

# Adaptive Gradient Methods at the Edge of Stability

Jeremy Cohen\* Behrooz Ghorbani Shankar Krishnan Naman Agarwal

Sourabh Medapati Michal Badura Daniel Suo David Cardoze

Zachary Nado George E. Dahl Justin Gilmer

Google Research, Brain Team

## Abstract

Very little is known about the training dynamics of adaptive gradient methods like Adam in deep learning. In this paper, we shed light on the behavior of these algorithms in the full-batch and sufficiently large batch settings. Specifically, we empirically demonstrate that during full-batch training, the maximum eigenvalue of the *preconditioned* Hessian typically equilibrates at a certain numerical value — the stability threshold of a gradient descent algorithm. For Adam with step size  $\eta$  and  $\beta_1 = 0.9$ , this stability threshold is  $38/\eta$ . Similar effects occur during minibatch training, especially as the batch size grows. Yet, even though adaptive methods train at the “Adaptive Edge of Stability” (AEoS), their behavior in this regime differs in a significant way from that of non-adaptive methods at the EoS. Whereas non-adaptive algorithms at the EoS are blocked from entering high-curvature regions of the loss landscape, adaptive gradient methods at the AEoS keep advancing into high-curvature regions, while adapting the preconditioner to compensate. Our findings can serve as a foundation for the community’s future understanding of adaptive gradient methods in deep learning.

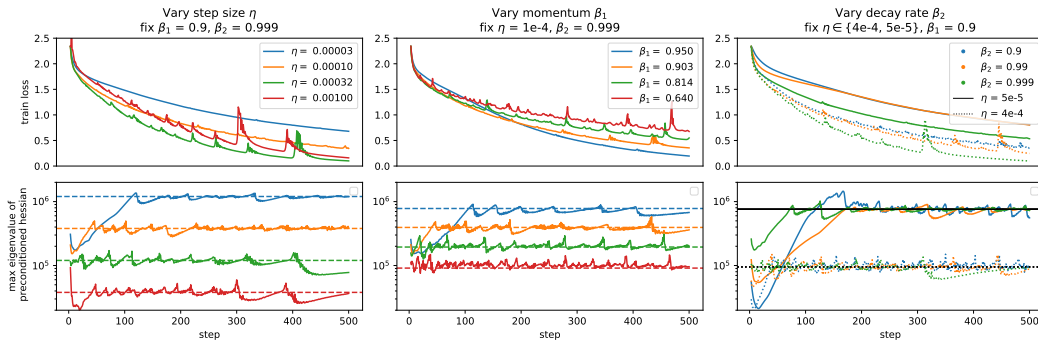


Figure 1: **Full-batch Adam trains at the Adaptive Edge of Stability (AEoS).** We train a fully-connected network on CIFAR-10 using full-batch Adam with various hyperparameters  $\eta, \beta_1, \beta_2$ . Observe that the maximum eigenvalue of the preconditioned Hessian equilibrates at the numerical value  $\frac{(2+2\beta_1)}{(1-\beta_1)}\eta$ , which is drawn as a dashed horizontal line. However, in contrast to the non-adaptive EoS, the maximum eigenvalue of the raw Hessian usually keeps rising at the AEoS (see Figure 6).

\*Work done as student researcher. Correspondence to {jeremycohen@cmu.edu, gilmer@google.com}.

# 1 Introduction

Neural networks are often trained using adaptive gradient methods such as Adam [23], rmsprop [38], or Adafactor [37]. These algorithms are variants of preconditioned gradient descent where the preconditioner is constantly *adapting* to gradients from recent steps. Despite the popularity of these algorithms, their behavior when training neural networks remains poorly understood.

Meanwhile, a line of recent work [33, 42, 20, 21, 7, 15] has shed light on the neural network training dynamics of *non-adaptive* gradient descent (both with and without momentum) in, at least, the full-batch and (sufficiently) large-batch regimes. These works have empirically demonstrated that *dynamical instability* plays a key role in the training process. When training neural networks, gradient descent is constantly attracted to regions of parameter space with increasingly high curvature [20, 21, 7]; yet, up to a quadratic Taylor approximation, gradient descent is an unstable dynamical system in regions where the curvature in any direction exceeds a certain threshold — the optimizer cannot linger for long in any such region without being expelled [33, 42, 21]. How is this tension resolved? In the full-batch special case, gradient descent spends the bulk of training in a regime called the Edge of Stability (EoS) [7] in which the *sharpness* — the maximum eigenvalue of the training Hessian — hovers right at, or just above, the stability threshold. At the EoS, gradient descent would still be moving into regions of higher curvature were it not being constantly repelled from these high-curvature regions by unstable dynamics. As we confirm below, these findings also generalize to preconditioned gradient descent (with a static preconditioner).

However, it has not been clear whether these findings have relevance for *adaptive* gradient methods. Because of adaptive preconditioning, adaptive gradient methods do not evolve as linear recurrences on the local quadratic Taylor approximation, and thus it is not clear why their local stability would be well-modeled by an eigenvalue condition.

In this paper, we demonstrate that the EoS phenomenon does, in fact, carry over to the adaptive setting. Our key empirical finding is that throughout training, the short-term stability behavior of, say, Adam is well-approximated by that of “frozen Adam” — a version of Adam in which the preconditioner is frozen at its current value. On the local quadratic Taylor approximation, “frozen Adam” *does* evolve as a linear recurrence, and is unstable whenever the maximum eigenvalue of the preconditioned Hessian (the *preconditioned sharpness*) exceeds the stability threshold of EMA-style heavy ball momentum, which is  $\frac{2+2\beta_1}{\eta(1-\beta_1)}$ . Indeed, we observe that during full-batch training by real Adam, the preconditioned sharpness equilibrates at this precise numerical value (Figure 1). During minibatch training, similar phenomena occur, paralleling the situation with non-adaptive optimizers.

