
High-probability bounds for Non-Convex Stochastic Optimization with Heavy Tails

Ashok Cutkosky
Boston University
ashok@cutkosky.com

Harsh Mehta
Google Research
harshm@google.com

Abstract

We consider non-convex stochastic optimization using first-order algorithms for which the gradient estimates may have heavy tails. We show that a combination of gradient clipping, momentum, and normalized gradient descent yields convergence to critical points in high-probability with best-known rates for smooth losses when the gradients only have bounded p th moments for some $p \in (1, 2]$. We then consider the case of second-order smooth losses, which to our knowledge have not been studied in this setting, and again obtain high-probability bounds for any p . Moreover, our results hold for arbitrary smooth norms, in contrast to the typical SGD analysis which requires a Hilbert space norm. Further, we show that after a suitable “burn-in” period, the objective value will *monotonically decrease* whenever the current iterate is not a critical point, which provides intuition behind the popular practice of learning rate “warm-up” and also yields a last-iterate guarantee.

1 SGD and Heavy Tails

The standard algorithm for training large machine learning models today is stochastic gradient descent (SGD) and its myriad variants [13, 11, 19, 39, 28]. SGD enjoys many properties that help explain its empirical success: it can be implemented very easily, can efficiently process training data in a streaming manner, and even achieves the optimal convergence rate for finding critical points of smooth non-convex objectives [2]. In particular, SGD can be understood as optimizing the following stochastic optimization problem:

$$\min_{\vec{w}} F(\vec{w}) = \mathbb{E}_z[f(\vec{w}, z)]$$

where $f(\vec{w}, z)$ is some loss function taking model parameters \vec{w} and a random example z , where z has some distribution P_z . Then SGD consists of the update:

$$\vec{w}_{t+1} = \vec{w}_t - \eta \nabla f(\vec{w}_t, z_t)$$

where z_1, \dots, z_T are i.i.d. random variables distributed according to P_z , and η is some real number called the learning rate. Assuming that F is L -smooth and that $\nabla f(\vec{w}, z)$ satisfies $\mathbb{E}[\|\nabla f(\vec{w}, z)\|^2] \leq G^2$ for some G for all \vec{w} , SGD satisfies [13]:

$$\frac{1}{T} \mathbb{E} \left[\sum_{t=1}^T \|\nabla F(\vec{w}_t)\|^2 \right] \leq O \left(\frac{L\Delta}{T} + \frac{G^2}{\sqrt{T}} \right)$$

where $\Delta = F(\vec{w}_1) - \inf F(\vec{w})$. This implies that SGD takes $O(\epsilon^{-4})$ iterations to find an ϵ -critical point, which is a point \vec{w} such that $\|\nabla F(\vec{w})\| \leq \epsilon$ [13]. When the stochastic gradients are further assumed to have light sub-gaussian tails, [21] shows that SGD with momentum can achieve this same result in high probability as well.

However, recent work [30, 31, 42, 41, 43] suggests that the assumptions of bounded variance or light tails may be too optimistic in practical deep learning problems. For example, Zhang et al. [43] provides empirical evidence that large NLP models based on attention and transformers [35, 10] have *heavy-tailed* gradients. In particular, the variance is extremely large (perhaps even essentially infinite for practical purposes), but the p th moment is bounded for some $p \in (1, 2]$. This setting is significantly more complicated, and SGD’s convergence properties are difficult to characterize even when assuming slightly stronger tail bounds [29] or various versions of convexity [36].

A standard tool when working with heavy-tailed random variables is truncation, or *clipping*. Succinctly, if X is a random variable such that $\mathbb{E}[\|X\|^2] = \infty$, X itself may be poorly behaved due to the heavy tail, but a truncated value $X^{\text{clip}} = \frac{X}{|X|} \min(|X|, \tau)$ will be much more benign so long as τ is chosen in an appropriate manner [6, 23]. Fortunately, gradient clipping has long been a standard technique to improve training of deep neural networks [26], and recent analysis [43, 14] suggests that clipping allows SGD to deal with heavy-tailed stochastic gradients. Specifically, [43] shows that SGD combined with clipping can find an ϵ -critical point in expectation in $O\left(\epsilon^{-\frac{3p-2}{p-1}}\right)$ iterations.

We propose to improve upon this backdrop in three important ways. First, the prior analysis of [43] for non-convex convergence of SGD under heavy tails only provides an *in-expectation* analysis. In-expectation results are unsatisfying because the quantity measured for convergence (i.e. gradient norm or function value) is itself a random variable, and so might *also* have a heavy tail. In this case individual training runs could be quite poor. Since training modern networks is quite expensive [27, 5], it is important to be confident that each and every training run will produce a high-quality result. We provide a variant of the normalized SGD with momentum algorithm of [9] that incorporates gradient clipping and show that the method finds an ϵ -critical point in $\tilde{O}\left(\epsilon^{-\frac{3p-2}{p-1}}\right)$ iterations with high probability.

Second, the vast majority of first-order optimization methods for stochastic non-convex optimization consider exclusively the standard L_2 or Hilbert-space norm. Such analyses can take advantage of the fact that $\mathbb{E}[\langle \nabla F(x), \vec{g} \rangle] = \|\nabla F(x)\|^2$ when \vec{g} is an unbiased estimate of $\nabla F(x)$. In contrast, our algorithms find critical points when the gradient is measured with any smooth norm (i.e. in a Banach space). In this case, the gradient estimate \vec{g} (which is a dual vector) must be converted into an appropriate primal vector. This operation is non-linear in general and so may add some bias. It turns out the key to operating with different norms is the popular practice of augmenting SGD iterates with momentum. When using momentum-based updates we do not rely on the same unbiasedness argument, and so the techniques can be extended to more general norms. This provides more concrete evidence of the theoretical advantages of momentum.

Third, recent theoretical advances in non-convex optimization have produced a number of new algorithms that avoid the lower-bounds of [2] by assuming extra structure on the objective F , such as *second-order smoothness*. In this case, [32, 1, 12, 9, 3] provide algorithms that achieve faster convergence rates. In particular, the algorithms of [12, 9] can find an ϵ -critical point in $O(\epsilon^{3.5})$ iterations using a first-order stochastic oracle, which is the optimal rate [3]. To our knowledge, no analysis even in expectation exists for second-order smooth losses under heavy tails. We fill this gap, providing an algorithm that identifies an ϵ -critical point in $\tilde{O}\left(\epsilon^{-\frac{5p-3}{2p-2}}\right)$ iterations with high probability. Further, in both smooth and second-order smooth cases we provide a convergence guarantee for the *last iterate* of the optimization algorithm: with high probability, the objective value for the final iterate will be smaller than the objective value at some approximate critical point.

Finally, we discuss an intriguing connection to the popular *warm-up* learning rate schedule [16, 15, 10], in which the learning rate starts small and increases over the first iterations of training. We show that for normalized SGD with momentum, after a brief “burn-in” period, the objective value decreases monotonically until an approximate critical point is reached. This provides an alternative motivation for warm-up: it would be reasonable to keep the learning rate small during the “burn-in” period. There is a long line of empirical research in understanding the need for warm-up. Warm-up has been shown to be connected to the use of layer-normalization, and that improved use of layer-norm can enable warm-up free training of transformer models [7, 37, 25]. Huang et al. [18] further build on these insights and proposes a weight initialization scheme which removes the need for both layer normalization and warm-up while training transformer models. Finally, Liu et al. [22] demonstrate that the variance of the learning rate in optimizers like Adam is abnormally high in the early stages of

the optimization process and further hypothesize that warm-up aids by reducing the variance in that stage. Our results provide a different view from an optimization perspective.

Paper Organization: in Section 1.1 we provide our concrete assumptions and setup. In Sections 2 and 3, we provide our algorithms and high-probability convergence bounds. In Section 4, we provide some empirical observations on the effect of the burn-in phase in our analysis.

1.1 Setup and assumptions

We assume that the objective $F(\vec{w}) = \mathbb{E}[f(\vec{w}, z)]$ is a Frechet-differentiable function mapping a Banach space \mathcal{B} to the reals \mathbb{R} . Let $\|\cdot\|$ be the norm in \mathcal{B} and $\|\cdot\|_*$ be the dual norm in the dual space \mathcal{B}^* . To preserve analogy and notation from the Hilbert space setting, we will write $\langle v, w \rangle$ to indicate the application of a dual vector $v \in \mathcal{B}^*$ to a vector $w \in \mathcal{B}$. We assume that $\|\cdot\|_*^p$ is Frechet differentiable and satisfies for some p, C for all $x, y \in \mathcal{B}$:

$$\|x + y\|_*^p \leq \|x\|_*^p + \langle \nabla \|x\|_*^p, y \rangle + C\|y\|_*^p \quad (1)$$

Such a Banach space is called (p, C) -smooth. Note that if $\|\cdot\|$ is the standard p norm in \mathbb{R}^n for $p \in (1, 2]$, then $\|\cdot\|$ is the corresponding q norm, and the dual space is $(2, \frac{1}{p-1})$ -smooth. In the common case that \mathcal{B} is a Hilbert space (e.g. \mathbb{R}^n with the euclidean norm), then $\mathcal{B} = \mathcal{B}^*$, $\|\cdot\| = \|\cdot\|_*$ and $p = 2$ and $C = 1$. In this paper, we will exclusively consider the case $p = 2$ and $C \geq 1$. Finally, given any $v \in \mathcal{B}^*$, we will write $d(v) \in \mathcal{B}$ to indicate a unit vector satisfying both $\|d(v)\| = 1$ and $\langle v, d(v) \rangle = \|v\|_*$, which we assume exists for all v . For more background on Banach spaces, see Appendix E.

We assume F is L -smooth, which means that $\|\nabla F(\vec{w}) - \nabla F(\vec{x})\|_* \leq L\|\vec{w} - \vec{x}\|_*$. In Section 3, we will also assume that F is ρ -second-order smooth, which means that $\|(\nabla^2 F(\vec{w}) - \nabla^2 F(\vec{x}))v\|_* \leq \rho\|\vec{w} - \vec{x}\|\|v\|$ for all $\vec{w}, \vec{x}, v \in \mathcal{B}$. We assume that $\nabla f(\vec{w}, z)$ satisfies both $\mathbb{E}[\|\nabla f(\vec{w}, z) - \nabla F(\vec{w})\|_*^p] \leq G^p$ and $\mathbb{E}[\|\nabla f(\vec{w}, z)\|_*^p] \leq G^p$ for some G and some $p \in (1, 2]$. We will design algorithms such that after T stochastic gradient evaluations, we can output a point \vec{w} such that $\|\nabla F(\vec{w})\|_*$ is small with high probability.

2 Normalized SGD with Momentum

In this section, we describe our algorithms for finding critical points with high probability in non-convex objectives. Our analysis makes use of gradient clipping, momentum, and normalized updates. Clipping is a standard technique for mitigating the effect of heavy tails, essentially by throwing out outlier data. Momentum is useful as it intuitively averages together many clipped gradients, resulting in a gradient estimate that concentrates about its mean with high probability. Normalizing the updates allows for significantly simplified analyses because the bias terms introduced by the moving average in momentum are now very precisely controlled by the learning rate, as originally analyzed for Hilbert spaces in [9]. In order to analyze this algorithm, we need the following critical Lemma, which characterizes the effect of taking a normalized SGD step (proved in appendix E).

Lemma 1. *Suppose $F : \mathcal{B} \rightarrow \mathbb{R}$ be a Frechet-differentiable L -Smooth function from a Banach space \mathcal{B} to the reals. Let $w \in \mathcal{B}$. Let $g^* \in \mathcal{B}^*$ and let $g \in \mathcal{B}$ be a unit-vector satisfying $\langle g^*, g \rangle = \|g\|_*$. Define $w' = w - \eta g$. Define $\epsilon = g^* - \nabla F(w)$. Then:*

$$F(w') \leq F(w) - \eta\|\nabla F(w)\|_* + 2\eta\|\epsilon\|_* + \frac{L\eta^2}{2}$$

Theorem 2. *Suppose $\mathbb{E}_z[\|\nabla f(\vec{w}, z)\|_*^p] \leq G^p$ for all \vec{x} for some G . Suppose F is L -smooth. Set $\beta = 1 - \alpha$, $\alpha = \frac{b}{T^{3p-2}}$ and $\eta = \frac{s}{T^{2p-1}}$ for arbitrary constant b and s satisfying $\alpha \leq 1$ and set $\tau = G/\alpha^{1/p}$. Then with probability at least $1 - \delta$:*

$$\frac{1}{T} \sum_{t=1}^T \|\nabla F(\vec{w}_t)\|_* \leq O\left(\frac{\log(T/\delta)}{T^{\frac{p-1}{3p-2}}}\right)$$

where the big-Oh hides constant that depend on L, G, b, s, C and $F(\vec{w}_1) - F(\vec{w}_{T+1})$, but not T or δ .

