb{E}[m_t^2] = \mathbb{E}[(\beta m_{t-1} + (1 - \beta)g_t)^2] \quad (10)$$

$$= \mathbb{E} \left[\left((1 - \beta) \sum_{i=0}^{t-1} \beta^i g_{t-i} \right)^2 \right] \quad (11)$$

$$= \mathbb{E} \left[(1 - \beta)^2 \sum_{i=0}^{t-1} \beta^{2i} g_{t-i}^2 + (1 - \beta)^2 \sum_{j=0}^{t-1} \sum_{\substack{k=0 \\ k \neq j}}^{t-1} \beta^{2t-j-k} g_{t-j} g_{g-k} \right] \quad (12)$$

$$= (1 - \beta)^2 \sum_{i=0}^{t-1} \beta^{2i} \mathbb{E}[g_{t-i}^2] + (1 - \beta)^2 \sum_{j=0}^{t-1} \sum_{\substack{k=0 \\ k \neq j}}^{t-1} \beta^{2t-j-k} \mathbb{E}[g_{t-j}] \mathbb{E}[g_{g-k}] \quad (13)$$

$$= (1 - \beta)^2 \sum_{i=0}^{t-1} \beta^{2i} \sigma^2 + 0 \quad (14)$$

$$= (1 - \beta)^2 \frac{1 - \beta^{2t}}{1 - \beta^2} \sigma^2 \quad (15)$$

$$= (1 - \beta)^2 \frac{1 - \beta^{2t}}{(1 - \beta)(1 + \beta)} \sigma^2 \quad (16)$$

$$= (1 - \beta^{2t}) \frac{1 - \beta}{1 + \beta} \sigma^2 \quad (17)$$

In the limit $t \rightarrow \infty$ we have $(1 - \beta^{2t}) \rightarrow 1$. We can derive the size of the second-moment v_t in AdamW in an analogous way, obtaining $\mathbb{E}[v_t] = (1 - \beta_2^t)\sigma^2$. For a random walk, the update size of Adam is scaled in a similar way. Since the update size of Lion is fixed and does not depend on β , we scale the update size to match that of AdamW for a random walk in a steady state, i.e. by $\gamma = \sqrt{\frac{1-\beta}{1+\beta}}$ as seen in algo. 2.

For **Nesterov momentum**, the update is modified to use:

$$u_t = \beta m_t + (1 - \beta)g_t \quad (18)$$

$$= \beta(\beta m_{t-1} + (1 - \beta)g_t) + (1 - \beta)g_t \quad (19)$$

$$= \beta^2 m_{t-1} + (1 - \beta)(1 + \beta)g_t \quad (20)$$

Note that m_{t-1} and g_t are independent and zero-mean, allowing us to use the previous result for:

$$\mathbb{E}[u_t^2] = \mathbb{E}\left[\left(\beta^2 m_{t-1} + (1 - \beta)(1 + \beta)g_t\right)^2\right] \quad (21)$$

$$= \beta^4 \mathbb{E}[m_{t-1}^2] + (1 - \beta^2)^2 \mathbb{E}[g_t^2] \quad (22)$$

$$= \beta^4(1 - \beta^{2t-2}) \frac{1 - \beta}{1 + \beta} \sigma^2 + (1 - \beta^2)^2 \sigma^2 \quad (23)$$

In the limit $t \rightarrow \infty$ this gives the Nesterov scaling factor used in LionA (algo. 2) to ensure that the update size corresponds to that of AdamW using an analogous Nesterov update.

