GRADIENT DESCENT ON NEURAL NETWORKS TYPICALLY OCCURS AT THE EDGE OF STABILITY

Anonymous authors
Paper under double-blind review

ABSTRACT

We empirically demonstrate that full-batch gradient descent on neural network training objectives typically operates in a regime we call the Edge of Stability. In this regime, the leading eigenvalue of the training loss Hessian hovers just above the value $2/(\text{step size})$, and the training loss behaves non-monotonically over short timescales, even as it consistently decreases over long timescales. Our results expose a large gap between theory and practice in non-convex optimization: many analyses of gradient descent (those based on L-smoothness or monotonic descent) do not apply at the Edge of Stability. Our results also shed significant light on the dynamics of gradient descent with a fixed step size, and how that algorithm differs from gradient flow.

1 INTRODUCTION

This paper empirically studies full-batch gradient descent with a fixed step size η, applied to neural network training. We demonstrate that, at all but the smallest learning rates, gradient descent eventually transitions into a regime we call the *Edge of Stability*, in which the sharpness (the largest eigenvalue of the training loss Hessian) hovers right at (or just past) the value $2/\eta$. This behavior is consistent across architectures, activation functions, tasks, loss functions, and batch normalization. An analogous phenomenon occurs with both Polyak and Nesterov momentum.

The Edge of Stability reflects a balance between two opposing processes: *progressive sharpening* and *instability*. On the one hand, we demonstrate empirically that gradient descent on neural network training objectives has an overwhelming tendency to continually move in a direction of increased sharpness. We call this phenomenon *progressive sharpening*. On the other hand, gradient descent

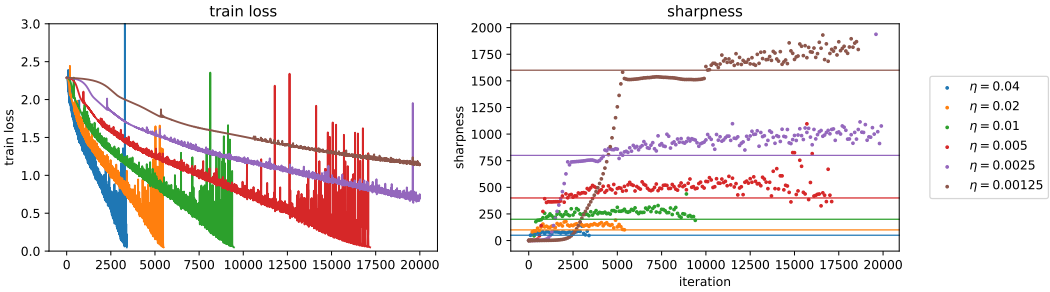

Figure 1: We train a VGG-11 (no batchnorm) on CIFAR-10 using full-batch gradient descent with a fixed learning rate η. The left pane plots the training loss; the right pane plots the sharpness, with a solid line marking the value $2/\eta$. Observe that the sharpness rises unil reaching the value $2/\eta$, at which point gradient descent transitions into the Edge of Stability, a regime where the sharpness hovers right at or above $2/\eta$, and the training loss behaves non-monotonically. We emphasize that the eventual transition to the Edge of Stability does not just occur at the largest trainable learning rates (e.g. blue), but also at much smaller learning rates that a practitioner would consider intolerably small (e.g. brown). Figure 6 shows that a batch-normalized VGG-11 exhibits the same behavior.

with step size η is, to quadratic order, unstable in regions of the loss landscape where the sharpness exceeds the value $2/\eta$: in such a region, gradient descent would diverge if run on the quadratic Taylor approximation. The joint consequence of sharpening and instability is that during gradient descent, the sharpness increases to, and then lingers at, the value $2/\eta$ — the Edge of Stability.

Since neural network training is poorly understood, we find it striking that across many different architectures and tasks, for both vanilla and momentum gradient descent, there is a simple rule that can *numerically* predict, to respectable accuracy, the value of some statistic (the sharpness) during training. Furthermore, our results (1) expose a significant gap between optimization theory and practice, and (2) shed light on the dynamics of gradient descent with a non-infinitesimal step size.

Optimization theory The conventional wisdom in optimization theory is that the step size η of gradient descent should be small enough that the sharpness along the optimization trajectory is less than $2/\eta$ (Bottou et al., 2018; Li & Arora, 2019). We refer to this assumption as the Bounded Curvature (BC) condition. The BC condition ensures both (1) that gradient descent will monotonically decrease the training objective, and (2) that gradient descent is a stable algorithm. In the case of neural network training, our paper complicates this conventional wisdom in three respects. First, because of progressive sharpening, in order for the BC condition to hold for the entire duration of training, the step size η must be set surprisingly small. Second, while neural networks can in principle be trained with these extremely small learning rates, we find that they can also be successfully trained at much larger learning rates. At these large learning rates, gradient descent eventually transitions into the Edge of Stability, a regime where gradient descent succeeds even though the BC condition does not hold. Third, gradient descent converges much faster at these larger learning rates; learning rates small enough for the BC condition to hold throughout training are highly suboptimal in terms of convergence speed — in a word, they are "unreasonable." Thus, during gradient descent training of neural networks at reasonable learning rates, the BC condition does not hold. An extreme (but not atypical) case is the network in Figure 1: while we know that there must exist learning rates small enough for the BC condition to hold throughout training, it is not feasible on an actual computer to train this network to completion at such learning rates.

Even analyses of gradient descent that do not explicitly assume the BC condition often assert that gradient descent will decrease the training objective monotonically (Allen-Zhu et al., 2019; Zhang et al., 2020). However, at the Edge of Stability, the training loss usually behaves non-monotonically over short timescales even as it consistently decreases over long timescales. Thus, in deep learning, gradient descent at reasonable learning rates does not decrease the training objective monotonically.

Dynamics of gradient descent Our paper sheds significant light on the dynamics of gradient descent on neural network training objectives, and specifically on how the trajectory of gradient descent differs from that of gradient flow (gradient descent with infinitesimal step sizes; see Bach (2020) for a primer). We show that gradient flow tends to continually head into sharper and sharper regions of the loss landscape (progressive sharpening). However, gradient descent with step size η can only track the gradient flow trajectory up till the point on that trajectory where the sharpness reaches $2/\eta$. After this point, gradient descent departs from the gradient flow trajectory and enters the Edge of Stability. At the Edge of Stability, gradient descent is constantly "trying" to follow gradient flow into a region of increased sharpness (and indeed, this is what would happen if we were to drop the learning rate), but is being (somehow) blocked from doing so by considerations of stability. Thus, because of progressive sharpening and instability, the optimization trajectory of gradient descent differs from that of gradient flow, and the trajectory of GD with a large learning rate differs from that of GD with a small learning rate. Note that while the trajectories of gradient descent at different stable step sizes can conceivably differ from the gradient flow trajectory, and from one another, due to discretization error (Nakkiran, 2020; Barrett & Dherin, 2020), our experiments demonstrate that in practice, the predominant impact of non-infinitesimal step sizes is rather due to instability.

Stochastic gradient descent We study full-batch gradient descent rather than SGD because only in the full-batch case are we (so far) able to precisely characterize the dynamics of the sharpness, and the effect of these dynamics on the optimization trajectory. Jastrzebski et al. (2020) recently studied similar questions for SGD, and drew much weaker conclusions (see Section 2). However, some of our results may have analogues for SGD (Section 5).

2 BACKGROUND AND RELATED WORK

Evolution of sharpness A number of recent works have empirically studied the evolution of the sharpness (and, more generally, the Hessian spectrum) during neural network training (Sagun et al., 2017; Ghorbani et al., 2019; Li et al., 2020; Papyan, 2018; 2019; 2020) Notably, several works have observed that the hyperparameters of SGD seem to affect the sharpness along the optimization trajectory: large learning rates and small batch sizes steer SGD into less sharp regions of parameter space (Keskar et al., 2016; Jastrzebski et al., 2017; 2019); in the case of full-batch gradient descent, our paper quantifies this.

Most similar to our paper is Jastrzebski et al. (2020), who postulated a simplified model for the beginning of SGD training which assumed, among other things, that the sharpness would increase until its increase caused SGD to become destabilized. However, while the stability properties of GD are clear (to quadratic order, the algorithm diverges iff the sharpness exceeds $2/\eta$) the stability properties of SGD are not yet well-understood; indeed, the stability of SGD only *relates to* the sharpness under a strong assumption not likely to be fully satisfied in practice. Consequently, Jastrzebski et al. (2020) could not (and did not) actually *verify* their assumption that the sharpness would increase as much as it possibly could. Moreover, whereas Jastrzebski et al. (2020) only conjectured the sharpness "wants to" increase at the very beginning of training, we find that this tendency exists throughout training. Finally, the strongest prediction in Jastrzebski et al. (2020) is that the maximum sharpness along the SGD optimization trajectory would be lower when the learning rate is larger or the batch size is smaller. This paper demonstrates the much stronger result that during full-batch gradient descent with step size η, the sharpness settles right at (or just past) the numerical value $2/\eta$.

Contrary to our results, in Ghorbani et al. (2019) a learning rate drop caused the sharpness to decrease rather than increase. We believe that this is because they used SGD rather than GD.

Bounded Curvature Our paper shows that gradient descent is perfectly successful even at the Edge of Stability, where the Bounded Curvature condition does not hold. Previously, Xing et al. (2018) observed that full-batch gradient descent eventually enters a regime where the training loss behaves non-monotonically, but did not describe when/why this occurs, and did not discuss the implications of this observation for optimization theory. Lewkowycz et al. (2020) demonstrated that an initial sharpness greater than $2/\eta$ is not necessarily fatal for gradient descent, since gradient descent can "catapult" out of that region into a much flatter region, rather than diverging entirely. However, the implication in Lewkowycz et al. (2020) is that the actual progress would occur in regions where the sharpness remains strictly less than $2/\eta$. Our experiments demonstrate otherwise.

Dynamical stability At the Edge of Stability, the stability properties of gradient descent are (somehow) constantly preventing it from following gradient flow into sharper regions of the loss landscape. Several previous works have argued that the stability properties of optimization algorithms could potentially serve as a form of implicit bias in deep learning (Nar & Sastry, 2018; Wu et al., 2018; Giladi et al., 2020; Jastrzebski et al., 2020; Mulayoff & Michaeli, 2020). Our contribution here is to demonstrate that due to progressive sharpening, gradient descent's implicit stability constraint "sharpness $\leq 2/\eta$" is typically "tight" throughout large portions of training.

L-smoothness The condition we call Bounded Curvature is an instance of a more general concept in optimization called L-smoothness. Formally, an objective function $f : \mathbb{R}^p \to \mathbb{R}$ is said to be (globally) L-smooth if the gradient map is globally L-Lipschitz, i.e. if $\|\nabla f(x) - \nabla f(y)\|_2 \leq L\|x - y\|_2$ for all $x, y \in \mathbb{R}^p$. For a twice-differentiable function, this is equivalent to the condition that the Hessian's eigenvalues are no greater than L in magnitude, i.e. $\|\nabla^2 f(x)\|_2 \leq L$ for all $x \in \mathbb{R}^p$ (Bubeck, 2014) The main consequence of L-smoothness is the *descent lemma*: for an L-smooth function f, a gradient descent step with learning rate η is guaranteed to decrease the training objective by at least $\eta\left(1 - \frac{1}{2}\eta L\right)\|\nabla f(x)\|^2$. In particular, if $\left(1 - \frac{1}{2}\eta L\right) > 0 \iff \eta < 2/L$, then a gradient step of size η is guaranteed to decrease (rather than increase) the training objective. Using the descent lemma, one can prove convergence results for gradient descent that only require the objective function to be L-smooth, and not convex. A famous example is the result stated in e.g. Nesterov (1998) Eq. 1.2.12 that gradient descent with step size $\eta = 1/L$ on an L-smooth function will eventually converge to a stationary point. The most minimal version of this kind of analysis would only assume the Bounded Curvature condition (i.e. that L-smoothness holds only near the

optimization trajectory, and the function is only L-smooth enough to ensure descent). Our paper, however, demonstrates that at reasonable learning rates, gradient descent eventually transitions into a regime where *even* the Bounded Curvature condition does not hold.

Since our experiments focus on gradient descent, they do not necessarily imply that L-smoothness assumptions are unjustified when analyzing other optimization algorithms. However, gradient descent is arguably the simplest optimization algorithm, so we believe that the fact that Bounded Curvature fails even there should raise serious questions about the suitability of the L-smoothness assumption in neural network optimization more generally. The "non-convex but L-smooth" setting is currently ubiquitous in optimization papers that are seemingly motivated by deep learning (Reddi et al., 2016; Agarwal et al., 2017; Zaheer et al., 2018; Zhou et al., 2018; Chen et al., 2018; Li & Orabona, 2019; Ward et al., 2019; You et al., 2020; Vaswani et al., 2019; Sankararaman et al., 2019; Reddi et al., 2020; Défossez et al., 2020; Xie et al., 2020; Liu et al., 2020). We think that authors may want to reconsider this practice in light of our results.

3 SHARPENING AND INSTABILITY

In this section, we characterize the trajectory traversed by gradient descent with a fixed step size η, with the aim of describing when and why gradient descent enters the Edge of Stability.

If the sharpness at initialization is greater than $2/\eta$, then either training will diverge entirely, or gradient descent will "locally diverge" out of the immediate neighborhood of initialization, but recover in a flatter area (e.g. see Figure 14 or 15). Lewkowycz et al. (2020) dubbed this a "catapult."