Algorithm 1 Normalized SGD with Clipping and Momentum

Input: Initial Point \vec{w}_1 , learning rate η , momentum parameter β , clipping parameter τ , time horizon T :
 Set $\vec{m}_0 = 0$.
for $t = 1 \dots T$ **do**
 Sample $z_t \sim P_z$.
 Set $\vec{g}_t^{\text{clip}} = \frac{\nabla f(\vec{w}_t, z_t)}{\|\nabla f(\vec{w}_t, z_t)\|_*} \min(\tau, \|\nabla f(\vec{w}_t, z_t)\|_*)$.
 Set $\vec{m}_t = \beta \vec{m}_{t-1} + (1 - \beta) \vec{g}_t^{\text{clip}}$.
 Set $\vec{w}_{t+1} = \vec{w}_t - \eta d(\vec{m}_t)$.
end for

This Theorem provides an analog of the standard in-expectation result for non-convex optimization, but now the result holds with high-probability. Note that this matches the in-expectation convergence rates from Zhang et al. [43] up to logarithmic factors. The next Theorem takes this theme one step further: it shows that after a brief “burn-in” period, the decrease in objective value will in fact be *monotonic* with high probability:

Theorem 3. *Under the assumptions of Theorem 2, define the constants:*

$$K = 10C \max(1, \log(3T/\delta)) + 4C^{1/2} \sqrt{\max(1, \log(3T/\delta))} + 1$$

$$Z = \frac{sL}{b} + GKb^{\frac{p-1}{p}}$$

Then, with probability at least $1 - \delta$, for all t satisfying:

$$t \geq \mathcal{T} = \frac{T^{\frac{p}{3p-2}}}{b} \left(\frac{p-1}{3p-2} \log(T) + \log(G) - \log(Z) \right)$$

we have

$$\|\vec{m}_t - \nabla F(\vec{w}_t)\|_* \leq \frac{2Z}{T^{\frac{p-1}{3p-2}}}$$

Further, so long as

$$\|\vec{m}_t\|_* \geq 2 \left(\frac{6Z}{T^{\frac{p-1}{3p-2}}} + \frac{Ls}{2T^{\frac{2p-1}{3p-2}}} \right) \quad (2)$$

we have $F(\vec{w}_{t+1}) < F(\vec{w}_t) - \frac{\eta}{2} \|\vec{m}_t\|_*$. Moreover, whenever (2) is not satisfied for $t \geq \mathcal{T}$, we must have

$$\|\nabla F(\vec{w}_t)\|_* \leq \frac{14Z}{T^{\frac{p-1}{3p-2}}} + \frac{Ls}{T^{\frac{2p-1}{3p-2}}}$$

Finally, if t is the iteration with smallest value of $\|\vec{m}_t\|_*$ such that $t \geq \mathcal{T}$, we have:

$$\|\nabla F(\vec{w}_t)\|_* \leq O\left(\frac{\log(T/\delta)}{T^{\frac{p-1}{3p-2}}}\right)$$

Theorem 3 shows that after a small burn-in period $\mathcal{T} = o(T)$, \vec{m}_t is very likely to be an extremely good estimate of the true gradient $\nabla F(\vec{x}_t)$. Moreover, we can empirically detect this event by monitoring the norm of the exponentially weighted moving average \vec{m} . This burn-in period is especially interesting given the success of learning rate schedules that incorporate a “warm-up” component, in which the learning rate actually increases during the first iterations. Our result seems to hint at a mechanism for this schedule: by keeping the learning rate small during the burn-in period, we can limit any increase in the objective value until we reach the stage at which it begins to decrease steadily.

Further, Theorem 3 provides a very concrete algorithm to find critical points: after the burn-in period of $\tilde{O}(T^{\frac{p}{3p-2}})$ iterates, we can simply return the point \vec{x}_t with the smallest value of $\|\vec{m}_t\|_*$. Intuitively, this works because the objective value can only increase by a total of $O(T^{\frac{p}{3p-2}} \eta) = O(T^{\frac{1-p}{3p-2}}) = O(1)$ during the burn-in period so even without a learning rate “warm-up”, we should not expect the burn-in period to have too significant an effect on the objective value. After the burn-in period, \vec{m}_t is always very close to $\nabla F(\vec{w}_t)$, so we make steady progress until a critical point is encountered.

Proof of Theorem 2. Define $\hat{\epsilon}_t = \bar{m}_t - \nabla F(\bar{w}_t)$, and $\epsilon_t = \bar{g}_t^{\text{clip}} - \nabla F(\bar{w}_t)$. Also, define $S(a, b) = \nabla F(a) - \nabla F(b)$. By smoothness, we have $\|S(\bar{w}_t, \bar{w}_{t+1})\|_* \leq L\|\bar{w}_t - \bar{w}_{t+1}\|_*$

From the definition of \bar{m}_{t+1} , we have the following recursive formulation for any $t \geq 1$:

$$\begin{aligned}\bar{m}_{t+1} &= (1 - \alpha)(\nabla F(\bar{w}_t) + \hat{\epsilon}_t) + \alpha \bar{g}_{t+1}^{\text{clip}} \\ &= \nabla F(\bar{w}_{t+1}) + (1 - \alpha)(S(\bar{w}_t, \bar{w}_{t+1}) + \hat{\epsilon}_t) + \alpha \epsilon_{t+1} \\ \hat{\epsilon}_{t+1} &= (1 - \alpha)S(\bar{w}_t, \bar{w}_{t+1}) + (1 - \alpha)\hat{\epsilon}_t + \alpha \epsilon_{t+1}\end{aligned}\quad (3)$$

Further, by defining $\bar{m}_0 = 0$ and $\bar{w}_0 = \bar{w}_1$ so that $\hat{\epsilon}_0 = -\nabla F(\bar{w}_1)$, we have that (3) holds for $t = 0$ as well. Now, we unravel the recursion for all iterations:

$$\hat{\epsilon}_{t+1} = (1 - \alpha)^{t+1}\hat{\epsilon}_0 + \alpha \sum_{t'=0}^t (1 - \alpha)^{t'} \epsilon_{t+1-t'} + (1 - \alpha) \sum_{t'=0}^t (1 - \alpha)^{t'} S(\bar{w}_{t-t'}, \bar{w}_{t+1-t'})$$

Next, take the magnitude of both sides and use triangle inequality (using the fact that $\|S(\bar{w}_t, \bar{w}_{t+1})\|_* \leq L\eta$):

$$\begin{aligned}\|\hat{\epsilon}_{t+1}\|_* &\leq (1 - \alpha)^{t+1}\|\hat{\epsilon}_0\|_* + \alpha \left\| \sum_{t'=0}^t (1 - \alpha)^{t'} \epsilon_{t+1-t'} \right\|_* + (1 - \alpha)\eta L \sum_{t'=0}^t (1 - \alpha)^{t'} \\ &\leq (1 - \alpha)^{t+1}G + \frac{(1 - \alpha)\eta L}{\alpha} + \alpha \left\| \sum_{t'=0}^t (1 - \alpha)^{t'} \epsilon_{t+1-t'} \right\|_*\end{aligned}$$

where in the second line we have used $\bar{m}_0 = 0$ and $\|\nabla F(\bar{w}_0)\|_* \leq G$, which follows since $\|\nabla F(\bar{w}_0)\|_* \leq \mathbb{E}_z[\|\nabla f(\bar{w}_0, z)\|_*^p]^{1/p} = G$.

Now, we invoke Lemma 9, a concentration inequality in Banach space (see appendix B). This implies that with probability at least $1 - \delta/T$:

$$\begin{aligned}\left\| \sum_{t'=0}^t (1 - \alpha)^{t'} \epsilon_{t+1-t'} \right\|_* &\leq 10C\tau \max(1, \log(3T/\delta)) + \sum_{t'=0}^t \frac{(1 - \alpha)^{t'} G^p}{\tau^{p-1}} \\ &\quad + 4 \left(C \sum_{t'=0}^t (1 - \alpha)^{2t'} G^p \tau^{2-p} \right)^{1/2} \sqrt{\max(1, \log(3T/\delta))} \\ &\leq 10C\tau \max(1, \log(3T/\delta)) + \frac{G^p}{\alpha\tau^{p-1}} + \frac{4(CG^p\tau^{2-p})^{1/2} \sqrt{\max(1, \log(3T/\delta))}}{\alpha^{1/2}}\end{aligned}$$

Therefore, with probability at least $1 - \delta/T$:

$$\begin{aligned}\|\hat{\epsilon}_{t+1}\|_* &\leq (1 - \alpha)^{t+1}G + \frac{(1 - \alpha)\eta L}{\alpha} + 10C\alpha\tau \max(1, \log(3T/\delta)) \\ &\quad + 4(\alpha CG^p\tau^{2-p})^{1/2} \sqrt{\max(1, \log(3T/\delta))} + \frac{G^p}{\tau^{p-1}}\end{aligned}$$

Now, with $\tau = \frac{G}{\alpha^{1/p}}$, and $D_\delta = \max(1, \log(3T/\delta))$ this becomes:

$$\begin{aligned}&\leq (1 - \alpha)^{t+1}G + \frac{(1 - \alpha)\eta L}{\alpha} + G\alpha^{1-1/p} \left[10CD_\delta + 4C^{1/p}\sqrt{D_\delta} + 1 \right] \\ &= (1 - \alpha)^{t+1}G + \frac{(1 - \alpha)\eta L}{\alpha} + G\alpha^{1-1/p}K\end{aligned}\quad (4)$$

where we define the constant $K = 10CD_\delta + 4C^{1/p}\sqrt{D_\delta} + 1$.

Now, we have from Lemma 1:

$$\sum_{t=1}^T \|\nabla F(\bar{w}_t)\|_* \leq \frac{F(\bar{w}_1) - F(\bar{w}_T)}{\eta} + \frac{LT\eta}{2} + 2 \sum_{t=1}^T \|\hat{\epsilon}_t\|_*$$

so with probability at least $1 - \delta$:

$$\begin{aligned} &\leq \frac{F(\vec{w}_1) - F(\vec{w}_T)}{\eta} + \frac{LT\eta}{2} + 2 \sum_{t=1}^T (1 - \alpha)^t G + \frac{(1 - \alpha)\eta L}{\alpha} + G\alpha^{1-1/p} K \\ &\leq \frac{F(\vec{w}_1) - F(\vec{w}_T)}{\eta} + \frac{LT\eta}{2} + \frac{2G(1 - \alpha)}{\alpha} + \frac{2(1 - \alpha)\eta TL}{\alpha} + 2GT\alpha^{1-1/p} K \end{aligned}$$

Now set $\alpha = \frac{b}{T^{\frac{p}{3p-2}}}$ and $\eta = \frac{s}{T^{\frac{2p-1}{3p-2}}}$. This yields:

$$\frac{1}{T} \sum_{t=1}^T \|\nabla F(\vec{w}_t)\| \leq O\left(T^{\frac{1-p}{3p-2}} \log(T/\delta)\right)$$

□

The proof of Theorem 3 follows essentially the same idea as that of Theorem 2. However, this time we use the explicit bound on $\|\hat{\epsilon}_t\|_*$ and observe that after the burn-in period, the contribution from $\|\epsilon_1\|_*$ has decreased exponentially to be insignificant, and so \vec{m}_t will be a very high-quality estimate of $\nabla F(\vec{x}_t)$ at every single iteration.

Proof of Theorem 3. Use the settings $\eta = \frac{s}{T^{\frac{2p-1}{3p-2}}}$ and $\alpha = \frac{b}{T^{\frac{p}{3p-2}}}$ and notice that from equation (4) that with probability $1 - \delta$, for all t :

$$\|\hat{\epsilon}_{t+1}\|_* \leq (1 - \alpha)^t G + \frac{sL}{bT^{\frac{p-1}{3p-2}}} + \frac{GKb^{\frac{p-1}{p}}}{T^{\frac{p-1}{3p-2}}} = (1 - \alpha)^t G + \frac{Z}{T^{\frac{p-1}{3p-2}}}$$

Further, we have

$$\begin{aligned} \log(1 - \alpha) &\leq -\alpha \\ \log(G(1 - \alpha)^t) &\leq -t\alpha + \log(G) \end{aligned}$$

Therefore, if $t \geq \mathcal{T}$, then we have with probability at least $1 - \delta$:

$$\|\hat{\epsilon}_{t+1}\|_* \leq \frac{2Z}{T^{\frac{p-1}{3p-2}}}$$

This implies the bound on $\|\nabla F(\vec{w}_t)\|_*$ when $\|\vec{m}_t\|_*$ is smaller than the threshold described in the Theorem statement. Next, again from Lemma 1,

$$\begin{aligned} F(\vec{w}_{t+1}) &\leq F(\vec{w}_t) - \eta \|\nabla F(\vec{w}_t)\|_* + 2\eta \|\hat{\epsilon}_t\|_* + \frac{L\eta^2}{2} \\ &\leq F(\vec{w}_t) - \eta \|\vec{m}_t\|_* + 3\eta \|\hat{\epsilon}_t\|_* + \frac{L\eta^2}{2} \\ &\leq F(\vec{w}_t) - \eta \|\vec{m}_t\|_* + \eta \left(\frac{6Z}{T^{\frac{p-1}{3p-2}}} + \frac{Ls}{2T^{\frac{2p-1}{3p-2}}} \right) \end{aligned}$$

from which we can conclude that if $\|\vec{m}_t\|_* \geq 2 \left(\frac{6Z}{T^{\frac{p-1}{3p-2}}} + \frac{Ls}{2T^{\frac{2p-1}{3p-2}}} \right)$, $F(\vec{w}_{t+1}) \leq F(\vec{w}_t) - \frac{\eta}{2} \|\vec{m}_t\|_*$.