However, even though adaptive gradient methods train at the “Adaptive Edge of Stability” (AEoS), their behavior in this regime differs in a significant way from that of non-adaptive methods in the non-adaptive EoS regime: whereas non-adaptive optimizers in the non-adaptive EoS regime are blocked from accessing high-curvature regions of the loss landscape, we find that adaptive gradient methods at the AEoS can and do enter these high-curvature regions via their ability to adapt the preconditioner. This is especially liable to occur if the step size is small or the preconditioner decay factor (e.g. Adam’s  $\beta_2$ ) is small. Thus, adaptive gradient methods sometimes lack the implicit inductive bias [34] that blocks non-adaptive methods from converging to high-curvature solutions.

## 2 Notation and Background

Optimization algorithms aim to minimize an objective function  $f : \mathbb{R}^p \rightarrow \mathbb{R}$  by producing a sequence of iterates  $\{\mathbf{x}_t\} \subset \mathbb{R}^p$ . We denote an algorithm “alg” with hyperparameters  $a, b$  as  $\text{ALG}(a, b)$ . We write  $\mathbf{H}(\mathbf{x})$  for the Hessian at parameter  $\mathbf{x}$ , and sometimes write  $\mathbf{H}_t$  as an abbreviation for  $\mathbf{H}(\mathbf{x}_t)$ , the Hessian at step  $t$ . We denote the largest positive eigenvalue of a matrix  $\mathbf{H}$  as  $\lambda_1(\mathbf{H})$ . We call the largest positive eigenvalue of a Hessian the *sharpness*, and the largest positive eigenvalue of a preconditioned Hessian the *preconditioned sharpness*.

**Gradient descent and momentum** The most basic optimization algorithm is  $\text{GD}(\eta)$ , defined by the parameter update  $\mathbf{x}_{t+1} = \mathbf{x}_t - \eta \mathbf{g}_{t+1}$ , where  $\mathbf{g}_{t+1}$  is a full-batch or minibatch gradient of  $f$  computed at  $\mathbf{x}_t$ . Vanilla gradient descent can be accelerated by the use of momentum, in which case the parameter update becomes  $\mathbf{x}_{t+1} = \mathbf{x}_t - \eta \mathbf{m}_{t+1}$ , where  $\mathbf{m}_{t+1}$  is a momentum vector. Momentum comes in two popular flavors: heavy ball (HB) and Nesterov. Both of these, in turn, can

be parameterized in the “standard” fashion or the “EMA” (exponential moving average) fashion. The momentum update for STANDARDHB( $\eta, \beta_1$ ) is  $\mathbf{m}_{t+1} = \beta_1 \mathbf{m}_t + \mathbf{g}_{t+1}$ , whereas the update for EMAHB( $\eta, \beta_1$ ) is  $\mathbf{m}_{t+1} = \beta_1 \mathbf{m}_t + (1 - \beta_1) \mathbf{g}_{t+1}$ .

Each of these algorithms, when run on any quadratic function (and in particular, when run on the local quadratic Taylor approximation to a neural training objective), evolves independently along each Hessian eigenvector as a linear recurrence relation (Appendix E). If any Hessian eigenvalue exceeds a certain algorithm-dependent threshold, then the linear recurrence relation for the corresponding eigenvector is *unstable* and diverges exponentially; this implies that the iterate will oscillate with exponentially increasing magnitude along this eigenvector. The stability thresholds are:  $2/\eta$  for GD( $\eta$ ),  $\frac{2+2\beta_1}{\eta}$  for STANDARDHB( $\eta, \beta_1$ ), and  $\frac{2+2\beta_1}{\eta(1-\beta_1)}$  for EMAHB( $\eta, \beta_1$ ).

We illustrate this threshold behavior in Figure 2, where we optimize the one-dimensional quadratic function  $f(x) = \frac{1}{2}x^2$  using EMA-style heavy ball momentum with  $\beta_1 = 0.9$  and varying learning rates. Observe that the iterates diverge whenever the learning rate  $\eta$  exceeds  $\frac{2+2\beta_1}{(1-\beta_1)} = 38$ .

**Preconditioned gradient descent** Gradient descent, or its momentum variants, can be preconditioned using a (static) preconditioner  $\mathbf{P}$ . In this case, the update rule becomes  $\mathbf{x}_{t+1} = \mathbf{x}_t - \eta \mathbf{P}^{-1} \mathbf{m}_{t+1}$ , where  $\mathbf{m}_t$  is updated according to one of the momentum rules described above. Often,  $\mathbf{P}$  is a diagonal matrix, in which case  $\mathbf{P}^{-1}$  can be interpreted as per-parameter learning rates. A special case of particular interest in this paper is preconditioned gradient descent with EMA-style heavy ball momentum, which we denote as PRECONDEMAHB( $\eta, \beta_1, \mathbf{P}$ ). On a quadratic objective with Hessian  $\mathbf{H}$ , preconditioned gradient descent will diverge if the maximum eigenvalue of the *preconditioned Hessian*  $\mathbf{P}^{-1/2} \mathbf{H} \mathbf{P}^{-1/2}$  (the *preconditioned sharpness*) exceeds the stability threshold of the non-preconditioned algorithm. In that event, the component of the iterate that is aligned with  $(\mathbf{P}^{1/2})^T \mathbf{v}_1$ , where  $\mathbf{v}_1$  is the top eigenvector of  $\mathbf{P}^{-1/2} \mathbf{H} \mathbf{P}^{-1/2}$ , will oscillate with exponentially increasing magnitude, and diverge. Note that  $\mathbf{P}^{-1/2} \mathbf{H} \mathbf{P}^{-1/2}$  shares eigenvalues with the similar matrix  $\mathbf{P}^{-1} \mathbf{H}$ , so the preconditioned sharpness can be written as both  $\lambda_1(\mathbf{P}^{-1/2} \mathbf{H} \mathbf{P}^{-1/2})$  or  $\lambda_1(\mathbf{P}^{-1} \mathbf{H})$ .