Inverse bias correction for momentum. Adam uses a bias correction to attempt to fix the update size over time. This scales early updates resulting in the contributions of the corresponding gradients being amplified. The relative representation change for those samples is increased as a result, similar to applying the same update multiple times. Removing the β_1 bias correction from AdamW removes this effect. LionA and LionAR similarly scale the update size, making it constant. We can counteract this by changing our scaling factors to use the time varying expressions based on the derivations above. Note however, that this assumed the gradients were uncorrelated so it only approximately undoes the scaling effect for real values with arbitrary alignment of successive gradients. To summarize, the inverse bias correction for momentum changes the momentum scaling factors (γ in algo. 3) to vary over time:

$$\text{Nesterov: } \gamma_t = \sqrt{(1 - \beta^2)^2 + (1 - \beta^{2t-2})\beta^4 \frac{1 - \beta}{1 + \beta}} \quad (24)$$

$$\text{Heavy-ball: } \gamma_t = \sqrt{(1 - \beta^{2t}) \frac{1 - \beta}{1 + \beta}} \quad (25)$$

B.2 Properties of Scale Invariance

Derivations for the gradient magnitude and curvature can be found in existing works, for example Lyu et al. [25]. When a scale invariant weight is scaled by a factor $c > 0$, the gradient is scaled by c^{-1} which scales the ratio of the gradient norm and weight norm, and therefore the angular updates, by c^{-2} . For normalized optimizers like Adam and Lion, where the update norm is not affected by the gradient magnitude, this factor is decreased to c^{-1} .

B.3 The Angular Update Size in LionAR

The scaling factor for the angular update size in algo. 3 is adopted directly from the AdamW value derived by Kosson et al. [19]. Since the Nesterov momentum does not change the total contribution of each gradient it does not affect the equilibrium magnitude. The expected angular updates are therefore scaled in the same way as the RMS update norm we derived in appx. B.1.

B.4 Relative Representation Change for Normalized Gradient Descent

Property (P1): For a dot product $y = \langle \mathbf{w}, \mathbf{x} \rangle$ and loss $\mathcal{L}(\mathbf{x}_b)$ that depends on y , we have:

$$\frac{\partial \mathcal{L}(\mathbf{x}_b)}{\partial \mathbf{w}} = \frac{\partial \mathcal{L}(\mathbf{x}_b)}{\partial y} \frac{\partial y}{\partial \mathbf{w}} = \frac{\partial \mathcal{L}(\mathbf{x}_b)}{\partial y} \mathbf{x}_b \quad (26)$$

where $\frac{\partial \mathcal{L}(\mathbf{x}_b)}{\partial y}$ is a scalar, ensuring that $\mathbf{g}_b := \frac{\partial \mathcal{L}(\mathbf{x}_b)}{\partial \mathbf{w}} \parallel \mathbf{x}_b$, assuming the vectors are not zero.

Lemma (L1): Consider two independent random vectors $\mathbf{a} \in \mathbb{R}^C$ and $\mathbf{b} \in \mathbb{R}^C$, whose elements are independent and identically distributed (IID). If at least one of the vectors has a zero-mean distribution, then the expected value of the squared inner product of \mathbf{a} and \mathbf{b} is given by:

$$\mathbb{E}[\langle \mathbf{a}, \mathbf{b} \rangle^2] = \frac{\mathbb{E}[\|\mathbf{a}\|^2]\mathbb{E}[\|\mathbf{b}\|^2]}{C} \quad (27)$$

Proof: Let $\mathbf{a} = (a_1, a_2, \dots, a_C)$ and $\mathbf{b} = (b_1, b_2, \dots, b_C)$. The inner product $\langle \mathbf{a}, \mathbf{b} \rangle$ is given by:

$$\langle \mathbf{a}, \mathbf{b} \rangle = \sum_{i=1}^C a_i b_i.$$

We need to find $\mathbb{E}[\langle \mathbf{a}, \mathbf{b} \rangle^2]$. Expanding the square of the inner product:

$$\langle \mathbf{a}, \mathbf{b} \rangle^2 = \left(\sum_{i=1}^C a_i b_i \right)^2 = \sum_{i=1}^C \sum_{j=1}^C a_i b_i a_j b_j.$$

Taking the expectation, we get:

$$\mathbb{E}[\langle \mathbf{a}, \mathbf{b} \rangle^2] = \mathbb{E} \left[\sum_{i=1}^C \sum_{j=1}^C a_i b_i a_j b_j \right] = \sum_{i=1}^C \sum_{j=1}^C \mathbb{E}[a_i b_i a_j b_j].$$