Now, for the final part of the Theorem, observe that since F is G -Lipschitz, we must have that

$$F(\vec{w}_{\mathcal{T}}) - F(\vec{w}_1) \leq G\mathcal{T}\eta = O(\log(T))$$

so that $F(\vec{w}_{\mathcal{T}}) - F(\vec{w}_{T+1}) \leq O(F(\vec{w}_1) - F(\vec{w}_{T+1}) + \log(T))$. Now consider two cases, either

$$\min_{t \geq \mathcal{T}} \|\vec{m}_t\|_* \leq 2 \left(\frac{6Z}{T^{\frac{p-1}{3p-2}}} + \frac{Ls}{2T^{\frac{2p-1}{3p-2}}} \right)$$

or not. In the first case, by our bound on $\|\hat{\epsilon}_t\|_*$, the desired bound on $\|\nabla F(\vec{w}_t)\|_*$ follows. In the latter case, we have

$$F(\vec{w}_{T+1}) \leq F(\vec{w}_{\mathcal{T}}) - \frac{\eta}{2} \sum_{t=\mathcal{T}}^T \|\vec{m}_t\|_*$$

So that

$$\min_{t \geq T} \|\bar{m}_t\|_* \leq O\left(\frac{F(\bar{w}_1) - F(\bar{w}_{T+1}) + \log(T)}{T\eta}\right) = \tilde{O}\left(\frac{1}{T^{\frac{2p-1}{3p-2}}}\right)$$

And so again the result follows from our bound on $\|\hat{\epsilon}_t\|_*$. \square

In addition to providing a method for identifying critical points, Theorem 3 also provides an intuitively desirable guarantee about the *last-iterate* of the algorithm \bar{w}_T . Specifically, one can easily check that the following Corollary:

Corollary 4. *Under the notation and assumptions of Theorem 3, with probability at least $1 - \delta$, there exists some \hat{w} such that both of the following inequalities hold:*

$$\begin{aligned} \|\nabla F(\hat{w})\|_* &\leq O\left(\frac{\log(T/\delta)}{T^{\frac{p-1}{3p-2}}}\right) \\ F(\bar{w}_T) &\leq F(\hat{w}) \end{aligned}$$

where \bar{w}_T is the last iterate of Algorithm 1

Proof. Set \hat{t} to be the last index such that $\bar{w}_{\hat{t}}$ does not satisfy (2). Set $\hat{w} = \bar{w}_{\hat{t}}$. Then Theorem 3 implies the result. \square

In most applications we are not actually interested in finding critical points, but instead wish to actually minimize the objective F . This observation tells us that our objective value is at least never worse than it would have been if had indeed searched specifically for a critical point. Moreover, by providing a guarantee about \bar{w}_T , we have a closer match to practical use-cases: many of the theoretical analyses of non-convex stochastic gradient methods [13, 38, 20] only show that a *randomly selected iterate* has a small gradient norm. In contrast, Corollary 4 provides a non-asymptotic guarantee for the final iterate, which is the iterate most likely to be deployed after training in practice.

3 Second-Order Smooth Losses

In this section, we provide our extension to losses that are *second order smooth*. The overall method is very similar: we employ normalized SGD with momentum. However, in order to take advantage of the second-order smoothness, we will need to use a more advanced form of momentum. Specifically, we use implicit gradient transport [4], as implemented by [9] in their NIGT algorithm. This algorithm replaces the standard momentum update with an extrapolation procedure:

$$\bar{m}_t = \beta \bar{m}_{t-1} + (1 - \beta) \nabla f\left(\bar{w}_t + \frac{\beta(\bar{w}_t - \bar{w}_{t-1})}{1 - \beta}, z_t\right)$$

By evaluating the gradient at the “shifted” point $\bar{w}_t + \frac{\beta(\bar{w}_t - \bar{w}_{t-1})}{1 - \beta}$ and performing a second-order Taylor expansion enabled by the second-order smoothness assumption, it is possible to show that \bar{m}_t is a less-biased estimate of $\nabla F(\bar{w}_t)$ than it would be using the standard momentum update. We augment this procedure with gradient clipping in Algorithm 2 below, and provide its analysis in Theorems 5 and 6, which are directly analogous to Theorems 2 and 3.

Theorem 5. *Suppose $\mathbb{E}_z[\|\nabla f(\bar{w}, z)\|^p] \leq G^p$ for all \bar{w} for some G . Suppose F is L -smooth and ρ -second-order smooth. Set $\beta = 1 - \alpha$, $\alpha = \frac{b}{T^{5p-3}}$ and $\eta = \frac{a}{T^{3p-1}}$ for arbitrary constants b and c satisfying $\alpha \leq 1$, and set $\tau = \frac{G}{\alpha^{1/p}}$. Then with probability at least $1 - \delta$:*

$$\frac{1}{T} \sum_{t=1}^T \|\nabla F(\bar{w}_t)\|_* \leq O\left(\frac{\log(T/\delta)}{T^{\frac{2p-2}{5p-3}}}\right)$$

where the big-Oh hides constant that depend on L , ρ , G , b , s , C , and $F(\bar{w}_1) - F(\bar{w}_{T+1})$, but not T or δ .

We also have a direct analog of Theorem 3

Algorithm 2 NIGT with Clipping

Input: Initial Point \vec{w}_1 , learning rate η , momentum parameter β , clipping parameter τ , time horizon T :
Set $\vec{m}_0 = 0$.
for $t = 1 \dots T$ **do**
 Sample $z_t \sim P_z$.
 $\vec{x}_t = \vec{w}_t + \frac{\beta(\vec{w}_t - \vec{w}_{t-1})}{1-\beta}$
 Set $\vec{g}_t^{\text{clip}} = \frac{\nabla f(\vec{x}_t, z_t)}{\|\nabla f(\vec{x}_t, z_t)\|_*} \min(\tau, \|\nabla f(\vec{x}_t, z_t)\|_*)$.
 Set $\vec{m}_t = \beta\vec{m}_{t-1} + (1-\beta)\vec{g}_t^{\text{clip}}$.
 Set $\vec{w}_{t+1} = \vec{w}_t - \eta d(\vec{m}_t)$.
end for

Theorem 6. *Under the assumptions of Theorem 5, define the constants:*

$$K = 10 \max(1, \log(3T/\delta)) + 4\sqrt{C \max(1, \log(3T/\delta))} + 1$$
$$Z = \frac{\rho s^2}{b^2} + GKb^{\frac{p-1}{p}}$$

Then, with probability at least $1 - \delta$, if

$$t \geq \mathcal{T} = \frac{T^{\frac{2p}{5p-3}}}{b} \left(\frac{2p-2}{5p-3} \log(T) + \log(G) - \log(Z) \right)$$

Then we have

$$\|\vec{m}_t - \nabla F(\vec{x}_t)\| \leq \frac{2Z}{T^{\frac{2p-2}{5p-3}}}$$

And so long as

$$\|\vec{m}_t\|_* \geq 2 \left(\frac{6Z}{T^{\frac{2p-2}{5p-3}}} + \frac{Ls}{2T^{\frac{3p-1}{5p-3}}} \right) \quad (5)$$

we have $F(\vec{w}_{t+1}) < F(\vec{w}_t) - \frac{\eta}{2} \|\vec{m}_t\|$. Moreover, if (5) is ever not satisfied, we must have

$$\|\nabla F(\vec{w}_t)\|_* \leq \frac{14Z}{T^{\frac{2p-2}{5p-3}}} + \frac{Ls}{T^{\frac{3p-1}{5p-3}}}$$

Finally, if t is the iteration with smallest value of $\|\vec{m}_t\|$ such that $t \geq \mathcal{T}$, we have with probability at least $1 - \delta$:

$$\|\nabla F(\vec{w}_t)\|_* \leq O\left(\frac{\log(T/\delta)}{T^{\frac{2p-2}{5p-3}}}\right)$$

Finally, we note that we have a statement about the last iterate of Algorithm 2 that is directly analogous to Corollary 4:

Corollary 7. *Under the notation and assumptions of Theorem 6, with probability at least $1 - \delta$, there exists some \hat{w} such that both of the following inequalities hold:*

$$\|\nabla F(\hat{w})\|_* \leq O\left(\frac{\log(T/\delta)}{T^{\frac{2p-2}{5p-2}}}\right)$$
$$F(\vec{w}_T) \leq F(\hat{w})$$

where \vec{w}_T is the last iterate of Algorithm 2

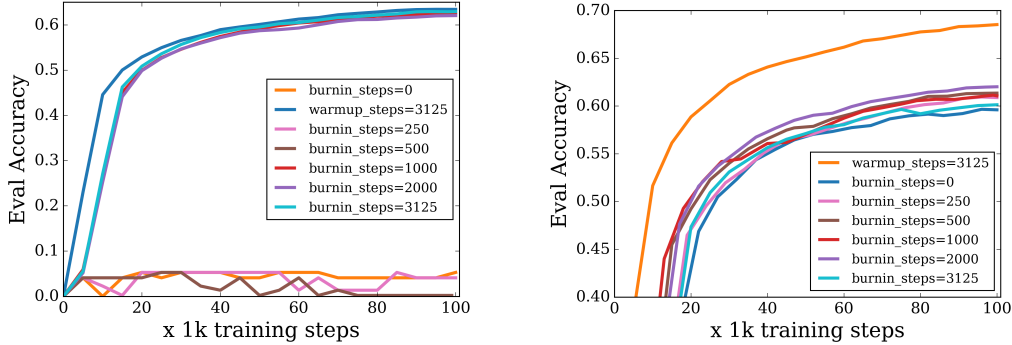


Figure 1: Comparison between warm-up and burn-in strategies on Masked language modeling BERT-Large pre-training task. Plot (a) shows the eval accuracy when trained with Algorithm 1 and plot (b) does the same when trained with LAMB.

4 Experiments

Our theoretical analysis suggests that there is a short “burn-in” phase in our algorithms during which the momentum estimate \vec{m}_t may be far from the true gradient $\nabla F(\vec{w}_t)$, and after this period, \vec{m}_t will always be close to $\nabla F(\vec{w}_t)$. As a result, during the burn-in phase it is possible for the optimizer to make slow or even negative progress, while afterwards we should expect more steady improvement. This suggests that one might set the learning rate $\eta = 0$ throughout the burn-in phase and only begin taking steps after \vec{m}_t has accumulated enough gradients to be a good estimate. This procedure is reminiscent of the common practice of learning rate “warm-up”, in which the learning rate starts at 0 and is linearly increased to some desired value during the initial phase of training, and then later decreased. Thus, we sought to answer how much of the success of the warm-up schedule can be attributed to burn-in.

To answer this question, we conducted an experimental study on the BERT pretraining task. We chose to experiment on BERT since it has been shown to have heavy tails empirically [43] and the common practice is already to clip the gradients. Figure 3 plots the masked language model accuracy of BERT-Large model when trained with Algorithm 1. We also include experimental results using the LAMB optimizer instead of Algorithm 1. LAMB optimizer has been shown to obtain state of the art performance on BERT task [40]. Moreover, both our analysis and the LAMB optimizer relies on normalized updates. Note that we also include experimental results using the smaller BERT-Base model in the appendix. All our experiments were conducted using the Tensorflow framework on TPUv3 architecture.