**Adaptive gradient methods** Adaptive gradient methods are variants of preconditioned gradient descent in which the preconditioner is changing rather than fixed, i.e. the parameter update is  $\mathbf{x}_{t+1} = \mathbf{x}_t - \eta \mathbf{P}_{t+1}^{-1} \mathbf{m}_{t+1}$ , where  $\mathbf{P}_{t+1}$  is the latest preconditioner. For example, RMSPROP( $\eta, \beta_2, \epsilon$ ) [38] updates its preconditioner according to:

$$\boldsymbol{\nu}_{t+1} = \beta_2 \boldsymbol{\nu}_t + (1 - \beta_2) \mathbf{g}_{t+1}^{\circ 2} \quad \mathbf{P}_{t+1} = \text{diag}(\boldsymbol{\nu}_{t+1}^{1/2}) + \epsilon \mathbf{I}. \quad (1)$$

A more popular algorithm is ADAM( $\eta, \beta_1, \beta_2, \epsilon$ ) [23], which adds EMA-style heavy ball momentum to rmsprop. Adam has an optional “bias correction” scheme. With no bias correction, Adam employs the rmsprop preconditioner rule in Eq (1); with bias correction, the Adam preconditioner is:

$$\mathbf{P}_{t+1} = \left( \frac{1}{1 - \beta_1^{t+1}} \right) \left[ \text{diag} \left( \sqrt{\frac{\boldsymbol{\nu}_{t+1}}{1 - \beta_2^{t+1}}} \right) + \epsilon \mathbf{I} \right]. \quad (2)$$

In contrast to gradient descent (and preconditioned gradient descent), adaptive gradient methods do not evolve as linear recurrence relations on quadratic functions. Thus, it is a priori unclear whether their local stability can be modeled using an eigenvalue condition.

### 3 Related Work

**Training dynamics of non-adaptive gradient descent** Recent empirical studies [21, 25, 7, 24] have shed substantial light on the training dynamics of non-adaptive gradient descent in deep learning. When training neural networks, gradient descent tends to continually move in a direction that increases the sharpness [20, 21, 7]. This phenomenon was dubbed *progressive sharpening* in [7], and remains poorly understood. In the special case of *full-batch* gradient descent, the sharpness usually rises past the optimizer’s stability threshold, at which point (the “breakeven point” [21]) the optimizer becomes destabilized by exponentially growing movement along the Hessian’s top eigenvector. On quadratic objective functions, this behavior would lead to divergence. However, neural network training objectives are not quadratic, and gradient descent typically does not diverge; instead, it

enters a regime called the Edge of Stability (EoS) [7] in which the sharpness hovers just above, or oscillates around, the stability threshold. At the EoS, gradient descent is constantly being repelled from regions of the loss landscape with sharpness exceeding the stability threshold, behavior which (author?) [10] have explained by appealing to a cubic Taylor expansion. Prior work has not discussed preconditioned gradient descent, but in the next section we confirm that the results of [7] carry over to the non-adaptive preconditioned setting. Note that these results only apply to *full-batch* gradient descent. In the more general case of SGD, similar effects seem to occur [20, 21], as we review in §B.

The fact that non-adaptive gradient descent is blocked from entering sharp (as quantified by maximum Hessian eigenvalue) regions of the loss landscape constitutes one implicit bias [34] of non-adaptive gradient descent. It is plausible that this implicit bias could impact generalization (e.g. see [32, 19]).

**Understanding adaptive gradient methods** There have been a number of convergence analyses of adaptive gradient methods [6, 47, 26, 40, 43, 11], none of which model the behavior that we observe here. Beyond formal convergence analyses, [3, 9] proposed to model Adam as an system of ordinary differential equations in continuous time; this approach also cannot explain the unstable dynamics that we observe. [2] argued that Adam should be viewed as a variant of sign gradient descent. Most relevant to our paper, [31] conducted a qualitative study of full-batch Adam during neural network training. They observed small bumps and large spikes in the training loss curve; our paper explains this phenomenon as arising from dynamical instability.

**Implicit bias of adaptive gradient methods** Absent well-tuned regularization, adaptive gradient methods have been reported to generalize worse than non-adaptive optimizers [41, 22, 28]. However, it has been far from clear why, in deep learning, adaptive optimizers should even find consistently different solutions than non-adaptive optimizers (let alone why these different solutions should generalize worse). [41] constructed a synthetic task where Adam provably generalizes worse than gradient descent. However, synthetic tasks also exists where the reverse is true [1]. [39] proved that rmsprop on homogenous neural networks is, like gradient descent [30], implicitly biased towards maximum margin solutions. [18] argued that adaptive methods take steps that are too large in small-Hessian-eigenvalue directions, which are more likely to be noise than signal.