Since \mathbf{a} and \mathbf{b} are independent and their elements are IID, we have:

$$\mathbb{E}[a_i b_i a_j b_j] = \mathbb{E}[a_i a_j] \mathbb{E}[b_i b_j].$$

Consider two cases:

1. When $i = j$:

$$\mathbb{E}[a_i b_i a_i b_i] = \mathbb{E}[a_i^2] \mathbb{E}[b_i^2].$$

2. When $i \neq j$:

$$\mathbb{E}[a_i b_i a_j b_j] = \mathbb{E}[a_i] \mathbb{E}[b_i] \mathbb{E}[a_j] \mathbb{E}[b_j].$$

Given that at least one of \mathbf{a} or \mathbf{b} has a zero-mean distribution, say \mathbf{a} without loss of generality, we have $\mathbb{E}[a_i] = 0$. Thus:

$$\mathbb{E}[a_i b_i a_j b_j] = 0.$$

So, the expectation simplifies to:

$$\mathbb{E}[\langle \mathbf{a}, \mathbf{b} \rangle^2] = \sum_{i=1}^C \mathbb{E}[a_i^2] \mathbb{E}[b_i^2].$$

Since a_i and b_i are IID, we have:

$$\mathbb{E}[a_i^2] = \mathbb{E}[a_1^2] \quad \text{and} \quad \mathbb{E}[b_i^2] = \mathbb{E}[b_1^2].$$

Therefore:

$$\mathbb{E}[\langle \mathbf{a}, \mathbf{b} \rangle^2] = C \mathbb{E}[a_1^2] \mathbb{E}[b_1^2].$$

Recognizing that:

$$\mathbb{E}[\|\mathbf{a}\|^2] = \mathbb{E}\left[\sum_{i=1}^C a_i^2\right] = C\mathbb{E}[a_1^2],$$

$$\mathbb{E}[\|\mathbf{b}\|^2] = \mathbb{E}\left[\sum_{i=1}^C b_i^2\right] = C\mathbb{E}[b_1^2],$$

we have:

$$\mathbb{E}[a_1^2] = \frac{\mathbb{E}[\|\mathbf{a}\|^2]}{C} \quad \text{and} \quad \mathbb{E}[b_1^2] = \frac{\mathbb{E}[\|\mathbf{b}\|^2]}{C}.$$

Thus:

$$\mathbb{E}[\langle \mathbf{a}, \mathbf{b} \rangle^2] = C \left(\frac{\mathbb{E}[\|\mathbf{a}\|^2]}{C} \right) \left(\frac{\mathbb{E}[\|\mathbf{b}\|^2]}{C} \right) = \frac{\mathbb{E}[\|\mathbf{a}\|^2]\mathbb{E}[\|\mathbf{b}\|^2]}{C}.$$

This completes the proof.

Assumption (A1): We assume the following vector pairs satisfy L1: $(\mathbf{x}_i, \tilde{\mathbf{g}}_b)$ when $i \neq b$, $(\bar{\mathbf{g}}, \tilde{\mathbf{g}}_b)$ and $(\mathbf{w}, \mathbf{x}_b)$.

Vector pairs of the type $(\mathbf{x}_i, \tilde{\mathbf{g}}_b)$ and $(\bar{\mathbf{g}}, \tilde{\mathbf{g}}_b)$ should be independent and $\tilde{\mathbf{g}}_b$ has a zero mean distribution. However, the elements of each vector are not necessarily IID. For $(\mathbf{w}, \mathbf{x}_b)$, this is an even stronger assumption. Generally, neither \mathbf{w} nor \mathbf{x}_b is guaranteed to be IID or zero mean, and their independence later in training does not necessarily hold. Applying weight standardization to \mathbf{w} or batch normalization to \mathbf{x} would suffice to make their mean zero. Overall, this assumption can be viewed as a simplifying approximation to obtain reasonable predictions without additional information about these vectors. Everett et al. [5] explore the behavior of $\langle \mathbf{w}, \mathbf{x}_b \rangle$ throughout training and find that it is more complicated than assumed here. This will lead to additional factors that may affect the RRC but we do not attempt to analyze.