On the other hand, if the initial sharpness is less than $2/\eta$, then at first, gradient descent will roughly[1] track the gradient flow trajectory. A crucial empirical property of gradient flow on neural network training objectives is that **gradient flow has an overwhelming tendency to continually head in a direction of increased sharpness.** We call this phenomenon *progressive sharpening*. By "overwhelming tendency," we mean that gradient flow does occasionally decrease the sharpness (see e.g. Figure 24), but these brief decreases always seem be followed by a return to continual increase. We do not know why progressive sharpening occurs; this is an important question for future work. In Appendix D we take a step towards reconciling progressive sharpening with NTK theory.

Progressive sharpening presents a complication for gradient descent. Whereas gradient flow can traverse arbitrarily sharp regions of the loss landscape without difficulty, gradient descent is unstable (to quadratic order) in regions where the sharpness is greater than $2/\eta$. By this, we mean that gradient descent would diverge if run on the quadratic Taylor approximation. In particular, suppose that starting from any iterate, we suddenly switched from running gradient descent on the real objective to running gradient descent on the Taylor approximation. Let Δ_t be the displacement along the leading Hessian eigenvector and let λ be the sharpness. Then Δ_t would evolve under gradient descent as $\Delta_t = c\left[(1-\eta\lambda)^t - 1\right]$, where c is a constant (Appendix B). This sequence diverges iff $|1-\eta\lambda| > 1$, which, since $\lambda, \eta > 0$, means $\lambda > 2/\eta$. Thus, if the sharpness exceeds $2/\eta$, then gradient descent run on the quadratic Taylor approximation would oscillate with increasing magnitude along the direction of highest curvature, leading to eventual divergence.

The learning rate η determines whether or not gradient descent remains stable for the entire duration of training. On the one hand, if η is so small that the sharpness along the entire gradient flow trajectory, from start to end, remains less than $2/\eta$, then gradient descent will stably track the gradient flow trajectory all the way to the end. (Put another way, this occurs if $\eta < 2/\lambda_{\max}$, where $\lambda_{\max}$ is the maximum sharpness along the gradient flow trajectory.) On the other hand, if $\eta > 2/\lambda_{\max}$, then gradient descent will eventually reach a point where the local sharpness crosses $2/\eta$, causing gradient descent to become destabilized. Once this occurs, we observe that gradient descent enters a new regime we call the Edge of Stability. In this regime, gradient descent continues to successfully optimize the loss, but in such a way that the sharpness hovers right at (or just beyond) $2/\eta$ instead of increasing further. In other words, at the Edge of Stability, gradient descent is essentially executing a local search subject to an active constraint that the sharpness remain $\leq 2/\eta$.

[1] As the step size tends to zero, the gradient descent trajectory will converge to the gradient flow trajectory, since GD is a forward Euler discretization of gradient flow, but if the step size is close to $2/\text{sharpness}$, these trajectories can conceivably differ. Nevertheless, we empirically find this to be a good approximation.

We now illustrate these points with a single example; the following section will show that they hold more generally. We consider training a network on a subset of CIFAR-10 using full-batch gradient descent. The network is a fully-connected architecture with two hidden layers of 200 units each (656,810 total parameters), and with ELU activations (Clevert et al., 2016). We train on a subset of CIFAR-10 of size 5,000. We use the square loss for classification (Hui & Belkin, 2020), encoding the target class with 1 and the other classes with 0; our findings also hold for cross-entropy loss, but with one minor twist. We consider training finished when the loss reaches 0.05.

We train this network at three small learning rates $\eta \in \{0.001, 0.0015, 0.002\}$, plotting the training loss in Figure 2(a). These three learning rates are small enough that gradient descent roughly follows the gradient flow trajectory from start to finish, traveling at a speed given by the learning rate. To illustrate that the gradient descent at these three learning rates is following the same path, in Figure 2(b) we plot the training loss by "time" = iteration $\times$ learning rate. This allows us to directly compare the loss after, say, 1000 iterations at training rate 1e-3 with the loss after 500 iterations of learning rate 2e-3 (both are time 1). When plotted by time, the three loss curves coincide exactly.

Figure 2(c) plots the evolution of the sharpness. Observe that sharpness continually rises, except for a brief downward blip at the beginning of training. This is progressive sharpening. Notice that the maximum sharpness along the gradient flow trajectory is just a bit under 1000, which means that learning rates of $\eta = 2/1000 = 0.002$ or smaller will be stable for the entire duration of training. Crucially, all three of our learning rates were at or below this threshold; this is the reason why gradient descent followed the gradient flow trajectory from start to finish.

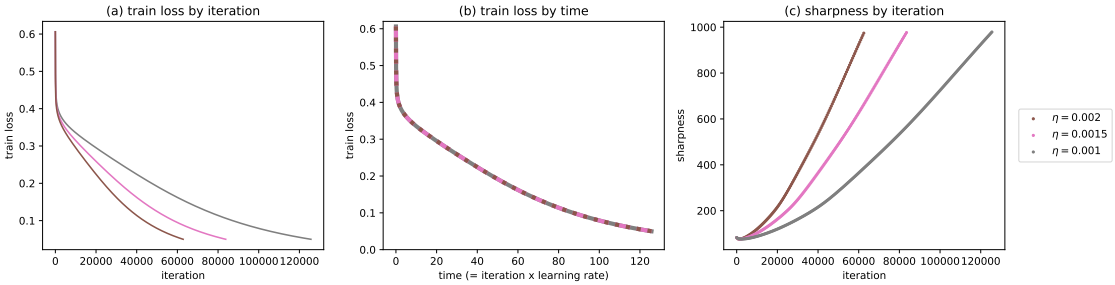

Figure 2: **At tiny learning rates, gradient descent tracks the gradient flow trajectory, the sharpness remains below $2/\eta$, and the train loss decreases monotonically.** (a) The train loss monotonically decreases. (b) When plotted by time, the three train loss curves exactly coincide. (c) The sharpness continually increases.

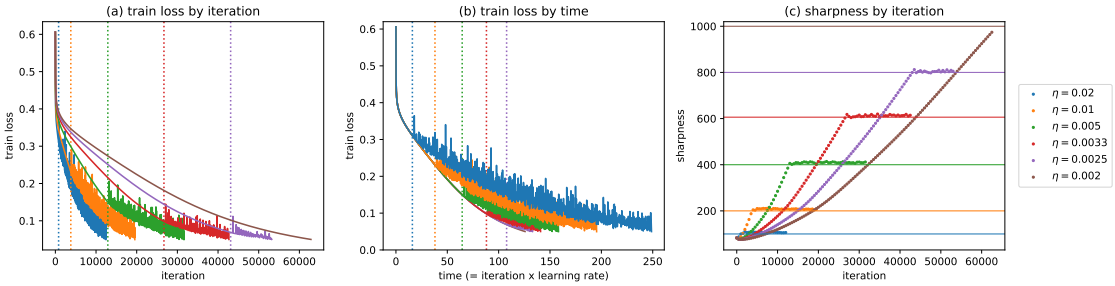

Figure 3: **At larger learning rates, the sharpness eventually hits $2/\eta$, at which point gradient descent departs from the gradient flow trajectory and enters the Edge of Stability.** (a) The vertical line marks the first iteration when the sharpness crosses $2/\eta$; shortly thereafter, the train loss starts to behave non-monotonically. (b) Gradient descent follows the gradient flow trajectory (i.e. the train loss curves coincide) only up to the point where sharpness reaches $2/\eta$. (c) We mark the value $2/\eta$ with a horizontal line. The sharpness does not increase beyond this value.

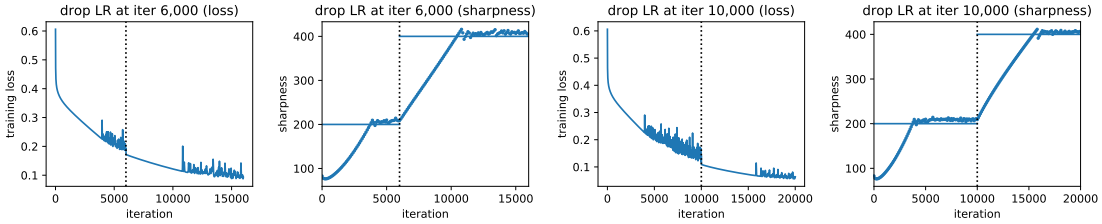

Figure 4: **After a learning rate drop, progressive sharpening resumes until gradient descent reaches the Edge of Stability once again.** For the $\eta = 0.01$ network, we halve the learning rate at iterations 6,000 (left two panes) and 10,000 (right two panes). See Figure 30 for more experiments.

We now train the same network at larger learning rates, plotting the training loss in Figure 3(a). For each learning rate, we draw a vertical dotted line at the iteration where the sharpness first crosses $2/\eta$. Observe that shortly after the sharpness crosses $2/\eta$, the training loss starts to behave non-monotonically. In Figure 3(b), we plot the training loss curves as a function of time. Observe that gradient descent closely tracks the gradient flow trajectory up until the point where the sharpness crosses $2/\eta$. After this point, gradient descent peels off from the gradient flow trajectory and follows a different route. (We include additional evidence for this claim in Appendix C.) In Figure 3(c), we plot the evolution of the sharpness over the course of training. For each learning rate η, we draw a horizontal line at the value $2/\eta$. Observe that once the sharpness hits $2/\eta$, it ceases to increase further, and instead hovers right at (or just a bit beyond) $2/\eta$ — at the Edge of Stability.

At the Edge of Stability, gradient descent is constantly "trying" to increase the sharpness, but is being somehow restrained from doing so. If we cut the learning rate, this restraint is removed, and gradient descent is freed to follow gradient flow into a region of increased sharpness. To demonstrate this, in Figure 4, we run gradient descent until it enters the Edge of Stability, then halve the learning rate and continue training. Observe that as soon as the learning rate is cut, the sharpness starts to increase. It only stops increasing once gradient descent is back at the Edge of Stability.

Momentum A natural question is whether the evolution of the sharpness can be precisely characterized for optimization algorithms other than gradient descent. Our experience suggests that the answer is no. We believe this is because even on quadratic functions, the stability of optimization algorithms other than gradient descent is not necessarily determined by a sharpness threshold.

There is one exception: gradient descent with momentum. Similar to vanilla gradient descent, on quadratic objectives, momentum gradient descent (with either Polyak or Nesterov momentum) acts independently along each Hessian eigenvector, and will diverge along the leading eigenvector if and only if the sharpness exceeds a certain threshold (the "maximum stable sharpness", or MSS). We

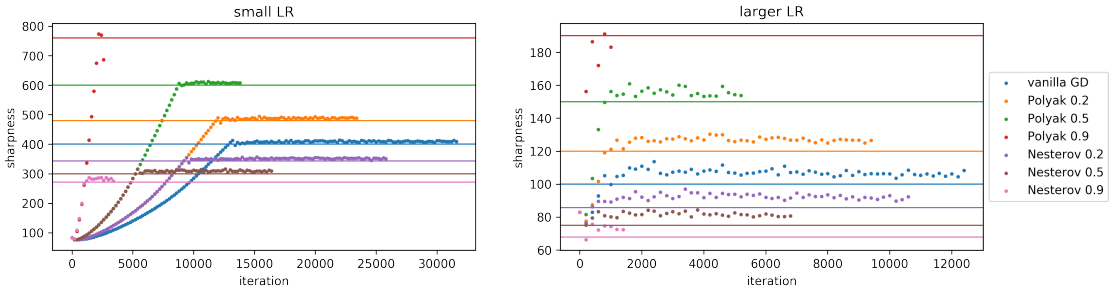

Figure 5: **Momentum gradient descent operates at the Edge of Stability.** We train the network to completion with fixed learning rate (left: $\eta = 0.005$; right: $\eta = 0.02$) but varying amounts β of either Polyak or Nesterov momentum (see legend). For each algorithm, the horizontal solid line of the appropriate color marks the maximum stable sharpness given by equation 1.

prove in Appendix B that for GD with learning rate η and momentum parameter β, the MSS is:

$$\mathrm{MSS}_{\mathrm{Polyak}}(\eta, \beta) = \frac{1}{\eta}(2 + 2\beta), \quad \mathrm{MSS}_{\mathrm{Nesterov}}(\eta, \beta) = \frac{1}{\eta}\left(\frac{2 + 2\beta}{1 + 2\beta}\right). \tag{1}$$

The Polyak result previously appeared in Goh (2017); the Nesterov one seems to be new. It is interesting to note that $\mathrm{MSS}_{\mathrm{Polyak}}(\eta, \beta)$ is increasing in β, whereas $\mathrm{MSS}_{\mathrm{Nesterov}}(\eta, \beta)$ is decreasing in β. That is, adding Polyak momentum gives gradient descent a higher tolerance for sharpness, whereas adding Nesterov momentum gives gradient descent a lower tolerance for sharpness.

In Figure 5, we train our network to completion with varying amounts β of both Polyak and Nesterov momentum, while keeping the learning rate η held fixed. Just like in the vanilla GD case, we see that the sharpness rises to the maximum stable sharpness and then remains approximately there.

Cross-entropy loss So far, we have focused on the squared loss. The situation is mostly the same for the commonly-used cross-entropy loss, but with one difference: for cross-entropy loss, at the end of training, we often observe that the sharpness decreases. This behavior seems to be caused by the well-known fact that at the end of training, gradient descent continually scales up the logits so as to drive the cross-entropy loss to zero (Soudry et al., 2018). When this occurs, the sharpness drops, because the second derivative of the cross-entropy goes to zero as its input increases.