Our paper gives evidence for a different implicit bias: adaptive gradient methods are liable to find higher-curvature solutions than non-adaptive algorithms, since whereas non-adaptive algorithms are blocked from high-curvature regions, adaptive optimizers can evade this restriction. [42] briefly speculated that this might be the case, but ran no experiments with adaptive optimizers (only L-BFGS). Meanwhile, [17] empirically observed that a VGG trained using an adaptive optimizer had 40x the sharpness than when trained using SGD, and [5] found that MLP-mixers and ViTs trained with Adam converged to far sharper solutions than a ResNet trained with SGD.

[48] and [44] previously argued theoretically that SGD can more quickly escape certain “sharp minima” than Adam, but using a different definition of sharpness, and a mechanism tied inextricably to stochasticity. In contrast, we demonstrate that even full-batch Adam converges to higher-curvature solutions than full-batch gradient descent, implying that stochasticity is not crucial.

## 4 Full-batch adaptive optimizers train at the Adaptive Edge of Stability

**Warm-up: “frozen Adam”** Before turning to adaptive preconditioning, we first confirm that prior work [7] generalizes to the setting of preconditioned gradient descent. We train a network using full-batch Adam (no bias correction), and extract the second moment accumulator  $\nu_t$  at some midpoint step  $t$  during training. We then train the same architecture from scratch at a range of learning rates using full-batch “frozen Adam,” i.e. Adam where the second-moment accumulator is held fixed at that particular  $\nu_t$  rather than updated via Eq (1) or Eq (2). Note that “frozen Adam” is equivalent to  $\text{PRECONDEMAHB}(\eta, \beta_1, \mathbf{P})$  where  $\mathbf{P} = \text{diag}(\nu_t^{1/2}) + \epsilon \mathbf{I}$ . Thus frozen Adam is dynamically unstable on the training objective’s quadratic Taylor approximation whenever the preconditioned sharpness  $\lambda_1(\mathbf{P}^{-1} \mathbf{H}_t)$  exceeds the stability threshold for EMA-style heavy ball momentum. This threshold is  $\frac{2+2\beta_1}{(1-\beta_1)\eta} = \frac{38}{\eta}$  for our choice of  $\beta_1 = 0.9$ . Assuming that the findings of [7] generalize from non-preconditioned to preconditioned gradient descent, one would expect that during full-batch training by frozen Adam, the preconditioned sharpness would rise until plateauing at that threshold value. In Figure 3(b), we confirm that this occurs. For each of several learning rates  $\eta$ , we plot the

evolution of the preconditioned sharpness  $\lambda_1(\mathbf{P}^{-1}\mathbf{H}_t)$  while visualizing the stability threshold of  $38/\eta$  as a horizontal dashed line. Observe that the preconditioned sharpness rises until plateauing at the dashed line. We also plot the evolution of the “raw” sharpness  $\lambda_1(\mathbf{H}_t)$  in Figure 3(c). For this network, the raw sharpness largely ceases to increase once training enters the EoS.

**Adaptive preconditioning: Adam** Adam is a more complex dynamical system than “frozen Adam,” because its preconditioner evolves in response to recent gradients; consequently, its dynamics on the local quadratic Taylor approximation does not reduce to a linear recurrence. Throughout physics and engineering, a time-tested way to understand complex systems is to approximate them by simpler ones that still capture some behavior of interest. Consider Adam at any particular step  $t_0$ , and let  $\mathbf{P}_{t_0}$  be Adam’s preconditioner at that step. Over short time horizons  $t > t_0$ , one can approximate Adam by  $\text{PRECONDEMAHB}(\eta, \beta_1, \mathbf{P}_{t_0})$ , i.e. “frozen Adam” where the preconditioner is frozen at  $\mathbf{P}_{t_0}$ . On the local quadratic Taylor approximation, this simpler algorithm is unstable if the preconditioned sharpness  $\lambda_1(\mathbf{P}_{t_0}^{-1}\mathbf{H}_{t_0})$  exceeds the stability threshold of heavy-ball EMA momentum, which is  $\frac{2+2\beta_1}{(1-\beta_1)\eta}$ . Of course, whether this approximation is useful is an empirical question. And empirically, we now show that Adam has similar stability properties as its frozen counterpart.

In Figure 1, we train a fully-connected network on the full CIFAR-10 dataset using full-batch Adam (no bias correction). In Figure 1(a), we vary the learning rate  $\eta$  while fixing  $\beta_2 = 0.999$  and  $\beta_1 = 0.9$ . Observe that for every learning rate  $\eta$ , the preconditioned sharpness  $\lambda_1(\mathbf{P}_t^{-1}\mathbf{H}_t)$  equilibrates right around the stability threshold of frozen Adam. In Figure 1(b), we vary  $\beta_1$  as we fix  $\eta = 0.0001$  and  $\beta_2 = 0.999$ . Observe that the same phenomenon occurs. Finally, in Figure 1(c), we vary  $\beta_2 \in \{0.9, 0.99, 0.999\}$  as we fix  $\beta_1 = 0.9$  and  $\eta \in \{5\text{e-}5, 4\text{e-}4\}$ . The same phenomenon occurs, even when  $\beta_2$  takes the relatively small value of 0.9, which is when one might expect Adam’s behavior to differ the most from frozen Adam’s, because the preconditioner adapts faster. Thus, we see that Adam becomes unstable whenever its frozen counterpart would become unstable. Nevertheless, in Appendix D, we will see that Adam’s behavior at the AEOs differs in other respects from that of frozen Adam at the EoS.

**Other architectures** In Figure 5, we verify that the full-batch phenomenon generalizes to other computer vision architectures. In particular, we consider (a) a CNN on CIFAR-10; (b) an unnormalized Wide ResNet (WRN) [46] on CIFAR-10; and (c) a (batch-normalized) WRN on CIFAR-100. Furthermore, Figure 10 verifies that our *minibatch* findings (see below) apply to transformers on WMT machine translation.