Deriving the Relative Representation Change: Applying L1 directly gives us the original expected square output :

$$\mathbb{E}[y_b^2] = \mathbb{E}[\langle \mathbf{w}, \mathbf{x}_b \rangle^2] = \frac{\|\mathbf{w}\|^2 \mathbb{E}[\|\mathbf{x}_b\|^2]}{C} \quad (28)$$

For the expected square representation change we get:

$$\mathbb{E}[(\Delta y_b)^2] \quad (29)$$

$$= \mathbb{E}[\langle -\eta \mathbf{g} / \sqrt{\mathbb{E}[\|\mathbf{g}\|^2]}, \mathbf{x}_b \rangle^2] \quad (30)$$

$$= \frac{\eta^2}{B^2} \frac{1}{\mathbb{E}[\|\mathbf{g}\|^2]} \mathbb{E} \left[\left(\sum_{i=1}^B \langle \mathbf{g}_i, \mathbf{x}_b \rangle \right)^2 \right] \quad (31)$$

$$= \frac{\eta^2}{B^2} \frac{1}{\mathbb{E}[\|\mathbf{g}\|^2]} \mathbb{E} \left[\left(\text{sign}(\langle \mathbf{x}_b, \mathbf{g}_b \rangle) \|\mathbf{g}_b\| \|\mathbf{x}_b\| + \sum_{i \neq B} \langle \mathbf{g}_i, \mathbf{x}_b \rangle \right)^2 \right] \quad (32)$$

$$= \frac{\eta^2}{B^2} \frac{1}{\mathbb{E}[\|\mathbf{g}\|^2]} \mathbb{E} \left[\left(\text{sign}(\langle \mathbf{x}_b, \mathbf{g}_b \rangle) \|\mathbf{g}_b\| \|\mathbf{x}_b\| + (B-1) \langle \bar{\mathbf{g}}, \mathbf{x}_b \rangle + \sum_{i \neq b} \langle \tilde{\mathbf{g}}_i, \mathbf{x}_b \rangle \right)^2 \right] \quad (33)$$

$$\quad (34)$$

where we have used the definitions from eq. (7) and D1. Using property P1, we can write:

$$\langle \bar{\mathbf{g}}, \mathbf{x}_b \rangle = \left\langle \bar{\mathbf{g}}, \text{sign}(\langle \mathbf{x}_b, \mathbf{g}_b \rangle) \frac{\|\mathbf{x}_b\|}{\|\mathbf{g}_b\|} \cdot (\bar{\mathbf{g}} + \tilde{\mathbf{g}}_b) \right\rangle \quad (35)$$

$$= \text{sign}(\langle \mathbf{x}_b, \mathbf{g}_b \rangle) \frac{\|\mathbf{x}_b\|}{\|\mathbf{g}_b\|} (\|\bar{\mathbf{g}}\|^2 + \langle \bar{\mathbf{g}}, \tilde{\mathbf{g}}_b \rangle) \quad (36)$$

Plugging this into the previous expression yields $\mathbb{E}[(\Delta y_b)^2]$

$$= \frac{\eta^2}{B^2} \frac{1}{\mathbb{E}[\|\mathbf{g}\|^2]} \mathbb{E} \left[\left(\text{sign}(\langle \mathbf{x}_b, \mathbf{g}_b \rangle) \left(\|\mathbf{g}_b\| \|\mathbf{x}_b\| + (B-1) \frac{\|\mathbf{x}_b\|}{\|\mathbf{g}_b\|} (\|\bar{\mathbf{g}}\|^2 + \langle \bar{\mathbf{g}}, \tilde{\mathbf{g}}_b \rangle) \right) + \sum_{i \neq b} \langle \tilde{\mathbf{g}}_i, \mathbf{x}_b \rangle \right)^2 \right] \quad (37)$$