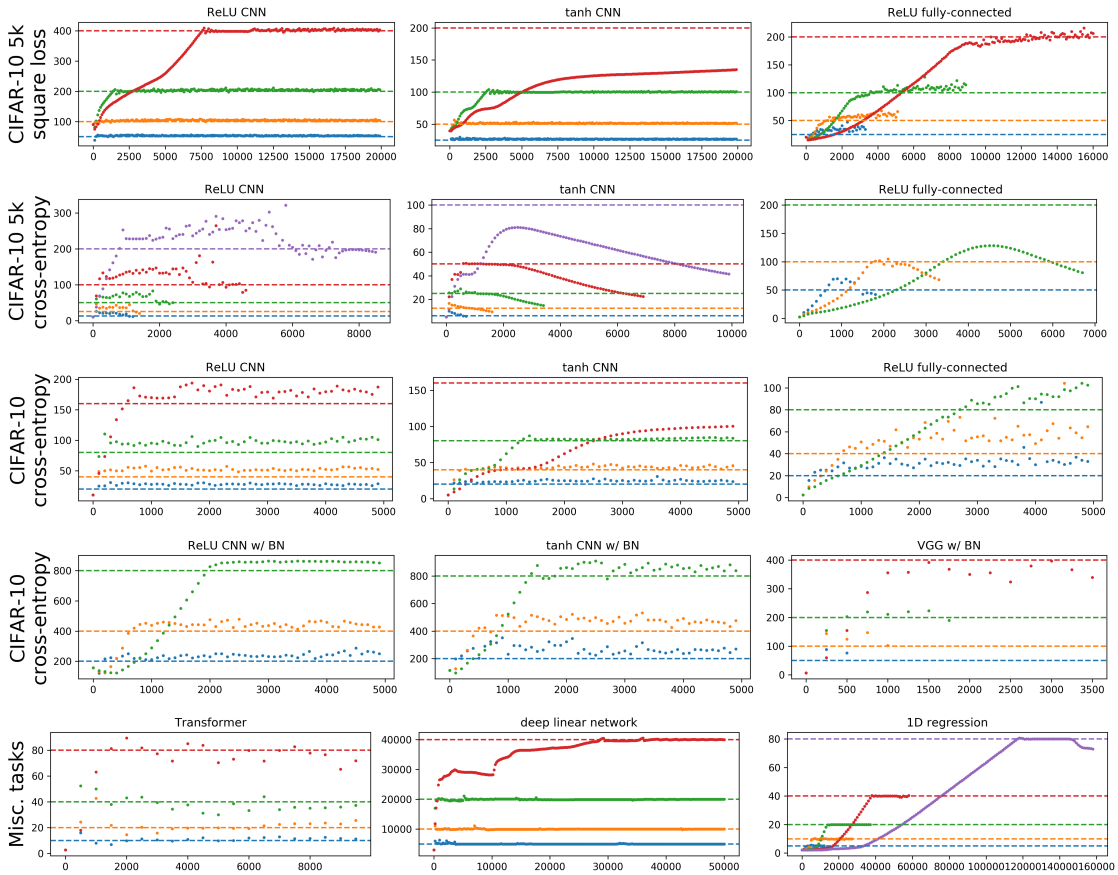

Figure 6: **Our findings hold across a diverse range of settings.** We train networks using gradient descent and measure the sharpness at regular intervals during training. Each color is a different learning rate η. The horizontal dashed line marks the value $2/\eta$. See Appendix G for plots of the train loss and train/test accuracies and additional commentary.

4 EXPERIMENTS

We now demonstrate our main point across different architectures, activation functions, loss functions, and tasks, and on networks both with and without batch normalization. Across all of these training setups, we will show that if the initial sharpness is less than $2/\eta$, then the sharpness will tend to continually rise during training (progressive sharpening) until gradient descent reaches the Edge of Stability, a regime where the sharpness hovers right at (or just above) $2/\eta$ and the training loss behaves non-monotonically. Training loss curves (and accuracies) are plotted in Appendix D.

In the first row of Figure 6, we demonstrate that our findings hold across several different architectures — a ReLU CNN, a tanh CNN, and a ReLU fully-connected network — trained using the square loss on a subset of 5k examples from CIFAR-10. In the second row, we demonstrate that our findings hold for the more commonly-used cross-entropy loss, the only difference being that the sharpness starts to decrease towards the end of training. (Here we cease training once the loss reaches 0.05, but if training were allowed to run longer, the sharpness would continue to drop.) In the third row, we verify that our findings hold on the full CIFAR-10 dataset, using three small-ish architectures (where we don't train to completion). Figure 1 shows the same point for the larger, realistic VGG-11 architecture (where we do). In the experiments on full CIFAR-10, we compute the sharpness on a subsampled set of 5k datapoints.

Batch normalization (Ioffe & Szegedy, 2015) is known to have unusual optimization properties (Li & Arora, 2019), so it is natural to wonder whether batch-normalized (BN) networks still operate at the Edge of Stability. The fourth row of Figure 6 demonstrates that the sharpness in BN networks behaves exactly the same under gradient descent as it does in non-BN networks. (As is standard practice, we use ghost batch normalization (Hoffer et al., 2017) to make full-batch gradient descent feasible with limited GPU memory.) Appendix E reconciles this finding with Santurkar et al. (2018).

Finally, in the fifth row of Figure 6, we demonstrate that our results hold on settings beyond image classification: a Transformer trained on the WikiText-2 language modeling task; a deep linear network trained on synthetic isotropic data; and a one-hidden-layer tanh network trained on a toy one-dimensional regression problem.

5 CONCLUSION

In this paper, we have empirically demonstrated that gradient flow (and gradient descent with sufficiently small step sizes) has an overwhelming tendency to continually move in a direction of increased sharpness. As a consequence of this *progressive sharpening* phenomenon, during gradient descent at reasonable learning rates, the sharpness eventually rises to $2/\eta$, whereupon gradient descent departs the gradient flow trajectory and enters the Edge of Stability, a regime where the sharpness hovers right at, or just above $2/\eta$. In this regime, gradient descent continues to successfully optimize the loss, in defiance of the conventional wisdom in optimization. We think that understanding precisely how gradient descent is able to succeed at the Edge of Stability should be viewed as an important open problem in optimization theory.

While we focused in this paper on full-batch gradient descent, we believe that some of our findings may also have analogues for SGD. First, while it is obviously normal for the training loss to behave non-monotonically during SGD, we present a surprising experiment in Appendix A which suggests that SGD on neural networks rapidly "acclimates" to the learning rate and batch size in such a way that subsequent SGD steps do not even consistently decrease the training loss *in expectation*, even though SGD steps with a smaller learning rate or larger batch size *would* consistently decrease the loss in expectation. We view this experiment as evidence that SGD on neural networks may behave more similar to gradient descent at the Edge of Stability than to classical analyses of SGD as in Bottou et al. (2018). Second, much like how the learning rate in gradient descent implicitly constrains the sharpness along the optimization trajectory, the experiments in Jastrzebski et al. (2020) indicate that a similar phenomenon may be at play during SGD, though SGD seems to implicitly constrain not just the sharpness but also the variance of the stochastic gradients. For SGD, we currently lack a precise quantitative understanding of this implicit bias. A fascinating open problem is whether SGD satisfies some "steady state" condition analogous to the sharpness $\approx 2/\eta$ rule that we have identified for full-batch gradient descent.

REFERENCES

Naman Agarwal, Zeyuan Allen-Zhu, Brian Bullins, Elad Hazan, and Tengyu Ma. Finding approximate local minima faster than gradient descent. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pp. 1195–1199, 2017.

Guillaume Alain, Nicolas Le Roux, and Pierre-Antoine Manzagol. Negative eigenvalues of the hessian in deep neural networks, 2019.

Zeyuan Allen-Zhu, Yuanzhi Li, and Zhao Song. A convergence theory for deep learning via over-parameterization. In *International Conference on Machine Learning*, pp. 242–252. PMLR, 2019.

Francis Bach. Effortless optimization through gradient flows, 2020. URL `https://francisbach.com/gradient-flows/`.

David G. T. Barrett and Benoit Dherin. Implicit gradient regularization, 2020.

Léon Bottou, Frank E. Curtis, and Jorge Nocedal. Optimization methods for large-scale machine learning. *Siam Reviews*, 60(2):223–311, 2018. URL `http://leon.bottou.org/papers/bottou-curtis-nocedal-2018`.

Sébastien Bubeck. Convex optimization: Algorithms and complexity, 2014.

Xiangyi Chen, Sijia Liu, Ruoyu Sun, and Mingyi Hong. On the convergence of a class of adam-type algorithms for non-convex optimization. *arXiv preprint arXiv:1808.02941*, 2018.

Djork-Arné Clevert, Thomas Unterthiner, and Sepp Hochreiter. Fast and accurate deep network learning by exponential linear units (elus), 2016.

Alexandre Défossez, Léon Bottou, Francis Bach, and Nicolas Usunier. On the convergence of adam and adagrad. *arXiv preprint arXiv:2003.02395*, 2020.

Saber Elaydi. *An introduction to difference equations*. Springer Science & Business Media, 2005.

Behrooz Ghorbani, Shankar Krishnan, and Ying Xiao. An investigation into neural net optimization via hessian eigenvalue density. volume 97 of *Proceedings of Machine Learning Research*, pp. 2232–2241, 2019.

Niv Giladi, Mor Shpigel Nacson, Elad Hoffer, and Daniel Soudry. At stability's edge: How to adjust hyperparameters to preserve minima selection in asynchronous training of neural networks? In *International Conference on Learning Representations*, 2020. URL `https://openreview.net/forum?id=Bkeb7lHtvH`.

Gabriel Goh. Why momentum really works. *Distill*, 2(4):e6, 2017.

Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. `http://www.deeplearningbook.org`.

Elad Hoffer, Itay Hubara, and Daniel Soudry. Train longer, generalize better: closing the generalization gap in large batch training of neural networks. In *Advances in Neural Information Processing Systems*, pp. 1731–1741, 2017.

Like Hui and Mikhail Belkin. Evaluation of neural architectures trained with square loss vs cross-entropy in classification tasks, 2020.

Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. pp. 448–456, 2015.

Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in neural information processing systems*, pp. 8571–8580, 2018.

Arthur Jacot, Franck Gabriel, and Clement Hongler. The asymptotic spectrum of the hessian of dnn throughout training. In *International Conference on Learning Representations*, 2020. URL `https://openreview.net/forum?id=SkgscaNYPS`.

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

Stanislaw Jastrzebski, Zachary Kenton, Nicolas Ballas, Asja Fischer, Yoshua Bengio, and Amost Storkey. On the relation between the sharpest directions of DNN loss and the SGD step length. In *International Conference on Learning Representations*, 2019. URL `https://openreview.net/forum?id=SkgEaj05t7`.

Stanislaw Jastrzebski, Maciej Szymczak, Stanislav Fort, Devansh Arpit, Jacek Tabor, Kyunghyun Cho*, and Krzysztof Geras*. The break-even point on optimization trajectories of deep neural networks. In *International Conference on Learning Representations*, 2020. URL `https://openreview.net/forum?id=r1g87C4KwB`.

Nitish Shirish Keskar, Dheevatsa Mudigere, Jorge Nocedal, Mikhail Smelyanskiy, and Ping Tak Peter Tang. On large-batch training for deep learning: Generalization gap and sharp minima, 2016.

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

Jaehoon Lee, Lechao Xiao, Samuel Schoenholz, Yasaman Bahri, Roman Novak, Jascha Sohl-Dickstein, and Jeffrey Pennington. Wide neural networks of any depth evolve as linear models under gradient descent. In *Advances in neural information processing systems*, pp. 8572–8583, 2019.

Aitor Lewkowycz, Yasaman Bahri, Ethan Dyer, Jascha Sohl-Dickstein, and Guy Gur-Ari. The large learning rate phase of deep learning: the catapult mechanism. *arXiv preprint arXiv:2003.02218*, 2020.

Xiaoyu Li and Francesco Orabona. On the convergence of stochastic gradient descent with adaptive stepsizes. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 983–992, 2019.

Xinyan Li, Qilong Gu, Yingxue Zhou, Tiancong Chen, and Arindam Banerjee. Hessian based analysis of sgd for deep nets: Dynamics and generalization. In *Proceedings of the 2020 SIAM International Conference on Data Mining*, pp. 190–198. SIAM, 2020.

Zhiyuan Li and Sanjeev Arora. An exponential learning rate schedule for deep learning. In *International Conference on Learning Representations*, 2019.

Chaoyue Liu, Libin Zhu, and Mikhail Belkin. Toward a theory of optimization for over-parameterized systems of non-linear equations: the lessons of deep learning, 2020.

Rotem Mulayoff and Tomer Michaeli. Unique properties of flat minima in deep networks, 2020.

Preetum Nakkiran. Learning rate annealing can provably help generalization, even for convex problems, 2020.

Kamil Nar and Shankar Sastry. Step size matters in deep learning. In *Advances in Neural Information Processing Systems*, pp. 3436–3444, 2018.

Yurii Nesterov. Introductory lectures on convex programming volume i: Basic course. 1998.

Yurii E Nesterov. A method for solving the convex programming problem with convergence rate o (1/kˆ 2). In *Dokl. akad. nauk Sssr*, volume 269, pp. 543–547, 1983.

Vardan Papyan. The full spectrum of deepnet hessians at scale: Dynamics with sgd training and sample size. *arXiv preprint arXiv:1811.07062*, 2018.

Vardan Papyan. Measurements of three-level hierarchical structure in the outliers in the spectrum of deepnet hessians. *arXiv preprint arXiv:1901.08244*, 2019.

Vardan Papyan. Traces of class/cross-class structure pervade deep learning spectra, 2020.

B.T. Polyak. Some methods of speeding up the convergence of iteration methods. *USSR Computational Mathematics and Mathematical Physics*, 1964.

Sashank Reddi, Zachary Charles, Manzil Zaheer, Zachary Garrett, Keith Rush, Jakub Konečný, Sanjiv Kumar, and H. Brendan McMahan. Adaptive federated optimization, 2020.