**Other full-batch adaptive gradient methods** In Figure 4, we train the same fully-connected network on CIFAR-10 using eight adaptive gradient algorithms in the full-batch regime: (a) Adam with bias correction, which essentially decreases the preconditioner on a schedule in the early stage of training; (b) AdamW [29], which employs decoupled weight decay; (c) Adafactor [37], which maintains a factored layer-wise preconditioner; (d) Amsgrad [36], which ensures that the preconditioner is entrywise non-decreasing over time; (e) Padam [4], which employs an exponent smaller than  $\frac{1}{2}$  (we use 0.25) in the definition of the preconditioner Eq (1); (f) Nadam [12], which uses Nesterov rather than heavy ball momentum; (g) rmsprop [38], which lacks momentum; and finally (h) Adagrad [13], which preconditions using the sum of past squared gradients, rather than an exponential moving average. In all cases, we observe that the preconditioned sharpness  $\lambda_1(\mathbf{P}_t^{-1}\mathbf{H}_t)$  equilibrates around the appropriate stability threshold (written in parentheses).

**A corner case** When running full-batch Adam at extremely small learning rates, we sometimes observe that the preconditioned sharpness equilibrates *short of* the stability threshold. We elaborate on this corner case, and offer an explanation, in Appendix H.

**The minibatch setting** So far we have discussed the idealized yet unrealistic setting of full-batch training. In Appendix B, we demonstrate that analogous effects in the more practical minibatch setting, paralleling the situation with non-adaptive optimizers. We consider both ResNet-50 on ImageNet as well as Transformers on WMT.

**Adam vs. momentum** In Appendix C, show that as a consequence of the phenomena described in this paper, Adam converges to sharper solutions than momentum SGD.

## 5 Conclusion

In a recent year, Adam [23] was the most highly cited scientific paper across all fields of science [8]. Yet, very little is known concretely about the training dynamics of Adam or other adaptive gradient methods in deep learning. Since the training dynamics affect both the optimization and generalization performance of neural networks in mysterious ways, any basic knowledge about the training dynamics could pay dividends for future research. This paper has demonstrated that in the special setting of full-batch training, there is an equilibrium rule that is typically maintained throughout training. We have provided evidence that a similar equilibrium may exist in the more practical minibatch setting, and we hope that an analogous equilibrium rule will eventually be found.

## References

- [1] Naman Agarwal, Rohan Anil, Elad Hazan, Tomer Koren, and Cyril Zhang. Disentangling adaptive gradient methods from learning rates. *arXiv preprint arXiv:2002.11803*, 2020.
- [2] Lukas Balles and Philipp Hennig. Dissecting adam: The sign, magnitude and variance of stochastic gradients. In *International Conference on Machine Learning*, pages 404–413. PMLR, 2018.
- [3] Anas Barakat and Pascal Bianchi. Convergence and dynamical behavior of the adam algorithm for non-convex stochastic optimization. *arXiv preprint arXiv:1810.02263*, 2018.
- [4] Jinghui Chen, Dongruo Zhou, Yiqi Tang, Ziyang Yang, Yuan Cao, and Quanquan Gu. Closing the generalization gap of adaptive gradient methods in training deep neural networks. *arXiv preprint arXiv:1806.06763*, 2018.
- [5] Xiangning Chen, Cho-Jui Hsieh, and Boqing Gong. When vision transformers outperform resnets without pre-training or strong data augmentations. *arXiv preprint arXiv:2106.01548*, 2021.
- [6] 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.
- [7] Jeremy M. Cohen, Simran Kaur, Yuanzhi Li, J. Zico Kolter, and Ameet S. Talwalkar. Gradient descent on neural networks typically occurs at the edge of stability. *ArXiv*, abs/2103.00065, 2021.
- [8] Bec Crew. Google scholar reveals its most influential papers for 2020, Jul 2020.
- [9] André Belotto da Silva and Maxime Gazeau. A general system of differential equations to model first-order adaptive algorithms. *J. Mach. Learn. Res.*, 21:129–1, 2020.
- [10] Alex Damian, Eshaan Nichani, and Jason D. Lee. Self-stabilization: The implicit bias of gradient descent at the edge of stability, 2023.
- [11] Alexandre Défossez, Léon Bottou, Francis Bach, and Nicolas Usunier. A simple convergence proof of adam and adagrad. *arXiv preprint arXiv:2003.02395*, 2020.
- [12] Timothy Dozat. Incorporating nesterov momentum into adam. *ICLR 2016, workshop track*, 2016.
- [13] John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12(61):2121–2159, 2011.
- [14] 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? *arXiv preprint arXiv:1909.12340*, 2019.
- [15] Justin Gilmer, Behrooz Ghorbani, Ankush Garg, Sneha Kudugunta, Behnam Neyshabur, David Cardoze, George Dahl, Zachary Nado, and Orhan Firat. A loss curvature perspective on training instability in deep learning. *arXiv preprint arXiv:2110.04369*, 2021.