When using Algorithm 1, we employ base learning rate η_0 of 0.3 and batch size of 512. We came up with this learning rate by performing grid search $\eta_0 \in [1.0, 0.5, 0.1, 0.05, 0.01, 0.005, 0.001, 0.0005, 0.0001]$ and choosing the one which attains the best eval accuracy. To obtain the optimal base learning rate of 0.3, we again ran a grid search $\eta_0 \in [0.1, 0.2, 0.3, 0.4, 0.5]$. Our warmup baseline employs the standard practice of linear warm-up for 3125 steps and polynomial decay of $\eta_t = \eta_0 * (1 - \frac{t}{T})$ for the rest of the steps. For a fair comparison with the baseline, we start the polynomial decay after 3125 steps regardless of how many steps were used for burn-in. We set the learning rate to the constant base learning rate value of 0.3 for BERT-Large for steps in between burn-in and polynomial decay. For the models trained using burn-in, instead of linear warm-up, we set the base learning rate η_0 to zero for the amount of steps we run the burn-in for. As shown in Fig 1, we perform a grid search across $[0, 250, 500, 1000, 2000, 3125]$ steps for the initial burn-in period and compare their results. Notice that, without warm-up or burn-in (i.e. 0 burn-in steps), we see very poor performance. Further, this poor performance persists for smaller values of burn-in, but with large enough burn-in values we suddenly see very good results. Once this critical level of burn-in is achieved, we continue to see small gains as the amount of burn-in is increased. This suggests that without burn-in or warm-up, the first iterations are taking extremely poor steps in the non-convex landscape that may destroy some favorable properties of the initialization point. With enough burn-in, each step is well-aligned with the true gradient and so we may avoid this bad behavior. Further, we see that employing burn-in recovers almost all of the advantage gained

from employing warm-up (final accuracy of 62.95% with burn-in vs. 63.45% with warm-up). These results lend empirical support to our theoretical intuition that the initial steps of the optimization may be unhelpful, or even hurtful.

For the LAMB optimizer, we reproduce all the hyperparameters from You et al. [40]. Most importantly, we train all our models using the base learning rate η_0 of 6.25e-4 and batch size of 512. As shown in Fig 1, again the warmup baseline employs the standard practice of linear warm-up for 3125 steps [40] and polynomial decay of $\eta_t = \eta_0 * (1 - \frac{t}{T})$ for the rest of the steps. Again, for a fair comparison with the baseline, we start the polynomial decay after 3125 steps regardless of how many steps were used for burn-in. We set the learning rate to the constant base learning rate value of 6.25e-4 for steps in between burn-in and polynomial decay. Note that when using the LAMB optimizer, warm-up baseline does considerably better than models where burn-in was employed. However, notice that the gap between employing warm-up and using *neither* warm-up nor burn-in is considerably smaller than with Algorithm 1, for which not using either technique results in extremely poor accuracy. This suggests that the additional heuristics employed by the LAMB optimizer over the simple normalized SGD approach of Algorithm 1 are providing some non-trivial robustness that remains to be understood. Further, there is still a noticeable gap between the linear learning-rate increase in warm-up and the coarser burn-in policy. It remains to be seen exactly what mechanism this smoother policy employs to aid optimization.

5 Conclusion

We have provided algorithms based on normalized SGD with momentum and gradient clipping that provide high probability-rates for finding critical points of non-convex objectives with heavy-tailed gradients. We also show that as a by-product of our high-probability bounds, the algorithms provide a pleasing result for the *last iterate* of the optimization procedure. Specifically, the last iterate has function value smaller than the function value of a critical point.

Finally, we identify a critical “burn-in” phase early in optimization in which the optimizer may be making poor decisions. Empirically, we show that keeping the learning rate zero for the first iterations indeed does help training compared to simply running the optimizer with no special treatment during this phase. It is an interesting question to precisely identify what other advantages are provided by the increasing schedule used by warm-up during this phase.

Acknowledgments: The authors thank the anonymous reviewers for many helpful suggestions. AC was a visiting researcher at Google when this work was completed.

References

- [1] Allen-Zhu, Z. Natasha 2: Faster non-convex optimization than sgd. In *Advances in neural information processing systems*, pp. 2675–2686, 2018.
- [2] Arjevani, Y., Carmon, Y., Duchi, J. C., Foster, D. J., Srebro, N., and Woodworth, B. Lower bounds for non-convex stochastic optimization. *arXiv preprint arXiv:1912.02365*, 2019.
- [3] Arjevani, Y., Carmon, Y., Duchi, J. C., Foster, D. J., Sekhari, A., and Sridharan, K. Second-order information in non-convex stochastic optimization: Power and limitations. In *Conference on Learning Theory*, pp. 242–299, 2020.
- [4] Arnold, S. M., Manzagol, P.-A., Babanezhad, R., Mitliagkas, I., and Roux, N. L. Reducing the variance in online optimization by transporting past gradients. *arXiv preprint arXiv:1906.03532*, 2019.
- [5] Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D. M., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., and Amodei, D. Language models are few-shot learners, 2020.
- [6] Bubeck, S., Cesa-Bianchi, N., and Lugosi, G. Bandits with heavy tail. *IEEE Transactions on Information Theory*, 59(11):7711–7717, 2013.

- [7] Chen, M. X., Firat, O., Bapna, A., Johnson, M., Macherey, W., Foster, G., Jones, L., Schuster, M., Shazeer, N., Parmar, N., and et al. The best of both worlds: Combining recent advances in neural machine translation. *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2018. doi: 10.18653/v1/p18-1008. URL <http://dx.doi.org/10.18653/v1/P18-1008>.
- [8] Cutkosky, A. *Algorithms and Lower Bounds for Parameter-free Online Learning*. PhD thesis, Stanford University, 2018.
- [9] Cutkosky, A. and Mehta, H. Momentum improves normalized sgd. *arXiv preprint arXiv:2002.03305*, 2020.
- [10] Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 4171–4186, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1423. URL <https://www.aclweb.org/anthology/N19-1423>.
- [11] Duchi, J., Hazan, E., and Singer, Y. Adaptive subgradient methods for online learning and stochastic optimization. In *Conference on Learning Theory (COLT)*, 2010.
- [12] Fang, C., Lin, Z., and Zhang, T. Sharp analysis for nonconvex sgd escaping from saddle points. In *Conference on Learning Theory*, pp. 1192–1234, 2019.
- [13] Ghadimi, S. and Lan, G. Stochastic first-and zeroth-order methods for nonconvex stochastic programming. *SIAM Journal on Optimization*, 23(4):2341–2368, 2013.
- [14] Gorbunov, E., Danilova, M., and Gasnikov, A. Stochastic optimization with heavy-tailed noise via accelerated gradient clipping. In *Neural Information Processing Systems*, 2020.
- [15] Goyal, P., Dollár, P., Girshick, R., Noordhuis, P., Wesolowski, L., Kyrola, A., Tulloch, A., Jia, Y., and He, K. Accurate, large minibatch sgd: Training imagenet in 1 hour, 2017.
- [16] He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778, 2015.
- [17] Howard, S. R., Ramdas, A., McAuliffe, J., and Sekhon, J. Time-uniform, nonparametric, nonasymptotic confidence sequences. *arXiv preprint arXiv:1810.08240*, 2018.
- [18] Huang, X. S., Perez, F., Ba, J., and Volkovs, M. Improving transformer optimization through better initialization. In III, H. D. and Singh, A. (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 4475–4483. PMLR, 13–18 Jul 2020. URL <http://proceedings.mlr.press/v119/huang20f.html>.
- [19] Kingma, D. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [20] Li, X. and Orabona, F. On the convergence of stochastic gradient descent with adaptive stepsizes. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 983–992. PMLR, 2019.
- [21] Li, X. and Orabona, F. A high probability analysis of adaptive sgd with momentum. *arXiv preprint arXiv:2007.14294*, 2020.
- [22] Liu, L., Jiang, H., He, P., Chen, W., Liu, X., Gao, J., and Han, J. On the variance of the adaptive learning rate and beyond. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=rkgz2aEKDr>.
- [23] Lugosi, G. and Mendelson, S. Mean estimation and regression under heavy-tailed distributions: A survey. *Foundations of Computational Mathematics*, 19(5):1145–1190, 2019.
- [24] Minsker, S. On some extensions of bernstein’s inequality for self-adjoint operators. *arXiv preprint arXiv:1112.5448*, 2011.

- [25] Nguyen, T. Q. and Salazar, J. Transformers without tears: Improving the normalization of self-attention. *ArXiv*, abs/1910.05895, 2019.
- [26] Pascanu, R., Mikolov, T., and Bengio, Y. Understanding the exploding gradient problem. *CoRR*, abs/1211.5063, 2(417):1, 2012.
- [27] Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., and Liu, P. J. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21(140):1–67, 2020. URL <http://jmlr.org/papers/v21/20-074.html>.
- [28] Reddi, S. J., Kale, S., and Kumar, S. On the convergence of adam and beyond. In *International Conference on Learning Representations*, 2018.
- [29] Scaman, K. and Malherbe, C. Robustness analysis of non-convex stochastic gradient descent using biased expectations. *Advances in Neural Information Processing Systems*, 33, 2020.
- [30] Şimşekli, U., Gürbüzbalaban, M., Nguyen, T. H., Richard, G., and Sagun, L. On the heavy-tailed theory of stochastic gradient descent for deep neural networks. *arXiv preprint arXiv:1912.00018*, 2019.
- [31] Simsekli, U., Sagun, L., and Gurbuzbalaban, M. A tail-index analysis of stochastic gradient noise in deep neural networks. In *International Conference on Machine Learning*, pp. 5827–5837. PMLR, 2019.
- [32] Tripuraneni, N., Stern, M., Jin, C., Regier, J., and Jordan, M. I. Stochastic cubic regularization for fast nonconvex optimization. In *Advances in neural information processing systems*, pp. 2899–2908, 2018.
- [33] Tropp, J. Freedman’s inequality for matrix martingales. *Electronic Communications in Probability*, 16:262–270, 2011.
- [34] Tropp, J. A. An introduction to matrix concentration inequalities. *arXiv preprint arXiv:1501.01571*, 2015.
- [35] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. *Advances in Neural Information Processing Systems*, 30:5998–6008, 2017.
- [36] Wang, H., Gürbüzbalaban, M., Zhu, L., Şimşekli, U., and Erdogdu, M. A. Convergence rates of stochastic gradient descent under infinite noise variance. *arXiv preprint arXiv:2102.10346*, 2021.
- [37] Wang, Q., Li, B., Xiao, T., Zhu, J., Li, C., Wong, D. F., and Chao, L. S. Learning deep transformer models for machine translation. *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 2019. doi: 10.18653/v1/p19-1176. URL <http://dx.doi.org/10.18653/v1/P19-1176>.
- [38] Ward, R., Wu, X., and Bottou, L. Adagrad stepsizes: Sharp convergence over nonconvex landscapes. In *International Conference on Machine Learning*, pp. 6677–6686. PMLR, 2019.
- [39] You, Y., Gitman, I., and Ginsburg, B. Large batch training of convolutional networks. *arXiv preprint arXiv:1708.03888*, 2017.
- [40] You, Y., Li, J., Reddi, S., Hseu, J., Kumar, S., Bhojanapalli, S., Song, X., Demmel, J., Keutzer, K., and Hsieh, C.-J. 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>.
- [41] Zhang, B., Jin, J., Fang, C., and Wang, L. Improved analysis of clipping algorithms for non-convex optimization. *arXiv preprint arXiv:2010.02519*, 2020.
- [42] Zhang, J., He, T., Sra, S., and Jadbabaie, A. Why gradient clipping accelerates training: A theoretical justification for adaptivity. In *International Conference on Learning Representations*, 2019.

- [43] Zhang, J., Karimireddy, S. P., Veit, A., Kim, S., Reddi, S., Kumar, S., and Sra, S. Why are adaptive methods good for attention models? *Advances in Neural Information Processing Systems*, 33, 2020.

A Proofs for Section 3

In this section we provide the missing proofs of Theorems 5 and 6.

Theorem 5. Suppose $\mathbb{E}_z[\|\nabla f(\bar{w}, z)\|^p] \leq G^p$ for all \bar{w} for some G . Suppose F is L -smooth and ρ -second-order smooth. Set $\beta = 1 - \alpha$, $\alpha = \frac{b}{T^{5p-3}}$ and $\eta = \frac{a}{T^{\frac{3p-1}{5p-3}}}$ for arbitrary constants b and c satisfying $\alpha \leq 1$, and set $\tau = \frac{G}{\alpha^{1/p}}$. Then with probability at least $1 - \delta$:

$$\frac{1}{T} \sum_{t=1}^T \|\nabla F(\bar{w}_t)\|_* \leq O\left(\frac{\log(T/\delta)}{T^{\frac{2p-2}{5p-3}}}\right)$$

where the big-Oh hides constant that depend on L , ρ , G , b , s , C , and $F(\bar{w}_1) - F(\bar{w}_{T+1})$, but not T or δ .