Sashank J Reddi, Ahmed Hefny, Suvrit Sra, Barnabas Poczos, and Alex Smola. Stochastic variance reduction for nonconvex optimization. In *International conference on machine learning*, pp. 314–323, 2016.

Levent Sagun, Utku Evci, V Ugur Guney, Yann Dauphin, and Leon Bottou. Empirical analysis of the hessian of over-parametrized neural networks. *arXiv preprint arXiv:1706.04454*, 2017.

Karthik A Sankararaman, Soham De, Zheng Xu, W Ronny Huang, and Tom Goldstein. The impact of neural network overparameterization on gradient confusion and stochastic gradient descent. *arXiv preprint arXiv:1904.06963*, 2019.

Shibani Santurkar, Dimitris Tsipras, Andrew Ilyas, and Aleksander Madry. How does batch normalization help optimization? In *Advances in Neural Information Processing Systems*, pp. 2483–2493, 2018.

Daniel Soudry, Elad Hoffer, Mor Shpigel Nacson, Suriya Gunasekar, and Nathan Srebro. The implicit bias of gradient descent on separable data. *The Journal of Machine Learning Research*, 19 (1):2822–2878, 2018.

Ilya Sutskever, James Martens, George Dahl, and Geoffrey Hinton. On the importance of initialization and momentum in deep learning. In *Proceedings of the 30th International Conference on Machine Learning*, 2013.

Sharan Vaswani, Aaron Mishkin, Issam Laradji, Mark Schmidt, Gauthier Gidel, and Simon Lacoste-Julien. Painless stochastic gradient: Interpolation, line-search, and convergence rates. In *Advances in Neural Information Processing Systems*, pp. 3732–3745, 2019.

Rachel Ward, Xiaoxia Wu, and Leon Bottou. Adagrad stepsizes: sharp convergence over nonconvex landscapes. In *International Conference on Machine Learning*, pp. 6677–6686, 2019.

Lei Wu, Chao Ma, and E Weinan. How sgd selects the global minima in over-parameterized learning: A dynamical stability perspective. In *Advances in Neural Information Processing Systems*, pp. 8279–8288, 2018.

Yuege Xie, Xiaoxia Wu, and Rachel Ward. Linear convergence of adaptive stochastic gradient descent. In *International Conference on Artificial Intelligence and Statistics*, pp. 1475–1485, 2020.

Chen Xing, Devansh Arpit, Christos Tsirigotis, and Yoshua Bengio. A walk with sgd. *arXiv preprint arXiv:1802.08770*, 2018.

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`.

Manzil Zaheer, Sashank Reddi, Devendra Sachan, Satyen Kale, and Sanjiv Kumar. Adaptive methods for nonconvex optimization. In *Advances in neural information processing systems*, pp. 9793–9803, 2018.

Jingzhao Zhang, Tianxing He, Suvrit Sra, and Ali Jadbabaie. Why gradient clipping accelerates training: A theoretical justification for adaptivity. In *International Conference on Learning Representations*, 2020. URL `https://openreview.net/forum?id=BJgnXpVYwS`.

Dongruo Zhou, Yiqi Tang, Ziyan Yang, Yuan Cao, and Quanquan Gu. On the convergence of adaptive gradient methods for nonconvex optimization. *arXiv preprint arXiv:1808.05671*, 2018.

A STOCHASTIC GRADIENT DESCENT

We have demonstrated that during full-batch gradient descent with step size η, the sharpness rises until reaching the maximum stable sharpness $2/\eta$, where it approximately remains for the duration of training. Unfortunately, prior work has shown that the evolution of the sharpness during SGD is not as easily characterized. On the one hand, just like with gradient descent, the learning rate seems to affect the sharpness along the optimization trajectory, with larger learning rates (as well as smaller batch sizes) steering SGD towards less sharp regions of the loss landscape (Jastrzebski et al., 2019; 2020). On the other hand, in behavior different from gradient descent, during SGD the sharpness sometimes does not equilibrate at any steady-state value, much less one that we can numerically predict; for example, in Figure 5(a) of Jastrzebski et al. (2020), one can see that the sharpness gradually drifts downwards during SGD training. We suspect that this is partly because the sharpness does not directly affect the stability of SGD unless one makes the strong assumption that the Hessian and the covariance matrix of per-example gradients share a leading eigenvector, as in Jastrzebski et al. (2020).

That said, we do suspect that there may exist some sense in which SGD, too, operates at the Edge of Stability. In particular, we now demonstrate that SGD seems to "acclimate" to the learning rate and batch size in such a way that an actual step sometimes increases and sometimes decreases the loss in expectation, yet a step with a larger learning rate or smaller batch size would almost always *increase* the loss in expectation, and a step with a smaller learning rate or a larger batch size would almost always *decrease* the loss in expectation.

In Figure 7, we train the ELU network from Section 3 using SGD with learning rate 0.01 and batch size 32. We periodically compute the training loss (over the full dataset) and plot these on the left pane of Figure 7. Observe that the training loss does not decrease monotonically, but of course this is not surprising — SGD is a random algorithm. However, what may be more surprising is that SGD is not even decreasing the training loss *in expectation*. On the right pane of Figure 7, every 500 steps during training, we use the Monte Carlo method to approximately compute the expected change in training loss that would result from an SGD step (the expectation here is over the randomness involved in selecting the minibatch). Observe that at many points during training, an SGD step would decrease the loss (as desired) in expectation, but at other points, and SGD step would increase the loss in expectation.

In Figure 8(a), during training of the same network, we compute the expected change in training loss that would result from taking an SGD step with the same learning rate used during training, but *half the batch size* used during training (i.e. 16). We observe that an SGD step with half the batch size would consistently cause an *increase* in the training loss in expectation. In Figure 8(b) we repeat this experiment, but with *twice* the batch size used during training (i.e. 64). Notice that an SGD step with twice the batch size would consistently cause a *decrease* in the training loss in expectation. In Figure 8(c) and (d) repeat this experiment with the learning rate; we observe that an SGD step with a larger learning rate would consistently increase the training loss in expectation, while an SGD step with a smaller learning rate would consistently decrease the training loss in expectation.

To be clear, there was nothing special in absolute terms about the hyperparameters we used. In Figure 9(a), we train the same network with learning rate 0.01 and batch size 16 and evaluate the expected loss change during training. Observe that some steps increase the loss in expectation, while others decrease the loss in expectation. Figure 9(b)-(d) repeats this experiment with: (b) learning rate 0.01 and batch size 64, (c) learning rate 0.02 and batch size 32, and (d) learning rate 0.005 and batch size 32.

We view this experiment as evidence that SGD on neural networks behaves more similar to gradient descent at the Edge of Stability, than to conventional analyses of SGD like the ones in Bottou et al. (2018).

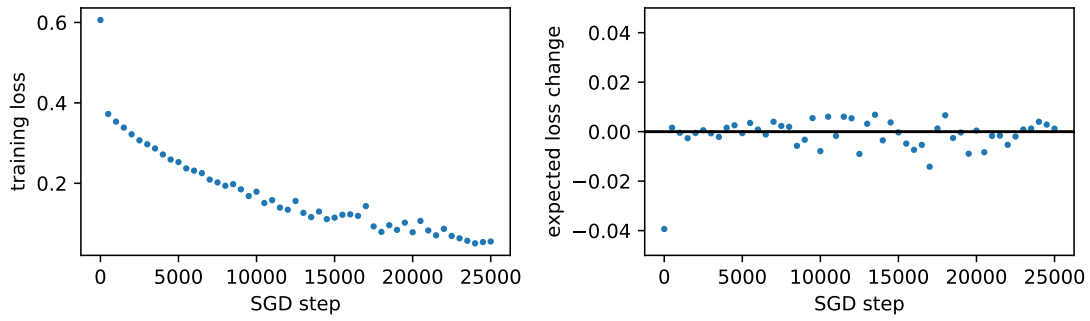

Figure 7: **SGD does not consistently decrease the training loss in expectation.** We train the ELU network from section 3 using SGD with learning rate 0.01 and batch size 32. At regular intervals during training, we compute (left) the full-batch training loss, and (right) the expected change in the full-batch training loss (where the expectation is over the randomness in sampling the minibatch). Strikingly, note that after the very beginning of training, the expected loss change is sometimes negative (as desired) but oftentimes positive. See Figure 8.

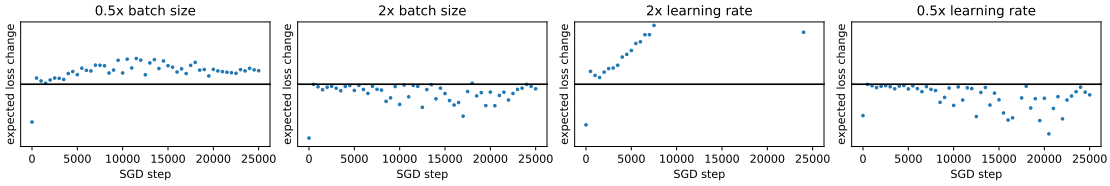

Figure 8: **An SGD step with a smaller learning rate or a larger batch size than the ones used during training *would* consistently decrease the loss in expectation.** At regular intervals during the training run depicted in Figure 7 (with $\eta = 0.01$ and batch size 32), we measure the expected change in the full-batch training loss that would result from an SGD step with a different learning rate or batch size. Observe that taking an SGD step with a smaller learning rate or a larger batch size would consistently have decreased the loss in expectation, while taking an SGD step with a larger learning rate or a smaller batch size would have consistently increased the loss in expectation.

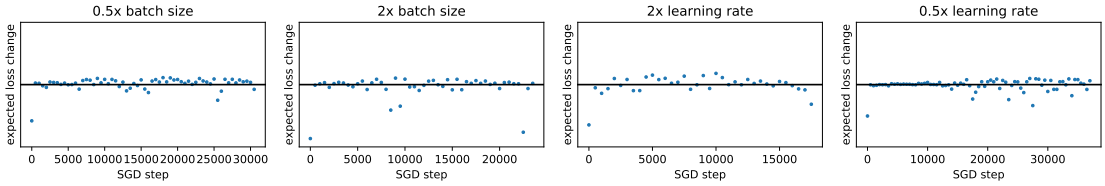

Figure 9: **There is nothing special about the hyperparameters we have used.** Figure 7 showed that if we train a network with $\eta = 0.01$ and batch size 32, the train loss does not consistently decrease in expectation, yet 8(b) showed that at every point during training that network, an SGD step with batch size 64 would have decreased the loss in expectation. A natural question is what would have happened had we trained with $\eta = 0.01$ and batch size 64. Figure 9(b) demonstrates that the train loss would not consistently decrease in expectation. The other panes are analogous.

B STABILITY OF GRADIENT DESCENT ON QUADRATIC FUNCTIONS

This appendix describes the stability properties of gradient descent (and its momentum variants) when optimizing the quadratic objective function

$$f(\mathbf{x}) = \frac{1}{2}\mathbf{x}^T\mathbf{H}\mathbf{x} + \mathbf{g}^T\mathbf{x} + c \tag{2}$$

starting from the initialization $\mathbf{x}_0 = \mathbf{0}$. Our notation is intentionally chosen to resemble a second-order Taylor approximation.

To review, vanilla gradient descent is defined by the iteration:

$$\mathbf{x}_{t+1} = \mathbf{x}_t - \eta\nabla f(\mathbf{x}_t).$$

Meanwhile, gradient descent with Polyak (also called "heavy ball") momentum (Polyak, 1964; Sutskever et al., 2013; Goodfellow et al., 2016) is defined by the iteration:

$$\begin{aligned}\mathbf{v}_{t+1} &= \beta\mathbf{v}_t - \eta\nabla f(\mathbf{x}_t)\\ \mathbf{x}_{t+1} &= \mathbf{x}_t + \mathbf{v}_{t+1}\end{aligned}$$

where $\mathbf{v}_t$ is a "velocity" vector and $0 \leq \beta < 1$ is the momentum coefficient. For $\beta = 0$ the algorithm reduces to vanilla GD.

Finally, Nesterov momentum Sutskever et al. (2013); Goodfellow et al. (2016) is an adaptation of Nesterov's accelerated gradient (Nesterov, 1983) for deep learning defined by the iteration:

$$\begin{aligned}\mathbf{v}_{t+1} &= \beta\mathbf{v}_t - \eta\nabla f(\mathbf{x}_t + \beta\mathbf{v}_t)\\ \mathbf{x}_{t+1} &= \mathbf{x}_t + \mathbf{v}_{t+1}\end{aligned}$$

where $\mathbf{v}_t$ is a "velocity" vector and $0 \leq \beta < 1$ is the momentum coefficient. For $\beta = 0$ the algorithm reduces to vanilla GD.

All three of these algorithms share a special property: on quadratic functions, they act independently along each Hessian eigenvector. That is, if we express the iterates in the Hessian eigenvector basis, then in this basis the coordinates evolve independent from one another under gradient descent.

Proposition 1. *Consider running vanilla gradient descent on the quadratic objective (2) starting from $\mathbf{x}_0 = \mathbf{0}$. Let $\mathbf{q}$ be the leading eigenvector of $\mathbf{H}$ and let λ be the sharpness. (Assume that $\lambda > 0$.) Then the sequence $\{\mathbf{q}^T\mathbf{x}_t\}$ will diverge iff $\lambda > 2/\eta$.*

Note that in the main text, we referred to $\mathbf{q}^T\mathbf{x}_t$ as Δ_t.

Proof. The update rule for gradient descent on this quadratic function is:

$$\mathbf{x}_{t+1} = (\mathbf{I} - \eta\mathbf{H})\mathbf{x}_t - \eta\mathbf{g}.$$

Therefore, the quantity $\mathbf{q}^T\mathbf{x}_t$ evolves under gradient descent as:

$$\mathbf{q}^T\mathbf{x}_{t+1} = (1 - \eta\lambda)\mathbf{q}^T\mathbf{x}_t - \eta\,\mathbf{q}^T\mathbf{g}$$

Define $\tilde{x}_t = \mathbf{q}^T\mathbf{x}_t + \frac{1}{\lambda}\mathbf{q}^T\mathbf{g}$, and note that $\{\mathbf{q}^T\mathbf{x}_t\}$ diverges if and only if $\{\tilde{x}_t\}$ diverges.