- [16] Gabriel Goh. Why momentum really works. *Distill*, 2017.
- [17] Diego Granzio. Flatness is a false friend. *arXiv preprint arXiv:2006.09091*, 2020.
- [18] Diego Granzio, Xingchen Wan, Samuel Albanie, and Stephen Roberts. Explaining the adaptive generalisation gap. *arXiv preprint arXiv:2011.08181*, 2020.
- [19] Stanislaw Jastrzebski, Devansh Arpit, Oliver Astrand, Giancarlo B Kerg, Huan Wang, Caiming Xiong, Richard Socher, Kyunghyun Cho, and Krzysztof J Geras. Catastrophic fisher explosion: Early phase fisher matrix impacts generalization. In *International Conference on Machine Learning*, pages 4772–4784. PMLR, 2021.
- [20] Stanislaw Jastrzebski, Zachary Kenton, Nicolas Ballas, Asja Fischer, Yoshua Bengio, and Amos Storkey. On the relation between the sharpest directions of dnn loss and the sgd step length. *arXiv preprint arXiv:1807.05031*, 2018.
- [21] 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. *arXiv preprint arXiv:2002.09572*, 2020.
- [22] Nitish Shirish Keskar and Richard Socher. Improving generalization performance by switching from adam to sgd. *arXiv preprint arXiv:1712.07628*, 2017.
- [23] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [24] Sungyoon Lee, Jinseong Park, and Jaewook Lee. Implicit jacobian regularization weighted with impurity of probability output, 2022.
- [25] 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.
- [26] 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*, pages 983–992. PMLR, 2019.
- [27] Liyuan Liu, Xiaodong Liu, Jianfeng Gao, Weizhu Chen, and Jiawei Han. Understanding the difficulty of training transformers. *arXiv preprint arXiv:2004.08249*, 2020.
- [28] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016.
- [29] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- [30] Kaifeng Lyu and Jian Li. Gradient descent maximizes the margin of homogeneous neural networks. *arXiv preprint arXiv:1906.05890*, 2019.
- [31] Chao Ma, Lei Wu, et al. A qualitative study of the dynamic behavior for adaptive gradient algorithms. *arXiv preprint arXiv:2009.06125*, 2020.
- [32] Rotem Mulayoff, Tomer Michaeli, and Daniel Soudry. The implicit bias of minima stability: A view from function space. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021.
- [33] Kamil Nar and Shankar Sastry. Step size matters in deep learning. *Advances in Neural Information Processing Systems*, 31, 2018.
- [34] Behnam Neyshabur, Ryota Tomioka, and Nathan Srebro. In search of the real inductive bias: On the role of implicit regularization in deep learning. *arXiv preprint arXiv:1412.6614*, 2014.
- [35] Martin Popel and Ondřej Bojar. Training tips for the transformer model. *arXiv preprint arXiv:1804.00247*, 2018.

- [36] Sashank J. Reddi, Satyen Kale, and Sanjiv Kumar. On the convergence of adam and beyond. In *International Conference on Learning Representations*, 2018.
- [37] Noam Shazeer and Mitchell Stern. Adafactor: Adaptive learning rates with sublinear memory cost. In *International Conference on Machine Learning*, pages 4596–4604. PMLR, 2018.
- [38] Tijmen Tieleman and Geoffrey Hinton. Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude. *Coursera: Neural Networks for Machine Learning*, 2012.
- [39] Bohan Wang, Qi Meng, Wei Chen, and Tie-Yan Liu. The implicit bias for adaptive optimization algorithms on homogeneous neural networks. In *International Conference on Machine Learning*, pages 10849–10858. PMLR, 2021.
- [40] Rachel Ward, Xiaoxia Wu, and Leon Bottou. Adagrad stepsizes: Sharp convergence over nonconvex landscapes. In *International Conference on Machine Learning*, pages 6677–6686. PMLR, 2019.
- [41] Ashia C Wilson, Rebecca Roelofs, Mitchell Stern, Nati Srebro, and Benjamin Recht. The marginal value of adaptive gradient methods in machine learning. *Advances in neural information processing systems*, 30, 2017.
- [42] Lei Wu, Chao Ma, and Weinan E. How sgd selects the global minima in over-parameterized learning: A dynamical stability perspective. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.
- [43] Yuege Xie, Xiaoxia Wu, and Rachel Ward. Linear convergence of adaptive stochastic gradient descent. In *International Conference on Artificial Intelligence and Statistics*, pages 1475–1485. PMLR, 2020.
- [44] Zeke Xie, Xinrui Wang, Huishuai Zhang, Issei Sato, and Masashi Sugiyama. Adai: Separating the effects of adaptive learning rate and momentum inertia. *arXiv preprint arXiv:2006.15815*, 2020.
- [45] Ruibin Xiong, Yunchang Yang, Di He, Kai Zheng, Shuxin Zheng, Chen Xing, Huishuai Zhang, Yanyan Lan, Liwei Wang, and Tieyan Liu. On layer normalization in the transformer architecture. In *International Conference on Machine Learning*, pages 10524–10533. PMLR, 2020.
- [46] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016.
- [47] Dongruo Zhou, Jinghui Chen, Yuan Cao, Yiqi Tang, Ziyan Yang, and Quanquan Gu. On the convergence of adaptive gradient methods for nonconvex optimization. *arXiv preprint arXiv:1808.05671*, 2018.
- [48] Pan Zhou, Jiashi Feng, Chao Ma, Caiming Xiong, Steven Chu Hong Hoi, et al. Towards theoretically understanding why sgd generalizes better than adam in deep learning. *Advances in Neural Information Processing Systems*, 33:21285–21296, 2020.

## A Figures

This appendix contains missing figures from the main paper.

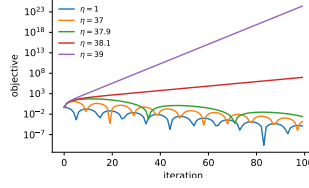


Figure 2: We optimize the quadratic objective  $f(x) = \frac{1}{2}x^2$  using  $\text{EMAHB}(\eta, 0.9)$  at various  $\eta$ . Observe that learning rates  $\eta$  above 38 diverge.