Squaring the expression results in various cross but all remaining dot products except the sign one are zero in expectation (due to the noise $\bar{\mathbf{g}}$) and independent from each other. The cross terms involving these thus all disappear under the expectation. We apply Lemma L1 to their squares and approximate the expected norms of \mathbf{x}_b and \mathbf{g}_b as being independent. This gives $\mathbb{E}[(\Delta y_b)^2]$

$$= \frac{\eta^2}{B^2} \frac{\mathbb{E}[\|\mathbf{x}_b\|^2]}{\mathbb{E}[\|\mathbf{g}\|^2]} \left(\mathbb{E}[\|\mathbf{g}_b\|^2] + \frac{(B-1)^2 \|\bar{\mathbf{g}}\|^2}{\mathbb{E}[\|\mathbf{g}\|^2]} \left(\|\bar{\mathbf{g}}\|^2 + \frac{\mathbb{E}[\|\tilde{\mathbf{g}}_b\|^2]}{C} \right) \right) \quad (38)$$

$$+ 2(B-1) \|\bar{\mathbf{g}}\|^2 + \frac{B-1}{C} \mathbb{E}[\|\tilde{\mathbf{g}}_i\|^2] \quad (39)$$

We can compute the expected magnitude of the batch gradient as:

$$\mathbb{E}[\|\mathbf{g}\|^2] = \mathbb{E}[\left\| \frac{1}{B} \sum_{i=1}^B (\bar{\mathbf{g}} + \tilde{\mathbf{g}}_i) \right\|^2] = \mathbb{E}[\left\| (\bar{\mathbf{g}} + \frac{1}{B} \sum_{i=1}^B \tilde{\mathbf{g}}_i) \right\|^2] = \|\bar{\mathbf{g}}\|^2 + \frac{1}{B} \mathbb{E}[\|\tilde{\mathbf{g}}_i\|^2] \quad (40)$$

and similarly $\mathbb{E}[\|\mathbf{g}_b\|^2] = \|\bar{\mathbf{g}}\|^2 + \mathbb{E}[\|\tilde{\mathbf{g}}_b\|^2]$. Using these facts we can further write $\mathbb{E}[(\Delta y_b)^2]$

$$= \frac{\eta^2}{B^2} \frac{\mathbb{E}[\|\mathbf{x}_b\|^2]}{\mathbb{E}[\|\bar{\mathbf{g}}\|^2] + \frac{1}{B} \mathbb{E}[\|\tilde{\mathbf{g}}_i\|^2]} \left(\|\bar{\mathbf{g}}\|^2 + \mathbb{E}[\|\tilde{\mathbf{g}}_b\|^2] + \frac{(B-1)^2 \|\bar{\mathbf{g}}\|^2}{\|\bar{\mathbf{g}}\|^2 + \mathbb{E}[\|\tilde{\mathbf{g}}_b\|^2]} \left(\|\bar{\mathbf{g}}\|^2 + \frac{\mathbb{E}[\|\tilde{\mathbf{g}}_b\|^2]}{C} \right) \right. \\ \left. + 2(B-1) \|\bar{\mathbf{g}}\|^2 + \frac{B-1}{C} \mathbb{E}[\|\tilde{\mathbf{g}}_i\|^2] \right) \quad (41)$$

Combining this with the previous expression for $\mathbb{E}[y_b^2]$ and the definition (D2) of the signal-to-noise ratio $\varphi := \mathbb{E}[\|\bar{\mathbf{g}}\|^2]/\mathbb{E}[\|\tilde{\mathbf{g}}_b\|^2]$ we obtain the expression in the main body:

$$\frac{\mathbb{E}[(\Delta y_b)^2]}{\mathbb{E}[y_b^2]} = \frac{\eta^2 C}{B^2 \|\mathbf{w}\|^2} \frac{1}{\mathbb{E}[\|\mathbf{g}\|^2]} \left(\mathbb{E}[\|\mathbf{g}_b\|^2] + \frac{B-1}{C} \mathbb{E}[\|\tilde{\mathbf{g}}_i\|^2] \right. \\ \left. + \frac{(B-1)^2}{\mathbb{E}[\|\mathbf{g}_b\|^2]} \left(\|\bar{\mathbf{g}}\|^4 + \frac{\|\bar{\mathbf{g}}\|^2 \mathbb{E}[\|\tilde{\mathbf{g}}_b\|^2]}{C} \right) + 2(B-1) \|\bar{\mathbf{g}}\|^2 \right) \quad (42)$$

$$= \frac{\eta^2 C}{B^2 \|\mathbf{w}\|^2} \frac{1}{\varphi + \frac{1}{B}} \left((\varphi + 1) + \frac{B-1}{C} + \left(\frac{(B-1)^2 \varphi}{\varphi + 1} \left(\varphi + \frac{1}{C} \right) + 2(B-1) \varphi \right) \right) \quad (43)$$