The quantity $\tilde{x}_t$ evolves under gradient descent according to the simple rule:

$$\tilde{x}_{t+1} = (1 - \eta\lambda)\tilde{x}_t$$

The sequence $\tilde{x}_t$ will diverge iff $|1-\eta\lambda| > 1$. Since $\eta > 0$ and $\lambda > 0$, it is impossible for $1-\eta\lambda > 1$, so $\tilde{x}_t$ diverges if and only if $1 - \eta\lambda < -1 \iff \lambda > 2/\eta$. □

Remark: to recover the formula given in the main text for the evolution of $\mathbf{q}^T\mathbf{x}_t$, note that $\mathbf{q}^T\mathbf{x}_t = \tilde{x}_t - \frac{1}{\lambda}\mathbf{q}^T\mathbf{g} = (1-\eta\lambda)^t\tilde{x}_0 - \frac{1}{\lambda}\mathbf{q}^T\mathbf{g} = ((1-\eta\lambda)^t - 1)(\frac{1}{\lambda}\mathbf{q}^T\mathbf{g})$.

We now prove analogous results for Nesterov and Polyak momentum.

Theorem 1. *Consider running Nesterov momentum on the quadratic objective (2) starting from* $\mathbf{x}_0 = \mathbf{0}$*. Let* $\mathbf{q}$ *be the leading eigenvector of* $\mathbf{H}$ *and let* λ *be the sharpness. (Assume that* $\lambda > 0$*.) Then the sequence* $\{\mathbf{q}^T\mathbf{x}_t\}$ *will diverge iff* $\lambda > \frac{1}{\eta}\left(\frac{2+2\beta}{1+2\beta}\right)$.

Proof. The update rules for Nesterov momentum on this quadratic function are:

$$\begin{aligned}\mathbf{v}_{t+1} &= \beta(\mathbf{I} - \eta\mathbf{H})\mathbf{v}_t - \eta\mathbf{g} - \eta\mathbf{H}\mathbf{x}_t \\ \mathbf{x}_{t+1} &= \mathbf{x}_t + \mathbf{v}_{t+1}.\end{aligned}$$

Define $\tilde{x}_t = \mathbf{q}^T\mathbf{x}_t + \frac{1}{\lambda}\mathbf{q}^T\mathbf{g}$ and $\tilde{v}_t = \mathbf{q}^T\mathbf{v}_t$. Note that $\mathbf{q}^T\mathbf{x}_t$ diverges iff $\tilde{x}_t$ diverges.

The quantities $\tilde{x}_t$ and $\tilde{v}_t$ evolve under Nesterov momentum as:

$$\begin{aligned}\tilde{v}_{t+1} &= \beta(1 - \eta\lambda)\tilde{v}_t - \eta\lambda\tilde{x}_t \\ \tilde{x}_{t+1} &= \tilde{x}_t + \tilde{v}_{t+1}\end{aligned}$$

By noting that $\tilde{v}_t = \tilde{x}_t - \tilde{x}_{t-1}$, we can rewrite this as a recurrence in $\tilde{x}$:

$$\begin{aligned}\tilde{x}_{t+1} &= \tilde{x}_t + \beta(1 - \eta\lambda)(\tilde{x}_t - \tilde{x}_{t-1}) - \eta\lambda\tilde{x}_t \\ &= (1 - \eta\lambda)(1 + \beta)\tilde{x}_t - \beta(1 - \eta\lambda)\tilde{x}_{t-1}\end{aligned}$$

This is a linear homogenous second-order difference equation. By Theorem 2.37 in Elaydi (2005), since $\lambda > 0$ and $\eta > 0$ and $\beta < 1$, this recurrence diverges if and only if $\lambda > \frac{1}{\eta}\left(\frac{2+2\beta}{1+2\beta}\right)$.

□

The following result previously appeared in Goh (2017).

Theorem 2. *Consider running Polyak momentum on the quadratic objective (2) starting from* $\mathbf{x}_0 = 0$*. Let* $\mathbf{q}$ *be the leading eigenvector of* $\mathbf{H}$ *and let* λ *be the sharpness. (Assume that* $\lambda > 0$*.) Then the sequence* $\{\mathbf{q}^T\mathbf{x}_t\}$ *will diverge iff* $\lambda > \frac{1}{\eta}(2 + 2\beta)$.

Proof. The update rules for Polyak momentum on this quadratic function are:

$$\begin{aligned}\mathbf{v}_{t+1} &= \beta\mathbf{v}_t - \eta\mathbf{H}\mathbf{x}_t - \eta\mathbf{g} \\ \mathbf{x}_{t+1} &= \mathbf{x}_t + \mathbf{v}_{t+1}.\end{aligned}$$

Define $\tilde{x}_t = \mathbf{q}^T\mathbf{x}_t + \frac{1}{\lambda}\mathbf{q}^T\mathbf{g}$ and $\tilde{v}_t = \mathbf{q}^T\mathbf{v}_t$. Note that $\mathbf{q}^T\mathbf{x}_t$ diverges iff $\tilde{x}_t$ diverges.

The quantities $\tilde{x}_t$ and $\tilde{v}_t$ evolve under Polyak momentum as:

$$\begin{aligned}\tilde{v}_{t+1} &= \beta\tilde{v}_t - \eta\lambda\tilde{x}_t \\ \tilde{x}_{t+1} &= \tilde{x}_t + \tilde{v}_{t+1}.\end{aligned}$$

By noting that $\tilde{v}_t = \tilde{x}_t - \tilde{x}_{t-1}$, we can rewrite this as a recurrence in $\tilde{x}$:

$$\begin{aligned}\tilde{x}_{t+1} &= \tilde{x}_t + \beta(\tilde{x}_t - \tilde{x}_{t-1}) - \eta\lambda\tilde{x}_t \\ &= (1 + \beta - \eta\lambda)\tilde{x}_t - \beta\tilde{x}_{t-1}.\end{aligned}$$

This is a linear homogenous second-order difference equation. By Theorem 2.37 in Elaydi (2005), since $\lambda > 0$ and $\eta > 0$ and $\beta < 1$, this recurrence diverges if and only if $\lambda > \frac{1}{\eta}(2 + 2\beta)$.

□

None of these three results have used the fact that $(\mathbf{q}, \lambda)$ are the *leading* eigenvector and eigenvalue; they would also apply to any eigenvector/eigenvalue pair $(\mathbf{q}_i, \lambda_i)$. Thus, along any direction of positive curvature (i.e. $\lambda_i > 0$), gradient descent (vanilla, or with Polyak/Nesterov momentum) diverges iff λ_i is strictly larger than some threshold, and converges iff λ_i is strictly smaller than that threshold. If the quadratic is convex (i.e. $\mathbf{H} \succ 0$), this means that gradient descent will converge to

the solution iff the sharpness is strictly less than that threshold, and will diverge iff the sharpness is strictly greater than that threshold.

If $\mathbf{H}$ has negative eigenvalues, then gradient descent (with a positive step size) will diverge along the directions of negative curvature. Of course, neural network training Hessians do have negative eigenvalues (Sagun et al., 2017; Ghorbani et al., 2019), yet gradient descent does not diverge during normal training. The reason may be that the directions of negative curvature are somewhat illusory: Alain et al. (2019) walked along directions of negative curvature, and found that the negative curvature did not last for very long, whereas the directions of most positive curvature lasted for longer. In other words, the quadratic Taylor approximation was accurate over greater distances for the directions of positive curvature than for the directions of negative curvature.

C ADDITIONAL FIGURES

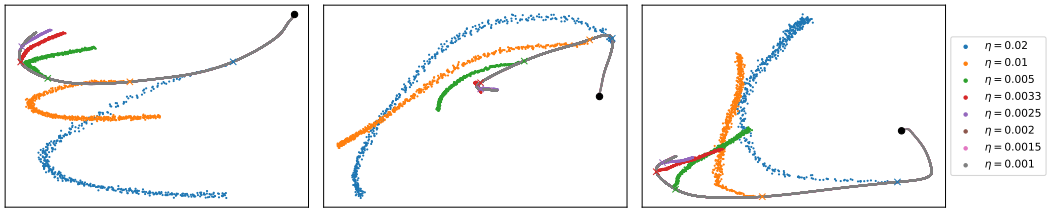

Figure 10: **Gradient descent with step size η roughly tracks the gradient flow trajectory until the sharpness reaches $2/\eta$, and then peels off from the gradient flow trajectory.** This figure visualizes the optimization trajectories of the ELU network that we trained in section 3. To reduce the iterates from dimension 656810 to dimension 2 so that they can be visualized, we randomly generated two vectors of length 656810 from the standard normal distribution, and we computed the inner product of these two vectors with the iterates (i.e. the random projection method). We repeated this procedure three times, resulting in the three panes in the figure above. For each learning rate, we draw an "x" at the spot where the sharpness first crosses $2/\eta$. The gray curve is approximately the gradient flow trajectory. Observe that at each learning rate η, gradient descent tracks the gradient flow trajectory until the point on that trajectory where the sharpness reaches $2/\eta$; after this point, gradient descent follows a different route. At learning rates 0.002, 0.0015, and 0.001, gradient descent terminates before reaching the point where the sharpness hits $2/\eta$, which is why these three curves are directly on top of one another, and only the gray curve ($\eta = 0.001$) is visible.

	0.001	**0.0015**	**0.002**	**0.0025**	**0.0033**	**0.005**	**0.01**	**0.02**
0.001	1.0000	0.9998	0.9997	0.9279	0.9048	0.7972	0.7794	0.6782
0.0015	0.9998	1.0000	0.9999	0.9281	0.9048	0.7973	0.7794	0.6784
0.002	0.9997	0.9999	1.0000	0.9280	0.9047	0.7974	0.7793	0.6783
0.0025	0.9279	0.9281	0.9280	1.0000	0.9201	0.7855	0.7728	0.6713
0.0033	0.9048	0.9048	0.9047	0.9201	1.0000	0.7721	0.8043	0.6885
0.005	0.7972	0.7973	0.7974	0.7855	0.7721	1.0000	0.7547	0.7019
0.01	0.7794	0.7794	0.7793	0.7728	0.8043	0.7547	1.0000	0.7951
0.02	0.6782	0.6784	0.6783	0.6713	0.6885	0.7019	0.7951	1.0000

Table 1: **At the three learning rates η small enough that sharpness along the gradient flow trajectory never reached $2/\eta$, gradient descent learned approximately the same function**. This table shows the similarities between the functions learned by the ELU architecture from section 3 when we train using gradient descent at different learning rates. To measure the similarity between two networks, we compute the fraction of the CIFAR-10 test set (size 10k) on which their class predictions agree. Observe that while the network trained at $\eta = 0.001$ is almost completely identical to the networks trained at $\eta = 0.0015$ and $\eta = 0.002$ (it agrees with them on 9998/10k and 9997/10k examples, respectively), it is quite different from even the network trained at $\eta = 0.0025$. This is because the $\eta = 0.0025$ network (as well as the networks with higher learning rates) departed the gradient flow trajectory and entered the Edge of Stability.

Note that even on the three examples where the $\eta = 0.001$ and $\eta = 0.002$ networks made different predictions, the differences in the class scores were slight. For instance, on one of these test examples, the $\eta = 0.001$ network predicted 4 while the $\eta = 0.002$ network predicted 7; the scores were:

class	0	1	2	3	4	5	6	7	8	9
$\eta = 0.001$	-0.3177	0.1639	-0.1745	-0.0889	0.8036	-0.3633	0.3287	0.8034	-0.0943	-0.1370
$\eta = 0.002$	-0.3178	0.1641	-0.1745	-0.0888	0.8033	-0.3635	0.3290	0.8033	-0.0943	-0.1370

D INFINITE WIDTH LIMITS

In this appendix, we conduct an experiment which aims to reconcile progressive sharpening with the theory of infinite-width neural networks. It is well-known that, under a certain network parameterization and initialization, when training infinitely wide neural networks using gradient flow or gradient descent with sufficiently small step sizes, the parameters move almost no distance from their initial position, and therefore the Hessian does not change from initialization (Jacot et al., 2018; Lee et al., 2019; Jacot et al., 2020). Accordingly, one might hypothesize that progressive sharpening might attenuate as networks (with NTK parameterization) become increasingly wide.

We consider fully-connected networks under the NTK parameterization (Lee et al., 2019) with two hidden layers and tanh activations. We consider networks with widths 100, 200, 400, and 800. We train these four networks on the Fashion-MNIST dataset using the cross-entropy loss, with a small learning rate (aiming to mimic gradient flow). We use the same learning rate for each network (this is sensible under the NTK parameterization). We stop each network when the loss reaches a certain threshold.

In Figure 11(b), we plot the sharpness by iteration. One can see that sharpening happens quicker for the narrow networks than for the wider networks. However, there is a confounding factor: as can be seen from Figure 11(a), the narrower networks train faster. To remove this confounding factor, in Figure 11(c), we plot sharpness as a function of train loss. Here one can clearly see that at every train loss, the narrower (resp: wider) networks have a higher (resp: lower) sharpness. This observation aligns with NTK theory.