Proof. The start of our proof is very similar to that of Theorem 3 in [9]. Similar to the proof of Theorem 2, define $\hat{\epsilon}_t = \bar{m}_t - \nabla F(\bar{w}_t)$ and $\epsilon_t = \hat{g}_t^{\text{clip}} - \nabla F(\bar{x}_t)$. Also (and as a departure from Theorem 2), define $S(a, b) = \nabla F(a) - \nabla F(b) - \nabla^2 F(b)(a - b)$. Then we have a recursion for all $t \geq 1$:

$$\begin{aligned} \bar{m}_{t+1} &= (1 - \alpha)(\nabla F(\bar{w}_t) + \hat{\epsilon}_t) + \alpha \hat{g}_{t+1}^{\text{clip}} \\ &= (1 - \alpha)(\nabla F(\bar{w}_{t+1}) + \nabla^2 F(\bar{w}_{t+1})(\bar{w}_t - \bar{w}_{t+1})) + (1 - \alpha)(S(\bar{w}_t, \bar{w}_{t+1}) + \hat{\epsilon}_t) \\ &\quad + \alpha(\nabla F(\bar{w}_{t+1}) + \nabla^2 F(\bar{w}_{t+1})(\bar{x}_{t+1} - \bar{w}_{t+1}) + \alpha(S(\bar{x}_{t+1}, \bar{w}_{t+1}) + \epsilon_{t+1})) \end{aligned}$$

Now, observe that by definition of \bar{x}_{t+1} , the hessian-dependent terms cancel:

$$\begin{aligned} &= \nabla F(\bar{w}_{t+1}) + (1 - \alpha)(S(\bar{w}_t, \bar{w}_{t+1}) + \hat{\epsilon}_t) + \alpha(S(\bar{x}_{t+1}, \bar{w}_{t+1}) + \epsilon_{t+1}) \\ \hat{\epsilon}_{t+1} &= (1 - \alpha)(S(\bar{w}_t, \bar{w}_{t+1}) + \hat{\epsilon}_t) + \alpha(S(\bar{x}_{t+1}, \bar{w}_{t+1}) + \epsilon_{t+1}) \end{aligned}$$

Further, as in the proof of Theorem 2, we set $\bar{w}_0 = \bar{w}_1$ and notice $\hat{\epsilon}_0 = -\nabla F(\bar{w}_1)$ so that the above recursion holds for $t = 0$ as well. Then we again unwind the recursion:

$$\begin{aligned} \hat{\epsilon}_{t+1} &= (1 - \alpha)^{t+1} \hat{\epsilon}_0 + \alpha \sum_{t'=0}^t (1 - \alpha)^{t'} \epsilon_{t+1-t'} + (1 - \alpha) \sum_{t'=0}^t (1 - \alpha)^{t'} S(\bar{w}_{t-t'}, \bar{w}_{t-t'+1}) \\ &\quad + \alpha \sum_{t'=0}^t (1 - \alpha)^{t'} S(\bar{x}_{t-t'+1}, \bar{w}_{t-t'+1}) \\ \|\hat{\epsilon}_{t+1}\|_* &\leq (1 - \alpha)^{t+1} \|\hat{\epsilon}_0\|_* + \alpha \left\| \sum_{t'=0}^t (1 - \alpha)^{t'} \epsilon_{t+1-t'} \right\|_* + (1 - \alpha) \left\| \sum_{t'=0}^t (1 - \alpha)^{t'} S(\bar{w}_{t-t'}, \bar{w}_{t-t'+1}) \right\|_* \\ &\quad + \alpha \left\| \sum_{t'=0}^t (1 - \alpha)^{t'} S(\bar{x}_{t-t'+1}, \bar{w}_{t-t'+1}) \right\|_* \end{aligned}$$

Now, by second-order smoothness, we have $\|S(a, b)\|_* \leq \frac{\rho}{2} \|a - b\|^2$ (see Proposition 17) and further $\frac{(1-\alpha)^2}{\alpha} + 1 - \alpha \leq 2\frac{1-\alpha}{\alpha}$, so that combining the last two sums yields:

$$\begin{aligned} &\leq (1 - \alpha)^{t+1} \|\hat{\epsilon}_0\|_* + \alpha \left\| \sum_{t'=0}^t (1 - \alpha)^{t'} \epsilon_{t+1-t'} \right\|_* + \frac{(1 - \alpha)\rho\eta^2}{\alpha} \sum_{t'=0}^t (1 - \alpha)^{t'} \\ &\leq (1 - \alpha)^{t+1} \|\hat{\epsilon}_0\|_* + \alpha \left\| \sum_{t'=0}^t (1 - \alpha)^{t'} \epsilon_{t+1-t'} \right\|_* + \frac{(1 - \alpha)\rho\eta^2}{\alpha^2} \end{aligned}$$

Now, by exactly the same argument as in Theorem 2, defining $D_\delta = \max(1, \log(1/\delta))$ we have with probability at least $1 - \delta/T$:

$$\|\hat{\epsilon}_{t+1}\|_* \leq (1 - \alpha)^{t+1} G + \frac{(1 - \alpha)\rho\eta^2}{\alpha^2} + 10C\alpha\tau D_\delta + 4\sqrt{\alpha C G^p \tau^{2-p} D_\delta} + \frac{G^p}{\tau^{p-1}}$$

Then, with $\tau = \frac{G}{\alpha^{1/p}}$, this becomes:

$$\begin{aligned} &\leq (1-\alpha)^{t+1}G + \frac{(1-\alpha)\rho\eta^2}{\alpha^2} + G\alpha^{1-1/p} \left[10CD_\delta + 4\sqrt{CD_\delta} + 1 \right] \\ &= (1-\alpha)^{t+1}G + \frac{(1-\alpha)\rho\eta^2}{\alpha^2} + G\alpha^{1-1/p}K \end{aligned} \quad (6)$$

where we again define the constant $K = 10D_\delta + 4\sqrt{CD_\delta} + 1$.

Now, we have from Lemma 1:

$$\sum_{t=1}^T \|\nabla F(\vec{w}_t)\|_* \leq \frac{F(\vec{w}_1) - F(\vec{w}_T)}{\eta} + \frac{LT\eta}{2} + 2 \sum_{t=1}^T \|\hat{\epsilon}_t\|_*$$

so with probability at least $1 - \delta$:

$$\begin{aligned} &\leq \frac{F(\vec{w}_1) - F(\vec{w}_T)}{\eta} + \frac{LT\eta}{2} + 2 \sum_{t=1}^T \left[(1-\alpha)^{t+1}G + \frac{(1-\alpha)\rho\eta^2}{\alpha^2} + G\alpha^{1-1/p}K \right] \\ &\leq \frac{F(\vec{w}_1) - F(\vec{w}_T)}{\eta} + \frac{LT\eta}{2} + \frac{2G(1-\alpha)}{\alpha} + \frac{2(1-\alpha)T\rho\eta^2}{\alpha^2} + 2GT\alpha^{1-1/p}K \end{aligned}$$

Now set $\alpha = \frac{b}{T^{\frac{2p}{5p-3}}}$ and $\eta = \frac{s}{T^{\frac{3p-1}{5p-3}}}$. This yields with probability at least $1 - \delta$:

$$\frac{1}{T} \sum_{t=1}^T \|\nabla F(\vec{w}_t)\|_* \leq O\left(T^{\frac{2-2p}{5p-3}} \log(T/\delta)\right)$$

□

Theorem 6. *Under the assumptions of Theorem 5, define the constants:*

$$\begin{aligned} K &= 10 \max(1, \log(3T/\delta)) + 4\sqrt{C \max(1, \log(3T/\delta))} + 1 \\ Z &= \frac{\rho s^2}{b^2} + GKb^{\frac{p-1}{p}} \end{aligned}$$

Then, with probability at least $1 - \delta$, if

$$t \geq \mathcal{T} = \frac{T^{\frac{2p}{5p-3}}}{b} \left(\frac{2p-2}{5p-3} \log(T) + \log(G) - \log(Z) \right)$$

Then we have

$$\|\vec{m}_t - \nabla F(\vec{x}_t)\| \leq \frac{2Z}{T^{\frac{2p-2}{5p-3}}}$$

And so long as

$$\|\vec{m}_t\|_* \geq 2 \left(\frac{6Z}{T^{\frac{2p-2}{5p-3}}} + \frac{Ls}{2T^{\frac{3p-1}{5p-3}}} \right) \quad (5)$$

we have $F(\vec{w}_{t+1}) < F(\vec{w}_t) - \frac{\eta}{2}\|\vec{m}_t\|$. Moreover, if (5) is ever not satisfied, we must have

$$\|\nabla F(\vec{w}_t)\|_* \leq \frac{14Z}{T^{\frac{2p-2}{5p-3}}} + \frac{Ls}{T^{\frac{3p-1}{5p-3}}}$$

Finally, if t is the iteration with smallest value of $\|\vec{m}_t\|$ such that $t \geq \mathcal{T}$, we have with probability at least $1 - \delta$:

$$\|\nabla F(\vec{w}_t)\|_* \leq O\left(\frac{\log(T/\delta)}{T^{\frac{2p-2}{5p-3}}}\right)$$

Proof. Use the settings $\eta = \frac{s}{T^{\frac{3p-1}{5p-3}}}$ and $\alpha = \frac{b}{T^{\frac{2p}{5p-3}}}$ and notice that from equation (6) that with probability $1 - \delta$, for all t :

$$\begin{aligned} \|\hat{\epsilon}_{t+1}\|_* &\leq (1 - \alpha)^t G + \frac{(1 - \alpha)\rho\eta^2}{\alpha^2} + G\alpha^{1-1/p}K \\ &= (1 - \alpha)^t G + \frac{\rho s^2}{b^2 T^{\frac{2p-2}{5p-3}}} + \frac{GKb^{\frac{p-1}{p}}}{T^{\frac{2p-2}{5p-3}}} \\ &= (1 - \alpha)^t G + \frac{Z}{T^{\frac{2p-2}{5p-3}}} \end{aligned}$$

Again, we have the identity:

$$\begin{aligned} \log(1 - \alpha) &\leq -\alpha \\ \log(G(1 - \alpha)^t) &\leq -t\alpha + \log(G) \end{aligned}$$

Therefore, if t satisfies:

$$t \geq \frac{T^{\frac{2p}{5p-3}}}{b} \left(\frac{2p-2}{5p-3} \log(T) + \log(G) - \log(Z) \right)$$

we have with probability at least $1 - \delta$:

$$\|\hat{\epsilon}_{t+1}\|_* \leq \frac{2Z}{T^{\frac{2p-2}{5p-3}}}$$

Next, again from Lemma 1:

$$\begin{aligned} F(\vec{w}_{t+1}) &\leq F(\vec{w}_t) - \eta \|\nabla F(\vec{w}_t)\|_* + 2\eta \|\hat{\epsilon}_t\|_* + \frac{L\eta^2}{2} \\ &\leq F(\vec{w}_t) - \eta \|\vec{m}_t\| + 3\eta \|\hat{\epsilon}_t\|_* + \frac{L\eta^2}{2} \\ &\leq F(\vec{w}_t) - \eta \|\vec{m}_t\| + \eta \left(\frac{6Z}{T^{\frac{2p-2}{5p-3}}} + \frac{Ls}{2T^{\frac{3p-1}{5p-3}}} \right) \end{aligned}$$

Now, analogously to Theorem 3, we argue that if $\|\vec{m}_t\| \geq 2 \left(\frac{6Z}{T^{\frac{2p-2}{5p-3}}} + \frac{Ls}{2T^{\frac{3p-1}{5p-3}}} \right)$, then $F(\vec{w}_{t+1}) \leq F(\vec{w}_t) - \frac{\eta}{2} \|\vec{m}_t\|$.

The last part of the Theorem follows from an identical case-work argument to that of Theorem 3. \square

B A Concentration Bound for Truncated Random Vectors

A key tool in proving our bounds is the following result, which provides a dimension-free concentration bound analogous to the concentration bounds for scalar random variables that can be found in [23, 6]. We present the simpler results for Hilbert space first, and then generalize to the Banach space case.

Lemma 8. *Suppose X_1, \dots, X_T are random vectors in a Hilbert space adapted to a filtration $\mathcal{F}_1, \mathcal{F}_2, \dots$ with means $\mathbb{E}_{t-1}[X_t] = \mu_t$. Suppose there is some $p \in (1, 2]$ such that for all t there is some $G_t < \infty$ such that $\mathbb{E}[\|X_t\|^p] \leq G_t^p$. Let b_1, \dots, b_T be fixed constants, with $B = \max_t b_t$ such that $B \leq 1$. For some $\tau > 0$, let $\hat{X}_t = \frac{X_t}{\|X_t\|} \min(\tau, \|X_t\|)$. Then with probability at least $1 - \delta$:*

$$\left\| \sum_{t=1}^T b_t (\hat{X}_t - \mu_t) \right\| \leq 4B\tau \log(3/\delta) + \sum_{t=1}^T \frac{b_t G_t^p}{\tau^{p-1}} + 2 \sqrt{\sum_{t=1}^T b_t^2 G_t^p \tau^{2-p} \max(1, \log(3/\delta))}$$

The proof has two main steps: first, we bound the bias and the variance of the truncated random vectors in terms of the truncation threshold. Then, we use a dimension-free Feedman-style bound to show that the average of these truncated vectors concentrates about the mean of the original un-truncated vectors.