B.5 Estimating the Signal-to-Noise Ratio

We use accumulation over the microbatches to estimate the SNR at a given time. Let's assume we have A microbatches of size M each, with the average gradient of a microbatch denoted \mathbf{g}_m and the average gradient of the whole batch denoted $\mathbf{g} = \frac{1}{A} \sum_m \mathbf{g}_m$.

We estimate the variance of the norm of a single gradient example, i.e. the noise power as:

$$P_N = \frac{A}{A-1} \cdot M \cdot \mathbf{1}^\top \left(\frac{1}{A} \sum_m \mathbf{g}_m^2 - \mathbf{g}^2 \right) \quad (44)$$

The signal power is estimated as:

$$P_S = \mathbf{1}^\top \mathbf{g}^2 - \frac{1}{AM} P_N \quad (45)$$

Our SNR estimate is then:

$$\varphi = P_S / P_N \quad (46)$$

B.6 RRC Correction Factor

The RRC correction is done based on eq. (9) and the SNR estimation eq. (46). We assume the learning rate was originally scaled with the square root of the batch size, which is derived for an SNR of zero, and downscale the step size to compensate for the measured SNR and batch size. We define:

$$\rho = \frac{1}{B(1+\varphi)} \left((\varphi+1) + \frac{B-1}{C} + \left(\frac{(B-1)^2 \varphi}{\varphi+1} \left(\varphi + \frac{1}{C} \right) + 2(B-1)\varphi \right) \right) \quad (47)$$

For numerical purposes, we clamp $1 \leq \rho \leq B$ which corresponds to $\varphi = 0$ and $\varphi = \infty$ for a large $C \rightarrow \infty$. The update scaling factor is the square root of an EMA of the inverse of this quantity. We use the same coefficient as for the momentum and compute this for the matrix of each linear layer independently. This form for the scaling factor is somewhat arbitrary, complicated by the fact that Lion-like algorithms fix the step size exactly, so scaling the gradient at each step size can not change the magnitude of the update. For Adam or SGD like algorithms we could scale the gradient contributions directly instead of scaling the update size.

B.7 Run-to-run Variance / Uncertainty Estimation

We do not quantify the uncertainty for every GPT2 configuration in our sweeps. This would require significantly more compute and our estimates of the uncertainty for select points indicate that this would not qualitatively change our results. For the baseline AdamW run the run-to-run differences in the validation loss over different seeds are around 0.05. However, the relative ranking of different runs remained the same.

B.8 Computational Requirements

Our experiments are performed on A100 GPUs with either 40GB or 80GB of RAM. One training run for our GPT2 setup takes around 4h, running on a single GPU. Reproducing the GPT2 experiments reported in the main body should take on the order of 1000 GPU hours. Including our preliminary experiments brings this up to around 3x this amount.

C Additional Experiments

C.1 Comparison with RAdam

RAdamW combines the variance reduction technique of Liu et al. [23] with the decoupled weight decay of Loshchilov and Hutter [24]. Since it is a well known technique for reducing the need for warmup it serves as a good comparison and contextualization of our work. Figure 9 shows that while RAdamW outperforms the AdamW baseline without warmup, it is unable to match longer warmups. Ma and Yarats [26] suggests that RAdamW is approximately equivalent to 4 steps of SGDM followed by AdamW with a special type of built-in warmup with an effective length of around $2/(1 - \beta_2) = 40$, which is likely too short in our setting. For comparison the 2% warmup shown corresponds to 100 steps but is too short for optimal performance.