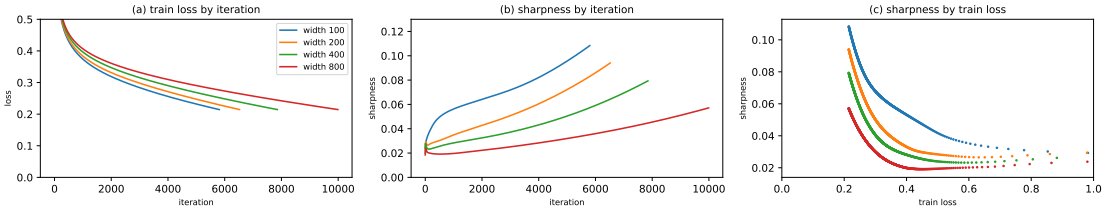

Figure 11: **Under the NTK parameterization, wider networks find a less-sharp solution for any given value of the training loss.**

E RECONCLING OUR RESULTS WITH SANTURKAR ET AL. (2018)

In this appendix, we reconcile the apparent contradiction between our results and those of the well-known work Santurkar et al. (2018).

Aiming to explain the effectiveness of batch normalization (BN), Santurkar et al. (2018) conjectured that BN smooths the optimization landscape, in three different senses of the word "smooth":

1. BN improves the Lipschitzness of the objective function along the optimization trajectory. To demonstrate this, at regular intervals during training, Santurkar et al. (2018) measured the change in loss that would arise from taking a step in the negative gradient direction. (Note that the step size used here was different from the training step size.)
2. BN improves the *gradient predictiveness* of the objective function along the optimization trajectory. Here, gradient predictiveness is defined as the ℓ_2 distance between the gradient at the current iterate, and the gradient after taking a very large step in the negative gradient direction. To demonstrate this, Santurkar et al. (2018) measured the gradient predictiveness at regular intervals during training.
3. BN improves the *effective smoothness* along the optimization trajectory, defined as the Lipschitz constant of the gradient in the direction of the negative gradient. Given an objective function f, an iterate θ, and a step size α, the effective smoothness is defined as
$$\sup_{\gamma \in [0,\alpha]} \frac{\|\nabla f(\theta) - \nabla f(\theta - \gamma \nabla f(\theta))\|_2}{\|\gamma \nabla f(\theta)\|_2}.$$
To demonstrate this, Santurkar et al. (2018) measured the effective smoothness at regular intervals during training.

Note that Santurkar et al. (2018) did *not* conjecture that BN improves the L-smoothness (a.k.a decrease the sharpness) along the optimization trajectory — only that BN improves the *effective* smoothness, which is the L-smoothness in the negative gradient direction.

Our main paper demonstrates that BN does not improve the L-smoothness along the optimization trajectory, at least for full-batch gradient descent. During gradient descent with step size η, for both BN and non-BN networks, if the sharpness is less than $2/\eta$, then the sharpness eventually rises to approximately $2/\eta$, where it remains for the duration of training. Thus, in the long run, for both BN and non-BN networks, the smoothness settles at approximately the same value — and moreover, this is precisely the value that voids the descent lemma (and arguments based on L-smoothness).

Moreover, we now demonstrate that BN does not improve the effective smoothness either (at least during full-batch gradient descent). In Figure 12, we train a ReLU CNN both with (top row) and without (bottom row) batch normalization, at the same grid of learning rates, using the cross-entropy loss on a subset of 5k examples from CIFAR-10. (For the BN network we use ghost batch normalization with two ghost batches of size 2,500, since a batch with all 5k examples does not fit into GPU memory.) We measure both the sharpness/smoothness and the effective smoothness at regular intervals during training. (We numerically evaluate the $\sup$ in the definition of effective smoothness using a grid of 10 evenly spaced points.) Observe that for each learning rate, the effective smoothness behaves similar to the smoothness: it rises to $2/\eta$ and then remains approximately there, until the end of training, when both values drop because we are using the cross-entropy loss. Thus, BN does not improve the effective smoothness along the optimization trajectory. Despite this, one can see that the BN network trains substantially faster than the non-BN network.

Note that this finding is actually consistent with Figure 4(c) and Figure 11(d) in Santurkar et al. (2018), which are meant to show, respectively, that both BN and their ℓ_p normalization schemes improve effective smoothness on a VGG network. These two plots show that for both BN and non-BN networks trained using SGD with step size $\eta = 0.1$, the effective smoothness hovers around the value $20 = 2/\eta$, which is exactly what we observe. Their plots do show that the effective smoothness behaves more *regularly* for the batch-normalized network than for the non-BN network. But we strongly disagree with their interpretation of these two figures as demonstrating that BN improves the effective smoothness during training.

The third and final piece of evidence in Santurkar et al. (2018) for the conjecture that BN improves the effective smoothness is their Figure 9(c), which shows that BN improves the effective smooth-

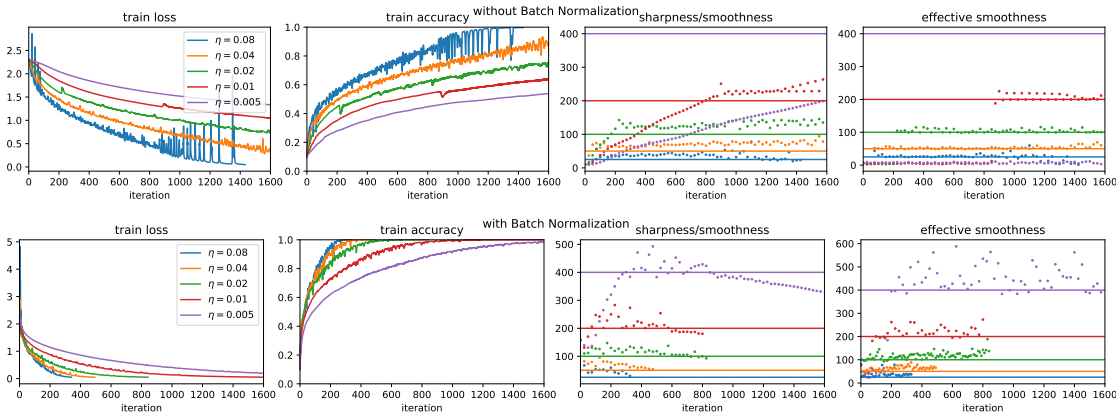

Figure 12: **Batch normalization does not improve the smoothness or the effective smoothness along the optimization trajectory.**

ness on a deep linear network. For this figure, the step size α used in their computation of effective smoothness was larger than the training step size η by a factor of $3 \cdot 10^7$. (For the VGG network, α was larger than η by a factor of 4.) The effective smoothness at this massive distance has no bearing on training.

F EXPERIMENTAL DETAILS

F.1 TASKS

We used the following tasks:

- **CIFAR-10** (Krizhevsky, 2009): we standardize each channel by subtracting the mean and dividing by the standard deviation.
- **CIFAR-10 5k**: this is a subset of CIFAR-10 of size 5,000. In particular, we retain every 10-th train example.
- **Wikitext-2**: this is a language modeling dataset. We use `bptt = 35`, i.e. we predict text in contiguous blocks of 35 words. We use only the first 4970 (it is divisible by 35) training examples in the dataset.
- **Deep linear network task**: The task is to map the matrix $\mathbf{X} \in \mathbb{R}^{n \times d} = [\mathbf{x}_1, \quad \ldots, \quad \mathbf{x}_n]^T$ to the matrix $\mathbf{Y} \in \mathbb{R}^{n \times d} = [\mathbf{y}_1, \quad \ldots, \quad \mathbf{y}_n]^T$ using a deep linear network $f : \mathbb{R}^d \to \mathbb{R}^d$ that consists of a stack of L matrices each of size $d \times d$. Error is measured using the squared loss: $\frac{1}{n} \sum_{i=1}^{n} \|f(\mathbf{x}_i) - \mathbf{y}_i\|_2^2$.

  We use $n = d = 50$. We generate $\mathbf{X}$ as a random whitened matrix (i.e. $\frac{1}{n}\mathbf{X}^T\mathbf{X} = \mathbf{I}$) and generate $\mathbf{Y} = \mathbf{X}\mathbf{A}^T$, where $\mathbf{A} \in \mathbb{R}^{d \times d}$ is a random matrix of standard Gaussians.

  To generate $\mathbf{X}$ as a random whitened matrix, we sample a 50×50 matrix of standard Gaussians, then we let $\mathbf{X}$ be $\sqrt{50}$ times the Q factor in the QR factorization of that matrix.
- **Toy regression task**: This is a one-dimensional regression task using the square loss. We sample 20 datapoints uniformly spaced between -1 and 1, and we label these datapoints noiselessly using the Chebyshev polynomial of degree 5. The dataset is plotted in Figure 13.

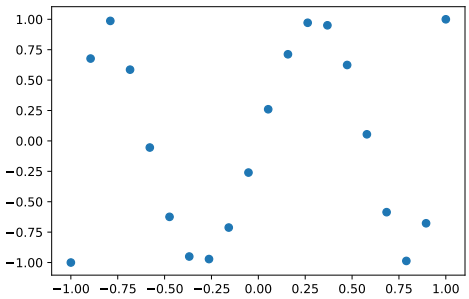

Figure 13: The dataset for the toy one-dimensional regression task.

F.2 ARCHITECTURES

- **VGG-11 (no BN)**: we took the VGG-11 sized for CIFAR-10 available at `https://github.com/chengyangfu/pytorch-vgg-cifar10`, and we removed the dropout layers. In Figure 1, we "warm-started" optimization by running GD for 400 iterations at learning rate 0.02. The reason we do this is that loss landscape is very flat at initialization, and if we did not do this, the network trained at, say, $\eta = 0.00125$ would basically not budge for the first several thousand iterations of training.
- **VGG-11 (BN)**: same as above (though we do not need to warm start training). Since GPUs do not have enough memory to do a full-batch step, we use ghost batch normalization (Hoffer et al., 2017) with 50 ghost batches of size 1,000.
- **ReLU CNN**: The PyTorch code for this network is:

```
nn.Sequential(
    nn.Conv2d(3, 32, bias=True, kernel_size=3, padding=1),
    nn.ReLU(),
    nn.MaxPool2d(2),
    nn.Conv2d(32, 32, bias=True, kernel_size=3, padding=1),
    nn.ReLU(),
    nn.MaxPool2d(2),
    nn.Flatten(),
    nn.Linear(2048, 10, bias=True)
)
```

  We use the default PyTorch initialization.
- **ReLU CNN (BN)** Same as the ReLU CNN, but we insert a `nn.BatchNorm2d` layer after each ReLU. We use 2 ghost batches of size 2,500 each.
- **Tanh CNN and Tanh CNN (BN)**: same as the ReLU CNN, except we use the tanh activation instead of ReLU.
- **ReLU fully-connected network** The PyTorch code for this network is:

```
nn.Sequential(
    nn.Linear(3072, 100, bias=True),
    nn.ReLU(),
    nn.Linear(100, 100, bias=True)
    nn.ReLU(),
    nn.Linear(100, 10, bias=True)
)
```

  We use the default PyTorch initialization.
- **Tanh fully-connected network**: Same as the ReLU fully-connected network, but with tanh activations.
- **ELU fully-connected network**: Same as the ReLU fully-connected network, but with ELU activations, and with hidden layer size 200 rather than 100. (There is no nefarious reason for the difference in hidden layer size; we just ran the ReLU and tanh experiments first, and then decided to make the network more overparameterized.)
- **Deep linear network**: this is a deep linear network with 20 layers, each a 50×50 matrix. We initialize each matrix using Xavier initialization: each entry is drawn from $\mathcal{N}(0, \frac{1}{50})$.
- **Toy 1d tanh network**: The PyTorch code for this network is:

```
nn.Sequential(
    nn.Linear(1, 100, bias=True),
    nn.Tanh(),
    nn.Linear(100, 1, bias=False),
 )
```

  We initialize using Xavier initialization.
- **Transformer**: we use the Transformer from the Pytorch language modeling tutorial: `https://github.com/pytorch/examples/tree/master/word_language_model` with ninp=200, nhead=2, nhid=200, nlayers=2, dropout=0.

G COMPLETE EXPERIMENT PLOTS

In this section, we present complete plots for the experiments described in Section 4. The purpose of these experiments is to demonstrate the following points:

1. When sharpness is less than $2/\eta$, gradient descent has an overwhelming tendency to increase the sharpness (*progressive sharpening*).
2. So long as the sharpness is less than $2/\eta$, gradient descent at different learning rates roughly follows the same path, just at different speeds. This is because gradient descent is roughly tracking the gradient flow trajectory (see the animation in Bach (2020) for a visualization of what we mean). The way we demonstrate this is by plotting the train loss and the sharpness by time; when plotted by time, the sharpness and training loss of GD at different learning rates almost exactly coincide.
3. If the sharpness rises to $2/\eta$ (see the caveat below), gradient decent subsequently enters the Edge of Stability, a regime where (1) the training loss typically behaves non-monotonically, and (2) the sharpness hovers right at, or just above, the value $2/\eta$, rather than increasing further.
4. For sufficiently small learning rates η, gradient descent can terminate before the sharpness reaches $2/\eta$. We will demonstrate that these small learning rates are highly suboptimal in terms of convergence speed — therefore, we say that they are not "reasonable."

   On some problems, it is feasible to actually train to completion at these small learning rates. In this case, our experiments demonstrate that larger learning rates (which enter the Edge of Stability) train faster. On other problems, sharpening seems to occur so rapidly that training to completion at these small learning rates is not a practical experiment for us to run; however, our gradient flow thought experiment from section 3 implies that these small learning rates should exist.

There is one slight caveat. On a few optimization problems (e.g. Figure 26), we observe that the gradient descent enters the Edge of Stability regime when the sharpness is a little bit smaller than $2/\eta$. This is because we are only measuring the sharpness right at the iterates themselves. However, if the sharpness *in between* the iterates (which we do not measure) exceeds $2/\eta$, then it is possible for a gradient step to increase the training loss, and for gradient descent to become destabilized.