Proof. First, we will bound $\|\mathbb{E}[\hat{X}_t] - \mu_t\|$, following analysis similar to that present in [43]:

$$\begin{aligned}\|\mathbb{E}[\hat{X}_t] - \mu_t\| &= \|\mathbb{E}[\hat{X}_t - X_t]\| \\ &\leq \mathbb{E}[\|\hat{X}_t - X_t\|] \\ &\leq \mathbb{E}[\|X_t\| \mathbf{1}[\|X_t\| \geq \tau]] \\ &\leq \mathbb{E}[\|X_t\|^p / \tau^{p-1}] \\ &\leq \frac{G_t^p}{\tau^{p-1}}\end{aligned}$$

Next, we bound the variance $\mathbb{E}[\|\hat{X}_t - \mathbb{E}[\hat{X}_t]\|^2]$:

$$\begin{aligned}\mathbb{E}[\|\hat{X}_t - \mathbb{E}[\hat{X}_t]\|^2] &\leq \mathbb{E}[\|\hat{X}_t\|^2] \\ &\leq \mathbb{E}[\|X_t\|^p \tau^{2-p}] \\ &\leq G_t^p \tau^{2-p}\end{aligned}$$

Now, invoke a vector-valued Feedman-inequality (Lemma ??), observing that $\|\hat{X}_t\| \leq \tau$ always so that $\|b_t(\hat{X}_t - \mathbb{E}[\hat{X}_t])\| \leq 2b_t\tau \leq 2B\tau$ where $B = \max_t b_t$. Thus, with probability $1 - \delta$:

$$\left\| \sum_{t=1}^T b_t(\hat{X}_t - \mathbb{E}[\hat{X}_t]) \right\| \leq 6B\tau \max(1, \log(3/\delta)) + 3\sqrt{\sum_{t=1}^T b_t^2 G_t^p \tau^{2-p} \max(1, \log(3/\delta))}$$

Therefore by triangle inequality:

$$\left\| \sum_{t=1}^T b_t(\hat{X}_t - \mu_t) \right\| \leq 6B\tau \max(1, \log(3/\delta)) + \sum_{t=1}^T \frac{b_t G_t^p}{\tau^{p-1}} + 3\sqrt{\sum_{t=1}^T b_t^2 G_t^p \tau^{2-p} \max(1, \log(3/\delta))}$$

□

Now, we present the generalization to Banach spaces.

Lemma 9. *Suppose X_1, \dots, X_T are random vectors in some Banach space \mathcal{B} satisfying (7) with $p \leq 2$ and $C \geq 1$ adapted to a filtration F_1, F_2, \dots . Let $\mathbb{E}_{t-1}[X_t] = \mu_t$. Suppose there is some $p \in (1, 2]$ such that for all t there is some $G_t < \infty$ such that $\mathbb{E}[\|X_t\|^p] \leq G_t^p$. Let b_1, \dots, b_T be fixed constants, with $B = \max_t b_t$ such that $B \leq 1$. For some $\tau > 0$, let $\hat{X}_t = \frac{X_t}{\|X_t\|} \min(\tau, \|X_t\|)$. Then with probability at least $1 - \delta$:*

$$\left\| \sum_{t=1}^T b_t(\hat{X}_t - \mu_t) \right\| \leq 10CB\tau \max(1, \log(3/\delta)) + \sum_{t=1}^T \frac{b_t G_t^p}{\tau^{p-1}} + 4 \left(C \sum_{t=1}^T b_t^p G_t^p \tau^{p-p} \right)^{1/p} \sqrt{\max(1, \log(3/\delta))}$$

The proof follows the same structure as the previous result for Hilbert spaces.

Proof. Note that the bound

$$\|\mathbb{E}[\hat{X}_t] - \mu_t\| \leq \frac{G_t^p}{\tau^{p-1}}$$

continues to hold as the arguments apply without modification in Banach spaces. For the variance, we have an analogous statement:

$$\begin{aligned}\mathbb{E}[\|\hat{X}_t - \mathbb{E}[\hat{X}_t]\|^p] &\leq \mathbb{E}[\|\hat{X}_t\|^p] \\ &\leq \mathbb{E}[\|X_t\|^p \tau^{p-p}] \\ &\leq G_t^p \tau^{p-p}\end{aligned}$$

Now, invoke a vector-valued Freedman-inequality (Lemma ??), observing that $\|\hat{X}_t\| \leq \tau$ always so that $\|b_t(\hat{X}_t - \mathbb{E}[\hat{\mu}_t])\| \leq 2b_t\tau \leq 2B\tau$ where $B = \max_t b_t$. Thus, with probability $1 - \delta$:

$$\left\| \sum_{t=1}^T b_t(\hat{X}_t - \mathbb{E}[\hat{X}_t]) \right\| \leq 10CB\tau \max(1, \log(3/\delta)) + 4 \left(C \sum_{t=1}^k b_t^p G_t^p \tau^{p-p} \right)^{1/p} \sqrt{\max(1, \log(3/\delta))}$$

Therefore by triangle inequality:

$$\left\| \sum_{t=1}^T b_t(\hat{X}_t - \mu_t) \right\| \leq 10CB\tau \max(1, \log(3/\delta)) + \sum_{t=1}^T \frac{b_t G_t^p}{\tau^{p-1}} + 4 \left(C \sum_{t=1}^k b_t^p G_t^p \tau^{p-p} \right)^{1/p} \sqrt{\max(1, \log(3/\delta))}$$

□

C Dimension-Free Martingale concentration from 1-d concentration

In this section, we show how to obtain a dimension-free Freedman-style martingale concentration bound in a Banach space via a reduction to the ordinary 1-dimensional Freedman inequality. The conversion is based on a construction outlined in [8] for converting 1-dimensional online linear optimization algorithms into dimension-free algorithms. It may be of independent interest as simple alternative proof for such concentration results, as well as providing a possibly more user-friendly result in comparison to advanced bounds available in [34, 24, 17]. However, we emphasize that the results in this section are essentially provided only for completeness: they do *not* improve upon previously available bounds.

Suppose B is a real Banach space whose norm $\|\cdot\|$ is Frechet differentiable and satisfies for some p and C for all x, y :

$$\|x + y\|^p \leq \|x\|^p + \langle \nabla \|x\|^p, y \rangle + C\|y\|^p \quad (7)$$

Notice that if B is a Hilbert space, we can take $p = 2$ and $C = 1$.

Given any sequence of random vectors X_1, \dots, X_T in B , consider the sequence of real numbers s_1, \dots, s_T defined recursively by:

1. $s_0 = 0$
2. If $\sum_{i=1}^{t-1} X_i \neq 0$, then we set:

$$s_t = \text{sign} \left(\sum_{i=1}^{t-1} s_i \right) \frac{\langle \nabla \|\sum_{i=1}^{t-1} X_i\|^p, X_t \rangle}{p \|\sum_{i=1}^{t-1} X_i\|^{p-1}},$$

where $\text{sign}(x) = 1$ if $x \geq 0$, -1 if $x < 0$ and 0 if $x = 0$.

3. If $\sum_{i=1}^{t-1} X_i = 0$, set $s_t = 0$.

The critical property of this sequence is the following, proved in [8] Theorem 5.3. We reproduce the proof for completeness below as we have mildly tightened the constants:

Lemma 10. *Suppose B , s_t and X_t are as described above and $p \geq 1$. Then $|s_t| \leq \|X_t\|$ for all t , and*

$$\left\| \sum_{t=1}^T X_t \right\| \leq \left| \sum_{t=1}^T s_t \right| + \left(\max_{t \leq T} \|X_t\|^p + C \sum_{t=1}^T \|X_t\|^p \right)^{1/p}$$

Proof. For the bound on s_t , observe that since $x \mapsto \|x\|$ is by definition 1-Lipschitz, $\|(\nabla \|x\|)\|_* \leq 1$. Therefore we have

$$\begin{aligned} |s_t| &\leq \frac{|\langle \nabla \|\sum_{i=1}^{t-1} X_i\|^p, X_t \rangle|}{p \|\sum_{i=1}^{t-1} X_i\|^{p-1}} \\ &\leq \langle \nabla \|\sum_{i=1}^{t-1} X_i\|, X_t \rangle \\ &\leq \|X_t\| \end{aligned}$$

Now, for the second statement we proceed by induction. Clearly the statement holds for $T = 1$. Suppose now that

$$\left| \sum_{t=1}^{T-1} X_t \right| \leq \left\| \sum_{t=1}^{T-1} s_t \right\| + \left(\max_{t \leq T-1} \|X_t\|^p + C \sum_{t=1}^{T-1} \|X_t\|^p \right)^{1/p}$$

Now, first observe that if $\sum_{i=1}^{T-1} X_i = 0$, then $\|\sum_{i=1}^T X_i\| = \|X_T\| \leq \max_{t \leq T} \|X_t\|$ and so the statement holds. Let us now assume $\sum_{i=1}^{T-1} X_i \neq 0$. Define $X_{1:t} = \sum_{i=1}^t X_i$. Then we have:

$$\left\| \sum_{i=1}^T X_i \right\| = (\|X_{1:T-1} + X_T\|^p)^{1/p} \leq (\|X_{1:T-1}\|^p + \langle \nabla \|X_{1:T-1}\|^p, X_T \rangle + C \|X_T\|^p)^{1/p}$$

Now we consider two cases, either $\|X_{1:T-1}\|^p + \langle \nabla \|X_{1:T-1}\|^p, X_T \rangle \leq 0$ or not. If $\|X_{1:T-1}\|^p + \langle \nabla \|X_{1:T-1}\|^p, X_T \rangle \leq 0$, then we have just shown:

$$\left\| \sum_{i=1}^T X_i \right\| \leq C^{1/p} \|X_T\| \leq \left| \sum_{i=1}^T s_i \right| + \left(\max_{t \leq T} \|X_t\|^p + C \sum_{i=1}^T \|X_t\|^p \right)^{1/p}$$

and so we are done.

Instead, let us suppose $\|X_{1:T-1}\|^p + \langle \nabla \|X_{1:T-1}\|^p, X_T \rangle > 0$. This implies $\|X_{1:T-1}\| + \frac{\langle \nabla \|X_{1:T-1}\|^p, X_T \rangle}{p \|X_{1:T-1}\|^{p-1}} \geq 0$ as well. Now, since the function $x \mapsto x^p$ is convex for any positive x , we have $x^p + pyx^{p-1} \leq (x+y)^p$ for any $x \geq 0$ and $x+y \geq 0$. Thus:

$$\|X_{1:T-1}\|^p + \langle \nabla \|X_{1:T-1}\|^p, X_T \rangle \leq \left(\|X_{1:T-1}\| + \frac{\langle \nabla \|X_{1:T-1}\|^p, X_T \rangle}{p \|X_{1:T-1}\|^{p-1}} \right)^p$$

Putting this together with our induction hypothesis:

$$\begin{aligned} \left\| \sum_{i=1}^T X_i \right\| &\leq \left(\left(\|X_{1:T-1}\| + \frac{\langle \nabla \|X_{1:T-1}\|^p, X_T \rangle}{p \|X_{1:T-1}\|^{p-1}} \right)^p + C \|X_T\|^p \right)^{1/p} \\ &\leq \left(\left(\left| \sum_{i=1}^{T-1} s_i \right| + \left(\max_{i \leq T-1} \|X_i\|^p + C \sum_{i=1}^{T-1} \|X_i\|^p \right)^{1/p} + \frac{\langle \nabla \|X_{1:T-1}\|^p, X_T \rangle}{p \|X_{1:T-1}\|^{p-1}} \right)^p + C \|X_T\|^p \right)^{1/p} \\ &\leq \left[\left(\left| \sum_{i=1}^{T-1} s_i \right| + \frac{\langle \nabla \|X_{1:T-1}\|^p, X_T \rangle}{p \|X_{1:T-1}\|^{p-1}} \right) + \left(\max_{i \leq T-1} \|X_i\|^p + C \sum_{i=1}^{T-1} \|X_i\|^p \right)^{1/p} \right]^p + C \|X_T\|^p \end{aligned}$$

Now, following [8], we observe that for any positive a, b and c , $(a+b)^p - b^p \leq (a+b+c)^p - (b+c)^p$.