The analysis of Liu et al. [23] is based on the idea that early in training the second-moment estimates are not accurate (noisy) and can therefore not be trusted to scale the update properly. This could in turn contribute to the need for warmup, although Ma and Yarats [26] question this interpretation. We first note that without momentum, perfect estimates of the second moment at the current time step would control the expected ℓ_2 -norm of the update. This relates our approach of looking at the update size to the adaptive learning rate view of Liu et al. [23]. Secondly, we note that counteracting noisy estimates of the second moment can not be the sole reason warmup is beneficial. This is supported by the fact that both SGD and Lion empirically need warmup in various settings but do not use the second moment at all, indicating there are additional factors that contribute to the need for warmup.

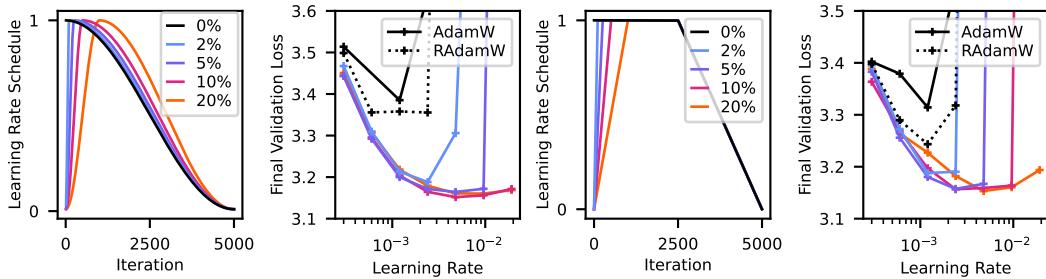


Figure 9: Comparison of the AdamW baseline with a cosine schedule and RAdamW. **Panel 1:** Cosine schedules with different warmup lengths. Note how the warmup shifts the curve, affecting the whole schedule including the decay portion. **Panel 2:** The final validation loss for GPT2-124M training on OpenWebText using the cosine schedules. Note that RAdamW helps but does not eliminate the need for warmup. **Panel 3:** The original trapezoidal schedules used in our experiments. **Panel 4:** Trapezoidal GPT2-124M OpenWebText results. The RAdamW results are similar to those in panel 2.

C.2 Model & Dataset Ablations

In this section we repeat some of our main experiments, varying the dataset and model architecture.

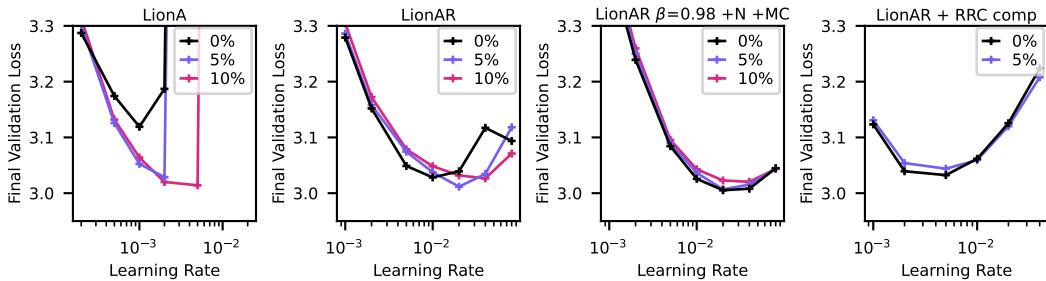


Figure 10: Dataset ablation study, GPT2-124M on SlimPajama [33]. Four update control approaches are shown (see titles). The overall results are similar to those for OpenWebText in the manuscript.

Figure 10 shows the effects of changing the dataset used in our experiments from OpenWebText [8] to SlimPajama [33]. Overall the results are similar as before, when controlling the ℓ_2 -norm via LionA warmup is still beneficial, controlling the angular updates via LionAR decreases the gap significantly. The higher momentum LionAR with Nesterov momentum and our momentum correction eliminates the gap fully. The RRC also seems to eliminate the benefit of warmup but still has the same practical limitations as we describe in §6.