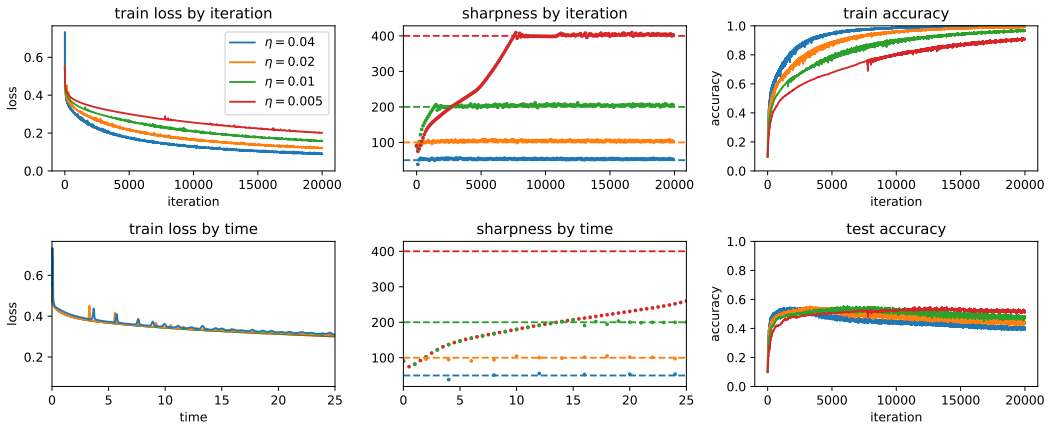

Figure 14: **ReLU CNN trained with square loss on 5k subset of CIFAR-10.** (1) For the blue learning rate, the initial sharpness was greater than $2/\eta$, so gradient descent catapulted at the beginning of training (which is why the blue train loss doesn't coincide with the others when plotted by time). (2) When we plot the sharpness by time, the red and green dots almost (but not exactly) overlap before the green one enters the Edge of Stability, the red and green learning rates are roughly following the same path.

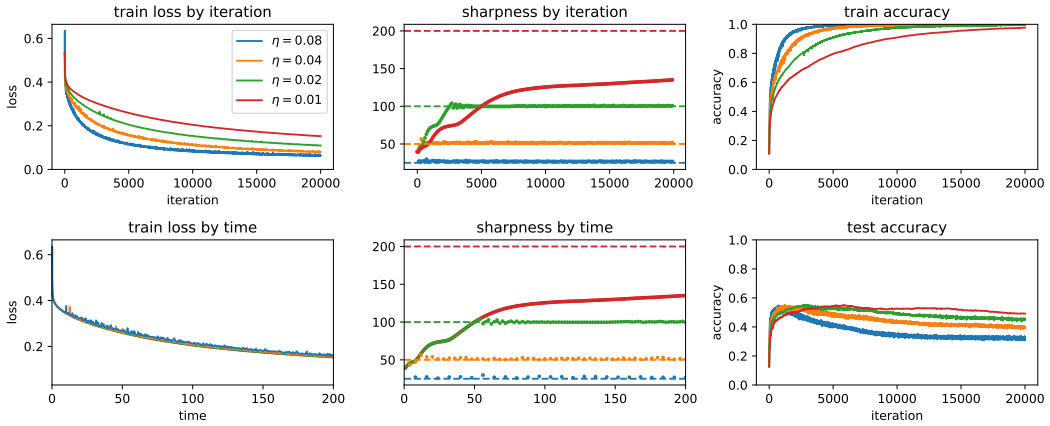

Figure 15: **Tanh CNN trained with square loss on 5k subset of CIFAR-10.** The blue learning rate is large enough that the initial sharpness was greater than $2/\eta$.

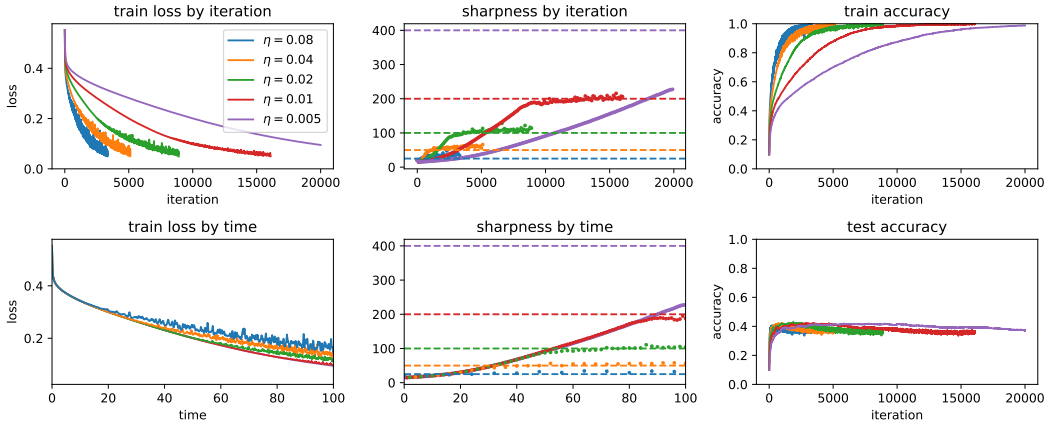

Figure 16: **ReLU fully-connected network trained with square loss on 5k subset of CIFAR-10.**. Here we use the cross-entropy loss, so the sharpness drops near the end of training.

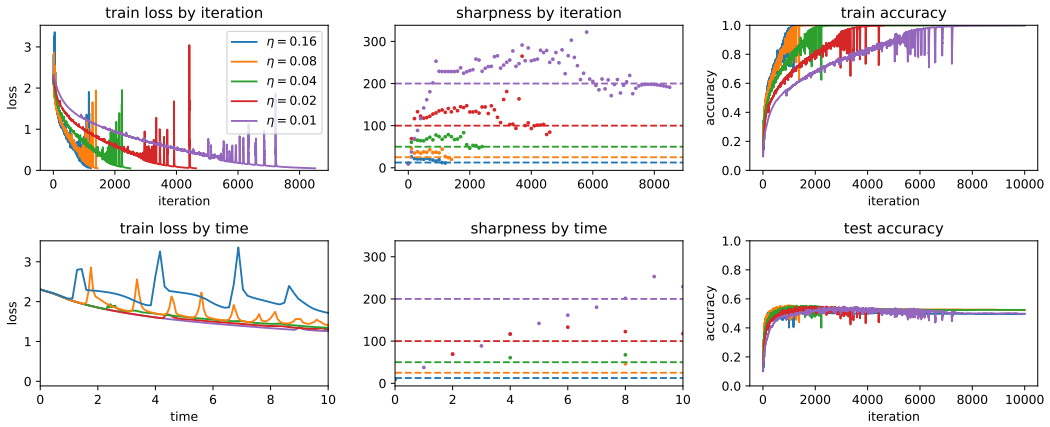

Figure 17: **ReLU CNN trained with cross-entropy on 5k subset of CIFAR-10.**

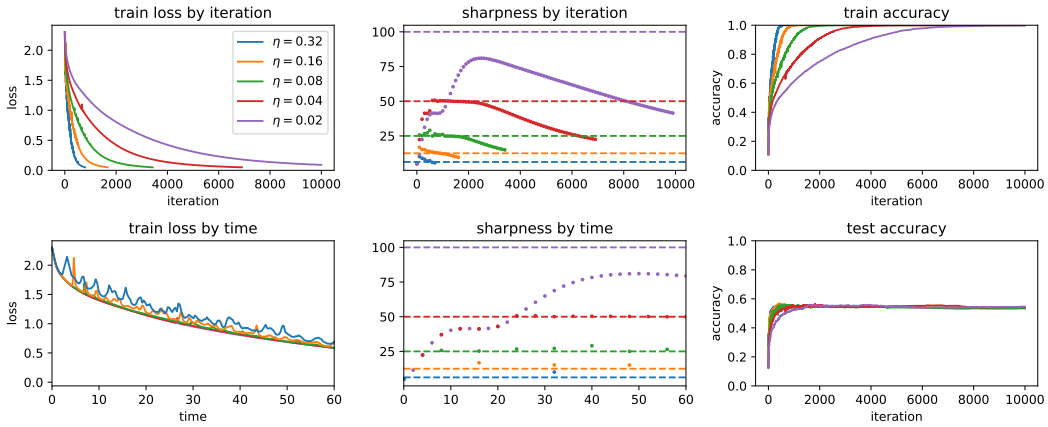

Figure 18: **Tanh CNN trained with cross-entropy on 5k subset of CIFAR-10.**

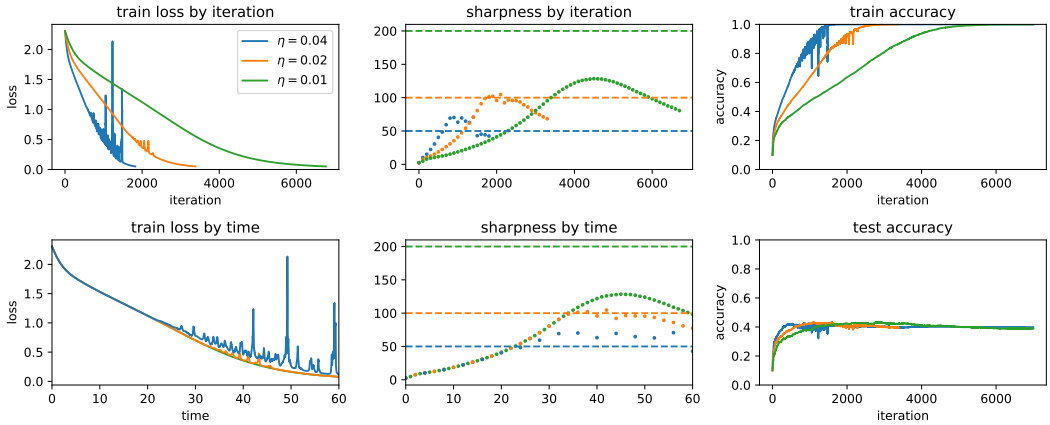

Figure 19: **ReLU fully-connected network trained with cross-entropy on 5k subset of CIFAR-10.**

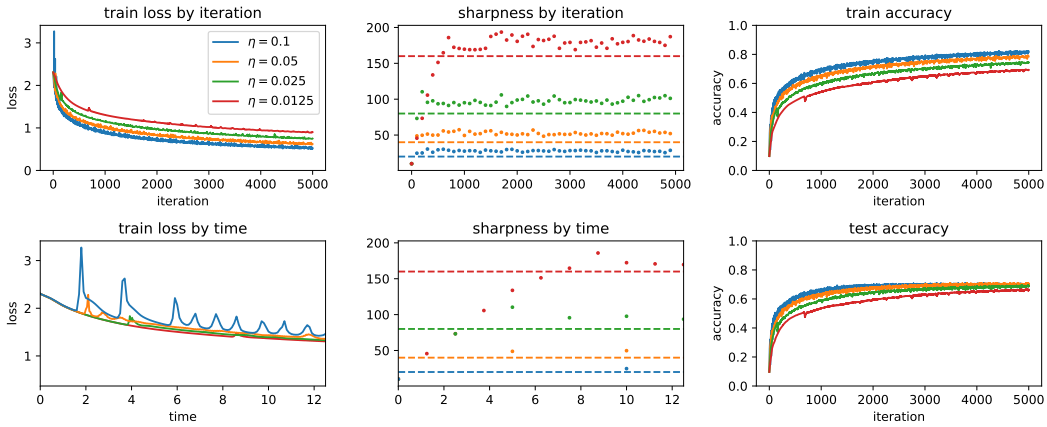

Figure 20: **ReLU CNN trained with cross-entropy on full CIFAR-10.**

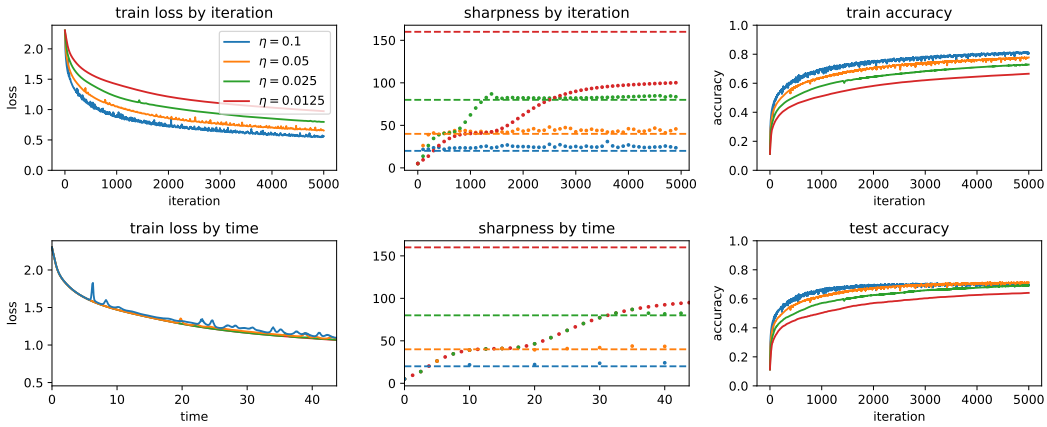

Figure 21: **Tanh CNN trained with cross-entropy on full CIFAR-10.**

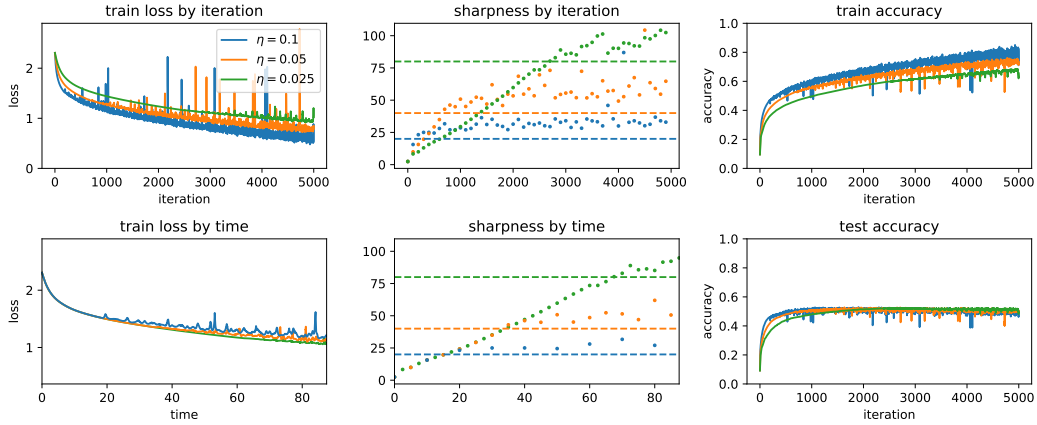

Figure 22: **ReLU fully-connected network trained with cross-entropy on full CIFAR-10.**