Thus, setting $a = \left| \sum_{i=1}^{T-1} s_i \right| + \frac{\langle \nabla \|X_{1:T-1}\|^p, X_T \rangle}{p \|X_{1:T-1}\|^{p-1}}$, $b = \left(\max_{i \leq T-1} \|X_i\|^p + C \sum_{i=1}^{T-1} \|X_i\|^p \right)^{1/p}$

and $b+c = \left(\max_{i \leq T} \|X_i\|^p + C \sum_{i=1}^T \|X_i\|^p \right)^{1/p}$, we obtain:

$$\begin{aligned} \left(\left| \sum_{i=1}^{T-1} s_i \right| + \frac{\langle \nabla \|X_{1:T-1}\|^p, X_T \rangle}{p \|X_{1:T-1}\|^{p-1}} \right) + \left(\max_{i \leq T-1} \|X_i\|^p + C \sum_{i=1}^{T-1} \|X_i\|^p \right)^{1/p} &\leq (a+b+c)^p + b^p - (b+c)^p \\ &= (a+b+c)^p - C \|X_T\|^p \end{aligned}$$

Plugging this identity back into our previous bound on $\left\| \sum_{i=1}^T X_i \right\|$, we have:

$$\begin{aligned} \left\| \sum_{i=1}^T X_i \right\| &\leq \left| \sum_{i=1}^{T-1} s_i \right| + \frac{\langle \nabla \|X_{1:T-1}\|^p, X_T \rangle}{p \|X_{1:T-1}\|^{p-1}} + \left(\max_{i \leq T} \|X_i\|^p + C \sum_{i=1}^T \|X_i\|^p \right)^{1/p} \\ &= \left| \sum_{i=1}^{T-1} s_i \right| + \text{sign} \left(\sum_{i=1}^{T-1} s_i \right) s_T + \left(\max_{i \leq T} \|X_i\|^p + C \sum_{i=1}^T \|X_i\|^p \right)^{1/p} \\ &= \left| \sum_{i=1}^T s_i \right| + \left(\max_{i \leq T} \|X_i\|^p + C \sum_{i=1}^T \|X_i\|^p \right)^{1/p} \end{aligned}$$

□

C.1 Warm-up: Hilbert space

Now, we show how to use Lemma 10 in conjunction with the standard 1-dimensional Freedman inequality to generate concentration bounds in Hilbert spaces. First, we state a (mildly weaker) version of the standard 1-dimensional inequality.

Lemma 11. *[Freedman's inequality] Suppose D_1, D_2, \dots, D_T is a martingale difference sequence adapted to a filtration F_1, F_2, \dots such that $D_i \leq R$ almost surely for all i . Let \mathbb{E}_i indicate expectation conditioned on F_i . Suppose further that for all t with probability 1,*

$$\sigma_t^2 \geq \mathbb{E}_{t-1}[D_t^2]$$

Then with probability at least $1 - \delta$, for all k we have

$$\sum_{t=1}^k D_t \leq \frac{2R \log(1/\delta)}{3} + \sqrt{2 \sum_{t=1}^k \sigma_t^2 \log(1/\delta)}$$

Proof. From the standard inequality (see e.g. [33] Theorem 1.1), we have

$$P \left[\exists k : \sum_{t=1}^k D_t \geq \epsilon \right] \leq \exp \left(- \frac{\epsilon^2/2}{\sum_{t=1}^k \sigma_t^2 + R\epsilon/3} \right)$$

Now, set the RHS equal to δ to obtain:

$$\frac{\epsilon^2/2}{\sum_{t=1}^k \sigma_t^2 + R\epsilon/3} = \log(1/\delta)$$

The result now follows by using the quadratic formula to bound ϵ . □

Now, we are in a position to describe our extension of Freedman's inequality to Hilbert spaces (which satisfy the hypotheses of Lemma 10 with $p = 2$ and $C = 1$):

Lemma 12. *Suppose X_1, \dots, X_T is a martingale difference sequence in a Hilbert space such that $\|X_t\| \leq R$ almost surely for some constant R . Further, assume $\mathbb{E}_{t-1}[\|X_t\|^2] \leq \sigma_t^2$ with probability 1 for some constants σ_t . Then with probability at least $1 - 3\delta$, for all k we have:*

$$\left\| \sum_{t=1}^k X_t \right\| \leq 3R \max(1, \log(1/\delta)) + 3 \sqrt{\sum_{t=1}^k \sigma_t^2 \max(1, \log(1/\delta))}$$

Proof. Define s_t as in Lemma 10. Then with probability 1 we have for all k :

$$\left\| \sum_{t=1}^k X_t \right\| \leq \left| \sum_{t=1}^k s_t \right| + \sqrt{R^2 + \sum_{t=1}^k \|X_t\|^2}$$

Further, notice that s_t is itself a martingale difference sequence, and satisfies $|s_t| \leq \|X_t\| \leq R$. Therefore by Lemma 11, with probability at least $1 - 2\delta$

$$\left| \sum_{t=1}^k s_t \right| \leq \frac{2R \log(1/\delta)}{3} + \sqrt{2 \sum_{t=1}^k \sigma_t^2 \log(1/\delta)}$$

Now, for the second term, set $Z_t = \|X_t\|^2 - \mathbb{E}_{t-1}[\|X_t\|^2]$. Then Z_t is also a martingale difference sequence. Since $\|X_t\|^2 \leq R^2$ with probability 1, $Z_t \leq R^2$ with probability 1. Further,

$$\mathbb{E}_{t-1}[\|Z_t\|^2] \leq \mathbb{E}_{t-1}[\|X_t\|^4] \leq \mathbb{E}_{t-1}[\|X_t\|^2] R^2 \leq \sigma_t^2 R^2$$

Thus, again by Lemma 11, with probability at least $1 - \delta$:

$$\sum_{t=1}^k Z_t \leq \frac{2R^2 \log(1/\delta)}{3} + \sqrt{2 \sum_{t=1}^T R^2 \sigma_t^2 \log(1/\delta)}$$

which implies:

$$\sum_{t=1}^k \|X_t\|^2 \leq \sum_{t=1}^k \sigma_t^2 + \frac{2R^2 \log(1/\delta)}{3} + \sqrt{2 \sum_{t=1}^k R^2 \sigma_t^2 \log(1/\delta)}$$

Applying Young inequality:

$$\begin{aligned} &\leq \sum_{t=1}^k \sigma_t^2 + \frac{2R^2 \log(1/\delta)}{3} + \sqrt{R^4 \log^2(1/\delta) + \left(\sum_{t=1}^k \sigma_t^2\right)^2} \\ &\leq 2 \sum_{t=1}^k \sigma_t^2 + \frac{5R^2 \log(1/\delta)}{3} \end{aligned}$$

Putting everything together, with probability at least $1 - 3\delta$:

$$\begin{aligned} \left\| \sum_{t=1}^k X_t \right\| &\leq \left| \sum_{t=1}^k s_t \right| + \sqrt{R^2 + \sum_{t=1}^k \|X_t\|^2} \\ &\leq \frac{2R(\log(1/\delta) + 3/2)}{3} + \sqrt{2 \sum_{t=1}^k \sigma_t^2 \log(1/\delta)} + \sqrt{2 \sum_{t=1}^k \sigma_t^2 + \frac{5R^2 \log(1/\delta)}{3}} \\ &\leq 3R \max(1, \log(1/\delta)) + 3 \sqrt{\sum_{t=1}^k \sigma_t^2 \max(1, \log(1/\delta))} \end{aligned}$$

□

C.2 Extension to Banach space

Having provided the concentration results in the more familiar Hilbert space setting, now we move to Banach spaces. We will also need the following useful observation:

Lemma 13. *For any $0 < p \leq q$ and any positive x_1, \dots, x_T ,*

$$\left(\sum_{t=1}^T x_t^q \right)^{1/q} \leq \left(\sum_{t=1}^T x_t^p \right)^{1/p}$$

Proof. We differentiate the expression $\left(\sum_{t=1}^T x_t^p \right)^{1/p}$ with respect to p and show that the derivative is always negative, which suffices to prove the Lemma.

$$\begin{aligned}
\frac{d}{dp} \left(\sum_{t=1}^T x_t^p \right)^{1/p} &= \frac{d}{dp} \exp \left\{ \frac{1}{p} \log \left[\sum_{t=1}^T \exp(p \log(x_t)) \right] \right\} \\
&= \left(\sum_{t=1}^T x_t^p \right)^{1/p} \left(\frac{-\log \left[\sum_{t=1}^T \exp(p \log(x_t)) \right]}{p^2} + \frac{\sum_{t=1}^T \exp(p \log(x_t)) \log(x_t)}{p \sum_{t=1}^T \exp(p \log(x_t))} \right) \\
&= \left(\sum_{t=1}^T x_t^p \right)^{1/p} \left(\frac{-\log \left[\sum_{t=1}^T x_t^p \right]}{p^2} + \frac{\sum_{t=1}^T x_t^p \log(x_t)}{p \sum_{t=1}^T x_t^p} \right) \\
&= \frac{\left(\sum_{t=1}^T x_t^p \right)^{1/p}}{p^2} \left(-\log \left[\sum_{t=1}^T x_t^p \right] + \frac{\sum_{t=1}^T x_t^p \log(x_t^p)}{\sum_{t=1}^T x_t^p} \right) \\
&= \frac{\left(\sum_{t=1}^T x_t^p \right)^{1/p-1}}{p^2} \left(\sum_{t=1}^T x_t^p \log(x_t^p) - \log \left[\sum_{t=1}^T x_t^p \right] \left(\sum_{t=1}^T x_t^p \right) \right)
\end{aligned}$$

since \log is an increasing function and $x_t > 0$:

$$\begin{aligned}
&\leq \frac{\left(\sum_{t=1}^T x_t^p \right)^{1/p-1}}{p^2} \left(\sum_{t=1}^T x_t^p \log \left[\sum_{t=1}^T x_t^p \right] - \log \left[\sum_{t=1}^T x_t^p \right] \left(\sum_{t=1}^T x_t^p \right) \right) \\
&= 0
\end{aligned}$$

□

Lemma 14. *Suppose X_1, \dots, X_T is a martingale difference sequence in a Banach space satisfying (7) for $p \in (1, 2]$ such that $\|X_t\| \leq R$ almost surely for some constant R . Further, assume $\mathbb{E}_{t-1}[\|X_t\|^2] \leq \sigma_t^2$ with probability 1 for some constants σ_t . Then with probability at least $1 - 3\delta$, for all k we have:*

$$\left\| \sum_{t=1}^k X_t \right\| \leq \frac{2R(\log(1/\delta) + 3/2)}{3} + \sqrt{2 \sum_{t=1}^k \sigma_t^2 \log(1/\delta)} + \left(2C \sum_{t=1}^k \sigma_t^p + \frac{7CR^p \log(1/\delta)}{3} \right)^{1/p}$$

Additionally, if $C \geq 1$ and $p \leq 2$:

$$5CR \max(1, \log(3/\delta)) + 4 \left(C \sum_{t=1}^k \sigma_t^p \right)^{1/p} \sqrt{\max(1, \log(3/\delta))}$$

Proof. First, note that by Jensen, we have $\mathbb{E}_{t-1}[\|X_t\|^p] \leq \mathbb{E}_{t-1}[\|X_t\|^{2p/2}] \leq \sigma_t^p$.

Define s_t as in Lemma 10. Then with probability 1 we have for all k :

$$\left\| \sum_{t=1}^k X_t \right\| \leq \left| \sum_{t=1}^k s_t \right| + \left(R^p + C \sum_{t=1}^k \|X_t\|^p \right)^{1/p}$$

Further, notice that s_t is itself a martingale difference sequence, and satisfies $|s_t| \leq \|X_t\| \leq R$. Therefore by Lemma 11, with probability at least $1 - 2\delta$

$$\left| \sum_{t=1}^k s_t \right| \leq \frac{2R \log(1/\delta)}{3} + \sqrt{2 \sum_{t=1}^k \sigma_t^2 \log(1/\delta)}$$

Now, for the second term, set $Z_t = \|X_t\|^p - \mathbb{E}_{t-1}[\|X_t\|^p]$. Then Z_t is also a martingale difference sequence. Since $\|X_t\|^p \leq R^p$ with probability 1, $Z_t \leq 2R^p$ with probability 1. Further,

$$\mathbb{E}_{t-1}[\|Z_t\|^2] \leq \mathbb{E}_{t-1}[\|X_t\|^{2p}] \leq \mathbb{E}_{t-1}[\|X_t\|^p] R^p \leq \sigma_t^p R^p$$

Thus, again by Lemma 11, with probability at least $1 - \delta$:

$$\sum_{t=1}^k Z_t \leq \frac{4R^p \log(1/\delta)}{3} + \sqrt{2 \sum_{t=1}^T R^p \sigma_t^p \log(1/\delta)}$$

$$\sum_{t=1}^k \|X_t\|^p \leq \sum_{t=1}^k \sigma_t^p + \frac{4R^p \log(1/\delta)}{3} + \sqrt{2 \sum_{t=1}^k R^p \sigma_t^p \log(1/\delta)}$$

Applying Young inequality:

$$\leq \sum_{t=1}^k \sigma_t^p + \frac{4R^p \log(1/\delta)}{3} + \sqrt{R^{2p} \log^2(1/\delta) + \left(\sum_{t=1}^k \sigma_t^p\right)^2}$$

$$\leq 2 \sum_{t=1}^k \sigma_t^p + \frac{7R^p \log(1/\delta)}{3}$$

Putting everything together, with probability at least $1 - 3\delta$:

$$\left\| \sum_{t=1}^k X_t \right\| \leq \left| \sum_{t=1}^k s_t \right| + \left(R^p + C \sum_{t=1}^k \|X_t\|^p \right)^{1/p}$$

$$\leq \frac{2R(\log(1/\delta) + 3/2)}{3} + \sqrt{2 \sum_{t=1}^k \sigma_t^p \log(1/\delta) + \left(2C \sum_{t=1}^k \sigma_t^p + \frac{7CR^p \log(1/\delta)}{3} \right)^{1/p}}$$

Now, if $C \geq 1$, we can apply Lemma 13 and over approximate to obtain with probability $1 - 3\delta$:

$$\left\| \sum_{t=1}^k X_t \right\| \leq \frac{2R(\log(1/\delta) + 3/2)}{3} + \left(\sum_{t=1}^k \sigma_t^p \right)^{1/p} \sqrt{2 \log(1/\delta)} + \left(2C \sum_{t=1}^k \sigma_t^p \right)^{1/p} + \left(\frac{7CR^p \log(1/\delta)}{3} \right)^{1/p}$$

$$\leq 5CR \max(1, \log(1/\delta)) + 4 \left(C \sum_{t=1}^k \sigma_t^p \right)^{1/p} \sqrt{\max(1, \log(1/\delta))}$$

□




Figure 2: Comparison between warm-up and burn-in strategies on Masked language modeling BERT-Base pre-training task. Plot (a) shows the eval accuracy when trained with Algorithm 1 and plot (b) does the same when trained with LAMB.