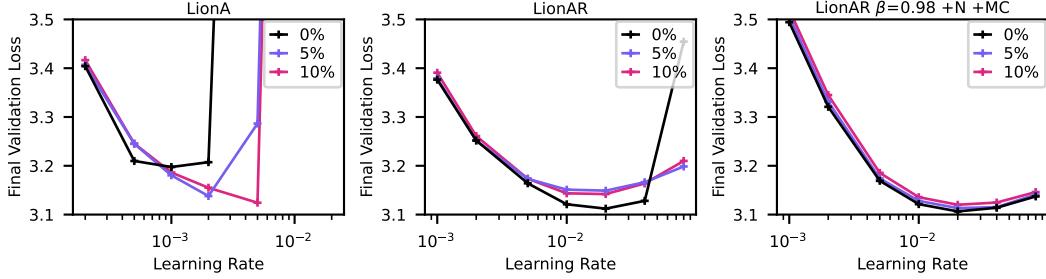


Figure 11: Architecture ablation, Llama2-124M on OpenWebText. Here LionAR is already sufficient to eliminate the need for warmup, no additional RRC compensation is needed. This could be due to the critical batch size being larger (for unclear reasons).

Figure 11 shows the effects of changing the architecture from GPT2 to a Llama style [36] while keeping the dataset and parameter count (124m) the same. This change consists of using SwiGLU activations, RoPE embeddings and RMSNorm. In this case LionAR is able to fully eliminate the need for warmup without any additional tricks like the RRC compensation or momentum corrections. Based on our analysis these additional tricks are likely only needed when the critical batch size is very small initially. We expect that using larger batch sizes could necessitate these additional tricks for Llama, but do not explore this further here.

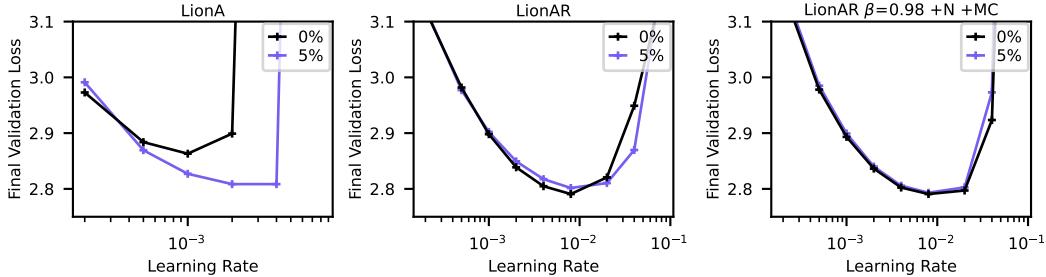


Figure 12: Larger llama2-209M training on SlimPajama. This experiment changes the architecture, dataset, model size, and training length (proportional to the model size). In this case LionAR suffices on its own again, no additional RRC correction needed as in the smaller llama experiments.

Figure 12 uses a larger Llama model with twice the depth. It also trains for twice as many iterations to keep the ratio of tokens to parameters similar. The overall results resemble those of the smaller llama experiments in fig. 11.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: **[Yes]**

Justification: We have aimed to write our abstract and introduction to accurately reflect the paper's contributions and scope.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: **[Yes]**

Justification: We include a limitation section §9 which discusses the main overarching limitations from our perspective.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: Our work is mostly empirical, aside from the analysis of the relative representation change where we state the assumptions upfront, with additional details delegated to the appendix.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We provide the main hyperparameters and other aspects of our training. We are not aware of any missing details that would significantly hinder reproducibility.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [No]

Justification: The datasets we use are freely available online but we have not released our code.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We have attempted include all relevant information.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: Our main results are qualitative and computationally expensive for us, so we do not estimate the variance for each point in our sweeps. We give an estimate based on select points in.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [\[Yes\]](#)

Justification: We do this for our main experiments in appx. B.8

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [\[Yes\]](#)

Justification: We read it and believe we fully conform to it.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [\[NA\]](#)

Justification: This paper is focused on understanding deep learning in general, and does not have an specific societal impacts beyond those of the field overall.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.

- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: We don't release anything that falls under this category.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We do this to the best of our ability.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.

- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: No new assets introduced

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: No crowdsourcing / human subjects used

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: No crowdsourcing / human subjects used

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.

- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.