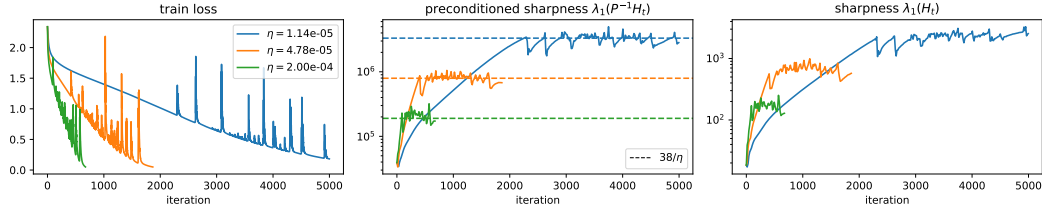


Figure 3: **“Frozen Adam” (preconditioned momentum) trains at the Edge of Stability.** We train a fully-connected network on CIFAR-10 using “frozen Adam,” i.e. preconditioned gradient descent with EMA-style heavy ball momentum and a fixed preconditioner  $\mathbf{P}$ . Consistent with [7], the preconditioned sharpness  $\lambda_1(\mathbf{P}^{-1} \mathbf{H}_t)$  rises until equilibrating at the stability threshold of  $38/\eta$ . The raw sharpness  $\lambda_1(\mathbf{H}_t)$  mostly ceases to increase at the EoS.

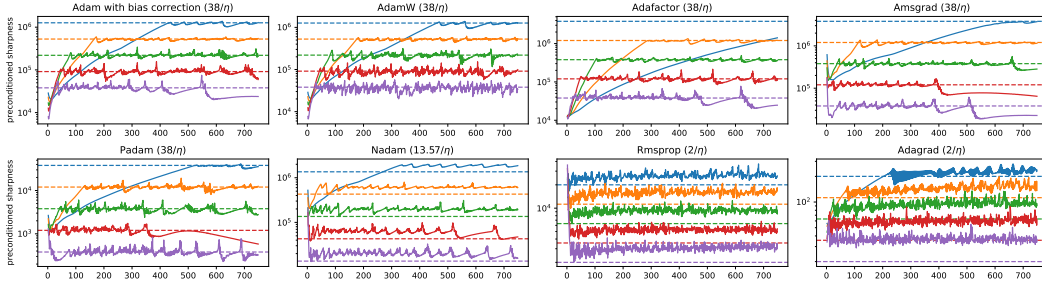


Figure 4: **Other adaptive gradient algorithms train at the AEoS.** We train a FC network on CIFAR-10 using eight adaptive optimizers in full-batch mode. We train each algorithm at five learning rates (colors). In each case, the preconditioned sharpness equilibrates at, or just above, the stability threshold (written in parentheses). Note that the qualitative behavior of the preconditioned sharpness depends on the presence and type of momentum; see Appendix G

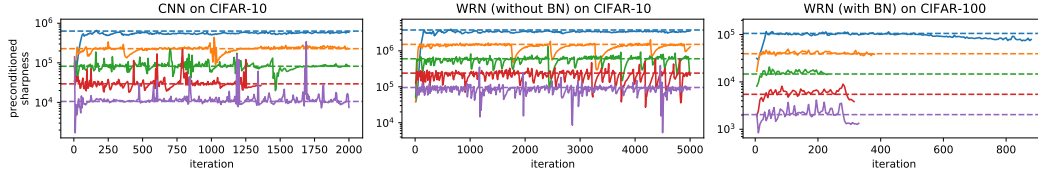


Figure 5: **The phenomenon generalizes to other architectures.** We train three vision architectures using full-batch Adam with  $\beta_1 = 0.9$  and  $\beta_2 = 0.999$  at a range of learning rates (colors). In each case, the preconditioned sharpness equilibrates at the threshold  $38/\eta$ . Each network was trained until either reaching a milestone train loss value, or until reaching a step limit.

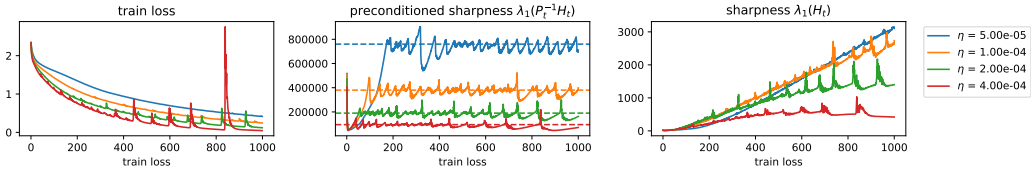


Figure 6: **The sharpness continues to rise even as Adam trains at the AEoS.** We train a fully-connected network on CIFAR-10 using full-batch Adam with  $\beta_1 = 0.9$ ,  $\beta_2 = 0.99$ , and four learning rates. Observe that while the *preconditioned* sharpness flatlines, the *raw* sharpness continues to rise.

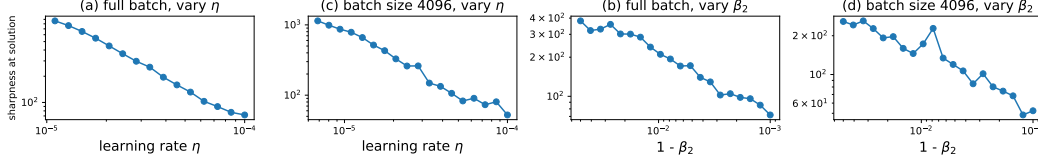


Figure 7: **Adam finds sharper solutions when  $\eta$  or  $\beta_2$  is small.** We train a WRN on CIFAR-100 until reaching train loss 0.1, and we record the sharpness  $\lambda_1(\mathbf{H}(\mathbf{x}^*))$  at the solution  $\mathbf{x}^*$ . In (a, b) we sweep the learning rate  $\eta$ , while in (c, d) we sweep  $\beta_2$ . In (a, c) we train in full-batch mode, while in (b, d) we use the large batch size 4096.