D Experiments

In this section, we run additional experiments on BERT-Base pre-training task. BERT-Base is a smaller version of BERT-Large model we experimented in Section 4. Again, all our experiments were conducted using the Tensorflow framework on TPUv3 architecture. When using Algorithm 1, we train all the models shown using the base learning rate η_0 of 0.5 and batch size of 512. We came up with this learning rate by performing grid search $\eta_0 \in [1.0, 0.5, 0.1, 0.05, 0.01, 0.005, 0.001, 0.0005, 0.0001]$ and choosing the one which attains the best eval accuracy. Again, shown in Fig 2, our warmup baseline employs the standard practice of linear warm-up for 3125 steps and polynomial decay of $\eta_t = \eta_0 * (1 - \frac{t}{T})$ for the rest of the steps. For a fair comparison with the baseline, we start the polynomial decay after 3125 steps regardless of how many steps were used for burn-in. We set the learning rate to the constant base learning rate value of 0.5 for for steps in between burn-in and polynomial decay. For the models trained using burn-in, instead of linear warm-up, we set the base learning rate η_0 to zero for the amount of steps we run the burn-in for. As shown in Fig 2, we perform a grid search across $[0, 250, 500, 1000, 2000, 3125]$ steps for the initial burn-in period and compare their results. When training the smaller BERT-Base model we observe that burn-in is less crucial compared to when training BERT-Large i.e. the model converges regardless of the amount of burn-in employed. Moreover, we don't see any convergence issues when we don't employ even any warm-up, this suggests that warm-up is also much less crucial (similar to burn-in) on the simpler BERT-Base problem.

E Smoothness Lemmas in Banach Spaces

In this section we provide a few technical lemmas relating to smoothness in Banach spaces. Most of these are fairly straightforward generalizations of standard results in Hilbert spaces.

We begin with a notations and definitions. Recall that a *Banach space* over the reals is a real vector space \mathcal{B} together with a norm $\|\cdot\| : \mathcal{B} \rightarrow \mathbb{R}$ such that \mathcal{B} is complete with respect to the topology induced by the norm. The dual space \mathcal{B}^* is the space of bounded linear functions $\mathcal{B} \rightarrow \mathbb{R}$. When \mathcal{B} is a Hilbert space, there is a natural isomorphism between $i : \mathcal{B}^* \rightarrow \mathcal{B}$ such that for any $v \in \mathcal{B}^*$ and $w \in \mathcal{B}$, $v(w) = \langle i(v), w \rangle$. In order to preserve notational familiarity, even when \mathcal{B} is *not* a Hilbert space, we therefore adopt the notation $\langle v, w \rangle$ to refer to $v(w)$, the application of the linear function $v \in \mathcal{B}^*$ to the vector $w \in \mathcal{B}$. Note that in this way the ordering of the arguments to $\langle \cdot, \cdot \rangle$ is important: $\langle w, v \rangle$ is meaningless. The dual space \mathcal{B}^* is equipped with the *dual norm* $\|v\|_* = \sup_{\|x\| \leq 1} \langle v, x \rangle$, and is itself a Banach space.

A function $F : \mathcal{B} \rightarrow \mathbb{R}$ is Frechet differentiable if for all $w \in \mathcal{B}$ there is a bounded linear operator $v \in \mathcal{B}^*$ such that:

$$\lim_{\delta \rightarrow 0} \frac{|F(w + \delta) - F(w) - \langle v, \delta \rangle|}{\|\delta\|} = 0$$

the operator v is denoted $\nabla F(w)$. Note that if F is Frechet differentiable, it also satisfies the familiar directional derivative identity: if $f(t) = F(w + t\delta)$, then

$$f'(t) = \langle \nabla F(w + t\delta), \delta \rangle$$

The second derivative or Hessian of F at a point w is a linear map $\nabla^2 F(w) : \mathcal{B} \rightarrow \mathcal{B}^*$ satisfying:

$$\lim_{\delta \rightarrow 0} \frac{\|\nabla F(w + \delta) - \nabla F(w) - \nabla^2 F(w)\delta\|_*}{\|\delta\|} = 0$$

In this paper we exclusively consider spaces such that \mathcal{B}^* satisfies:

$$\|x + y\|_*^p \leq \|x\|_*^p + \langle \nabla \|x\|_*^p, y \rangle + C\|y\|_*^p \quad (8)$$

for some p and C for all $x, y \in \mathcal{B}^*$. We will also assume that $C \geq 1$.

Proposition 15. *Suppose that $F : \mathcal{B} \rightarrow \mathbb{R}$ is Frechet differentiable and the gradient satisfies $\|\nabla F(x) - \nabla F(y)\|_* \leq L\|x - y\|$ for all $x, y \in \mathcal{B}$ (i.e. F is L -smooth). Then for all $x, y \in \mathcal{B}$:*

$$F(x) \leq F(y) + \langle \nabla F(y), x - y \rangle + \frac{L}{2}\|x - y\|^2$$

Proof. Consider the function $g : \mathbb{R} \rightarrow \mathbb{R}$ given by $g(t) = F(y + t(x - y))$. Since F is Frechet differentiable, g is differentiable and $g'(t) = \langle \nabla F(y + t(x - y)), (x - y) \rangle$. Therefore by the fundamental theorem of calculus:

$$\begin{aligned}
g(1) &= g(0) + \int_0^1 g'(t) dt \\
&= g(0) + g'(0) + \int_0^1 (g'(t) - g'(0)) dt \\
&\leq g(0) + \langle \nabla F(y), x - y \rangle + \int_0^1 |g'(t) - g'(0)| dt \\
&\leq g(0) + \langle \nabla F(y), x - y \rangle + \int_0^1 |\langle \nabla F(y + t(x - y)) - \nabla F(y), x - y \rangle| dt \\
&\leq g(0) + \langle \nabla F(y), x - y \rangle + \int_0^1 tL \|x - y\|^2 dt \\
&= g(0) + \langle \nabla F(y), x - y \rangle + \frac{L}{2} \|x - y\|^2
\end{aligned}$$

Now plug in the definition of $g(1)$ and $g(0)$ to conclude the argument. \square

Proposition 16. *Suppose that $F : \mathcal{B} \rightarrow \mathbb{R}$ is twice Frechet differentiable and the Hessian satisfies $\|(\nabla^2 F(x) - \nabla^2 F(y))v\|_* \leq \rho \|v\| \|x - y\|$ for all $x, y, v \in \mathcal{B}$ (i.e. F is ρ -second-order smooth). Then for all $x, y \in \mathcal{B}$:*

$$F(x) \leq F(y) + \langle \nabla F(y), x - y \rangle + \frac{\langle \nabla^2 F(y)(x - y), x - y \rangle}{2} + \frac{\rho}{6} \|x - y\|^3$$

Proof. Consider the function $g : \mathbb{R} \rightarrow \mathbb{R}$ given by $g(t) = F(y + t(x - y))$. Since F is twice Frechet differentiable, g is twice differentiable and $g'(t) = \langle \nabla F(y + t(x - y)), x - y \rangle$ and $g''(t) = \langle \nabla^2 F(y + t(x - y))(x - y), x - y \rangle$. Therefore by the fundamental theorem of calculus:

$$\begin{aligned}
g(1) &= g(0) + \int_0^1 g'(t) dt \\
&= g(0) + \int_0^1 \left(g'(0) + \int_0^t g''(k) dk \right) dt \\
&= g(0) + g'(0) + \int_0^1 \int_0^t g''(k) dk dt \\
&= g(0) + g'(0) + \int_0^1 \int_0^t g''(0) dk dt + \int_0^1 \int_0^t (g''(k) - g''(0)) dk dt \\
&= g(0) + g'(0) + \frac{g''(0)}{2} + \int_0^1 \int_0^t (g''(k) - g''(0)) dk dt \\
&= g(0) + g'(0) + \frac{g''(0)}{2} + \int_0^1 \int_0^t \langle (\nabla^2 F(y + k(x - y)) - \nabla^2 F(y))(x - y), x - y \rangle dk dt \\
&\leq g(0) + g'(0) + \frac{g''(0)}{2} + \int_0^1 \int_0^t \rho \|x - y\|^3 k dk dt \\
&\leq g(0) + g'(0) + \frac{g''(0)}{2} + \frac{\rho \|x - y\|^3}{6}
\end{aligned}$$

Now plug in the definition of $g(1)$ and $g(0)$ to conclude the argument. \square

Proposition 17. *Suppose that $F : \mathcal{B} \rightarrow \mathbb{R}$ is twice Frechet differentiable and the Hessian satisfies $\|(\nabla^2 F(x) - \nabla^2 F(y))v\|_* \leq \rho \|v\| \|x - y\|$ for all $x, y, v \in \mathcal{B}$ (i.e. F is ρ -second-order smooth). Then for all $x, y \in \mathcal{B}$:*

$$\|\nabla F(y) - \nabla F(x) - \nabla^2 F(x)(y - x)\|_* \leq \frac{\rho}{2} \|x - y\|^2$$

Proof. For any fixed $z \in \mathcal{B}$, define $g_z(t) = \langle \nabla F(x + t(y - x)), z \rangle$. Then we have $g'_z(t) = \langle \nabla^2 F(x + t(y - x))(y - x), z \rangle$. Then, by the fundamental theorem of calculus:

$$\begin{aligned} g_z(1) &= g_z(0) + \int_0^1 g'_z(t) dt \\ &\leq g_z(0) + g'_z(0) + \int_0^1 |g'_z(t) - g'_z(0)| dt \\ &\leq g_z(0) + g'_z(0) + \int_0^1 \rho t \|x - y\|^2 \|z\| dt \\ &\leq g_z(0) + g'_z(0) + \frac{\rho \|x - y\|^2 \|z\|}{2} \end{aligned}$$

rearrange and apply the definition of g_z :

$$\langle \nabla F(y) - \nabla F(x) - \nabla^2 F(x)(y - x), z \rangle \leq \frac{\rho \|x - y\|^2 \|z\|}{2}$$

Since z was arbitrary, by definition of dual norm we have

$$\|\nabla F(y) - \nabla F(x) - \nabla^2 F(x)(y - x)\|_* \leq \frac{\rho}{2} \|x - y\|^2$$

□

The following important Lemma generalizes Lemma 2 of [9] to more general norms, and improves constants.

Lemma 1. *Suppose $F : \mathcal{B} \rightarrow \mathbb{R}$ be a Frechet-differentiable L -Smooth function from a Banach space \mathcal{B} to the reals. Let $w \in \mathcal{B}$. Let $g^* \in \mathcal{B}^*$ and let $g \in \mathcal{B}$ be a unit-vector satisfying $\langle g^*, g \rangle = \|g\|_*$. Define $w' = w - \eta g$. Define $\epsilon = g^* - \nabla F(w)$. Then:*

$$F(w') \leq F(w) - \eta \|\nabla F(w)\|_* + 2\eta \|\epsilon\|_* + \frac{L\eta^2}{2}$$

Proof. First, by smoothness (Proposition 15) we have

$$F(w') \leq F(w) + \langle \nabla F(w), w' - w \rangle + \frac{L}{2} \|w' - w\|^2$$

Next, since $\|g\| = 1$,

$$\begin{aligned} F(w') &\leq F(w) - \eta \langle \nabla F(w), g \rangle + \frac{L\eta^2}{2} \\ &= F(w) - \eta \langle g^* - \epsilon, g \rangle + \frac{L\eta^2}{2} \\ &= F(w) - \eta (\|g^*\|_* - \langle \epsilon, g \rangle) + \frac{L\eta^2}{2} \\ &\leq F(w) - \eta \|\nabla F(w) + \epsilon\|_* + \eta \|\epsilon\|_* \|g\| + \frac{L\eta^2}{2} \\ &\leq F(w) - \eta \|\nabla F(w)\|_* + 2\eta \|\epsilon\|_* + \frac{L\eta^2}{2} \end{aligned}$$

□