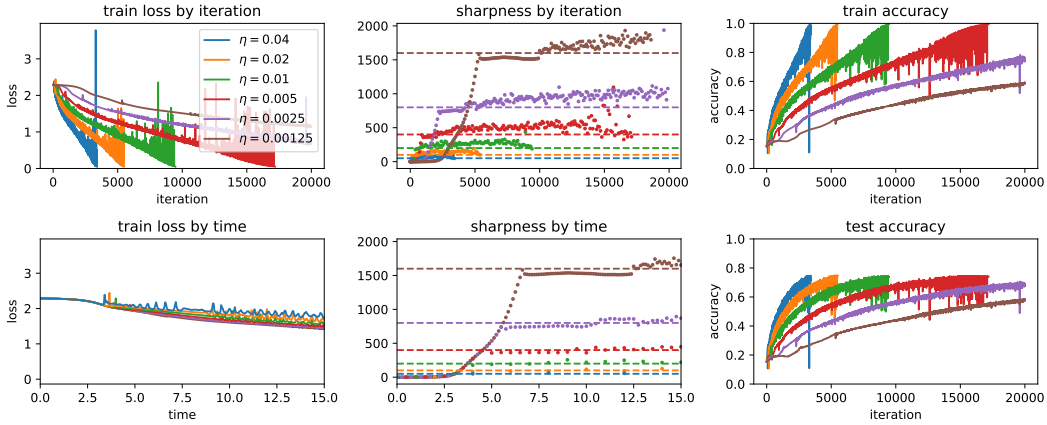

Figure 23: **VGG-11 (no BN) trained with cross-entropy on full CIFAR-10.**

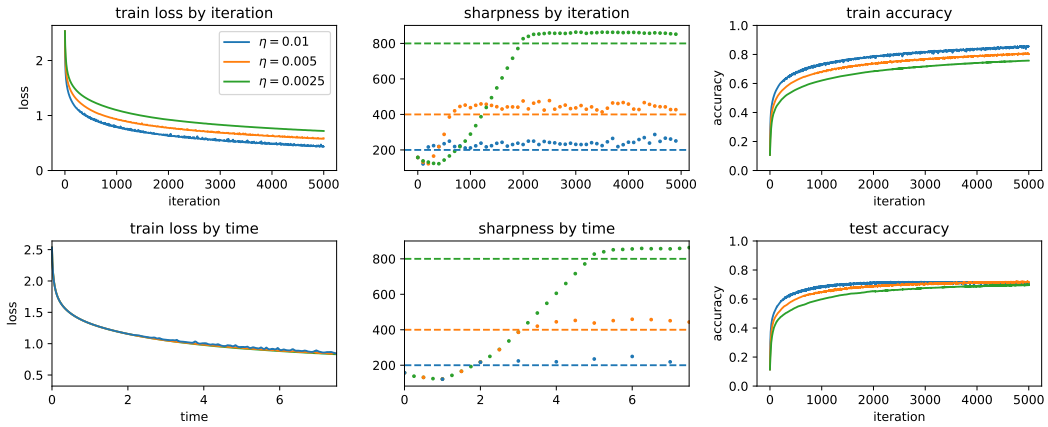

Figure 24: **ReLU CNN (with BN) trained with cross-entropy on full CIFAR-10.** Note that at the beginning of training, the sharpness decreases for a short while. This is why we phrase progressive sharpening as an "overwhelming tendency" for the sharpness to increase during gradient flow, not a hard-and-fast rule.

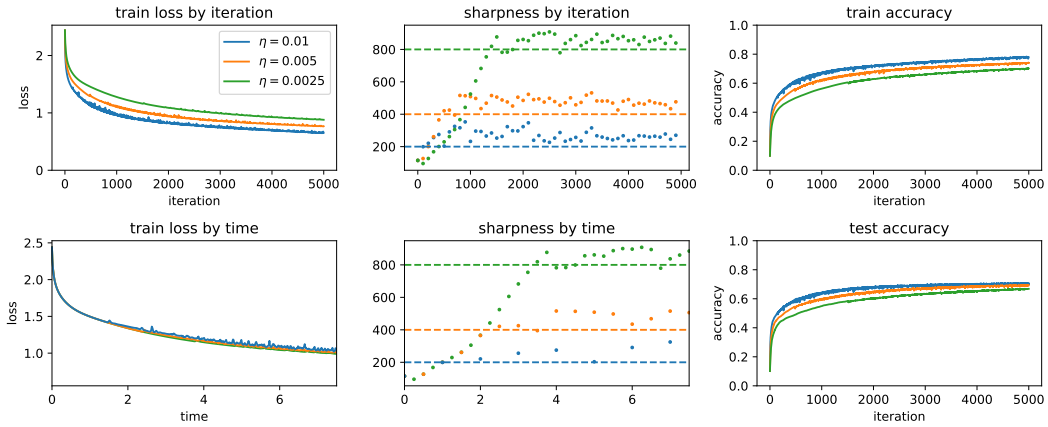

Figure 25: **Tanh CNN (with BN) trained with cross-entropy on full CIFAR-10.**

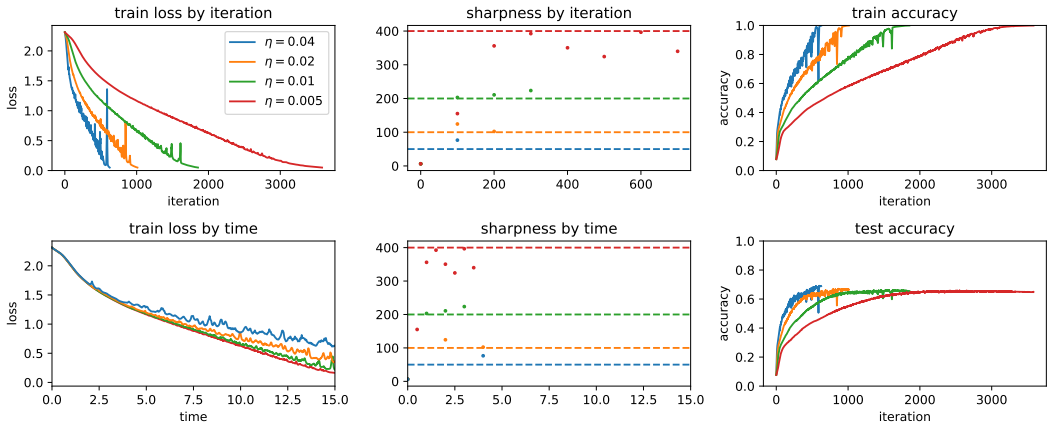

Figure 26: **VGG-11 (with BN) trained with cross-entropy on full CIFAR-10.** The red learning rate enters the Edge of Stability when the sharpness directly on the gradient descent trajectory is slightly less $2/\eta$; we have verified that during this period, the sharpness in *between* iterates exceeds $2/\eta$.

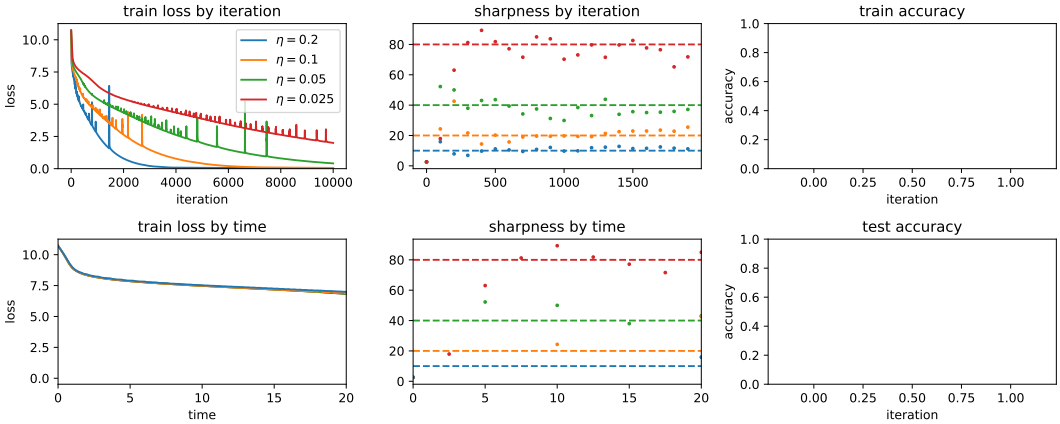

Figure 27: **Transformer trained on a subset of the WikiText-2 language modeling dataset.**

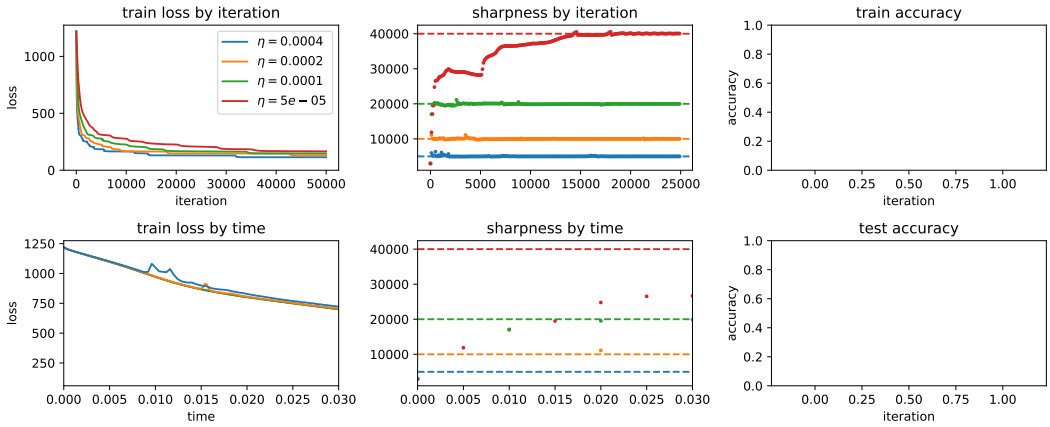

Figure 28: **Deep linear network trained on isotropic Gaussian data.**

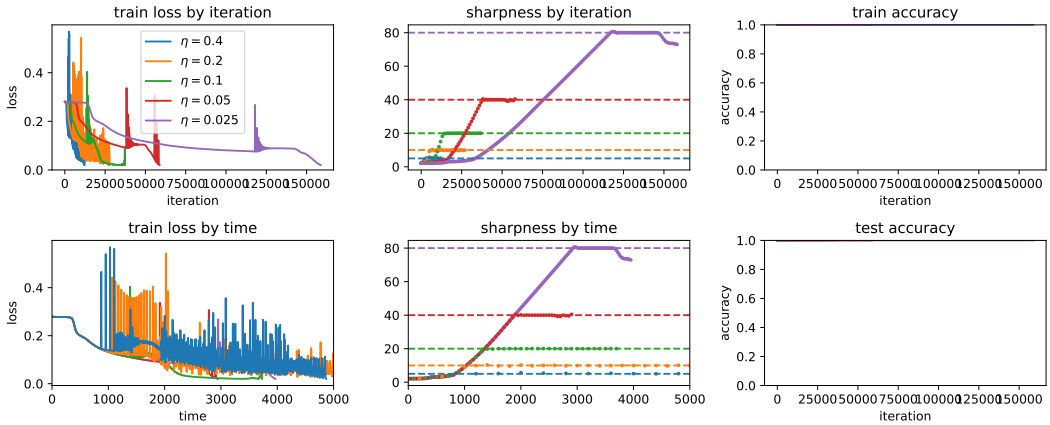

Figure 29: **Tanh fully-connected network with one hidden layer trained on toy 1-dimensional regression task.**

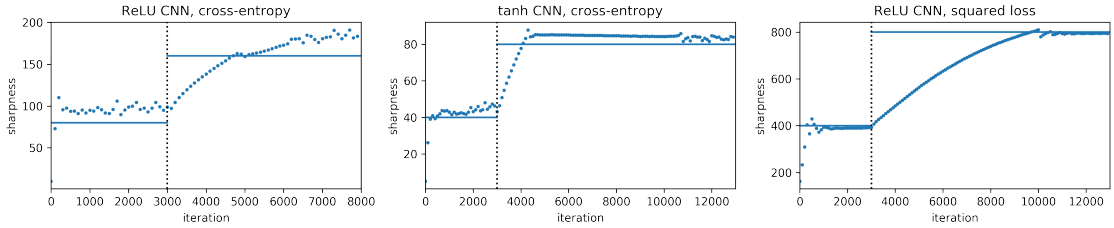

Figure 30: **Additional learning rate drop experiments.** If we cut the learning rate while the network is at the Edge of Stability, then progressive sharpening resumes until the network is once again at the Edge of Stability. (These networks are trained on the full CIFAR-10 dataset.)