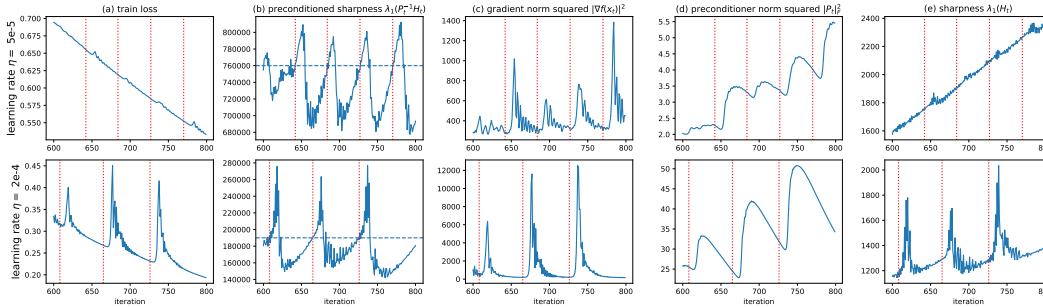


Figure 8: **Zoom in on 200 steps of training.** For two learning rates from Figure 6 (top row =  $5e-5$ , bottom row =  $2e-4$ ), we zoom in on 200 steps of training, plotting five important quantities. The red vertical lines mark the moments at which the preconditioned sharpness crosses the stability threshold.

## B The minibatch setting

We now move beyond the full-batch setting to the more general setting of minibatch training. In the case of gradient descent and minibatch SGD, it is clear from prior work [20, 21] that during minibatch training, the sharpness is subject to similar effects as during full-batch training. For one, provided that training is successful, the sharpness never ventures more than a bit above the stability threshold of the corresponding full-batch algorithm [21]. This can be explained by the fact that SGD is unstable in expectation whenever GD is unstable [14]. Yet additional factors seem to also be at play in the case of SGD. Specifically, it has been observed that at small batch sizes (where there is more gradient noise), the mid-training sharpness is smaller. One hypothesized explanation [42, 21] is that in the presence of gradient noise, SGD becomes unstable at lower sharpnesses.

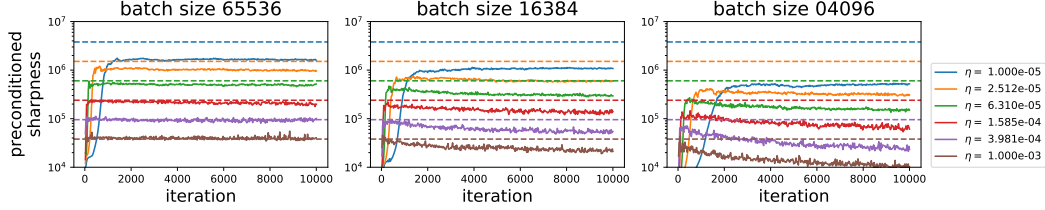


Figure 9: **During minibatch Adam, the preconditioned sharpness behaves analogous to the sharpness during minibatch SGD.** We train a Resnet-50 on ImageNet using minibatch Adam. The preconditioned sharpness is (1) below the full-batch stability threshold (pictured as a horizontal line), (2) smaller when the batch size is smaller, and (3) smaller when the learning rate is larger.

We now demonstrate that the behavior of the preconditioned sharpness during minibatch Adam parallels that of the sharpness during minibatch SGD. Namely, we observe that during minibatch Adam, the preconditioned sharpness (1) never rises more than a bit beyond the stability threshold of the full-batch algorithm, provided that training succeeds; (2) tends to be smaller, mid-training, when the learning rate is large; and (3) tends to be smaller, mid-training, when the batch size is small.

In Figure 9, we train a Resnet-50 on ImageNet at three batch sizes. Observe that for the largest batch size considered, the preconditioned sharpness behaves similar to the full-batch setting (almost rising to the full-batch stability threshold); yet, as the batch sizes decreases, the trajectory of the preconditioned sharpness shifts downwards.

In Figure 10, we train both a pre-LayerNorm (Pre-LN) and a post-LayerNorm (Post-LN) Transformer on the En->De WMT translation task. Following standard practice [35], we train using a linear warmup schedule lasting 40k steps. We train at both a smaller batch size of 1024 and a larger batch size of 10240. The dotted line depicts the stability threshold of  $38/\eta$ ; this threshold is constantly decreasing as the learning rate is warmed up. Observe that the preconditioned sharpness initially rises until reaching either the stability threshold (larger batch size), or a value just short of the stability threshold (smaller batch size). Then, as learning rate warmup continues, the preconditioned sharpness decreases to track the stability threshold. In effect, the preconditioned sharpness is being “pushed down” by Adam’s learning rate warmup. Note that [15] previously showed that a warmup schedule for momentum SGD likewise “pushed down” the (unpreconditioned) sharpness in Transformer training. The Post-LN Transformer is known to be less stable than the Pre-LN variant [45, 27]; this instability can be seen in the spike in  $\lambda_1(\mathbf{P}^{-1}\mathbf{H})$  (and training loss) towards the end of the warmup period.

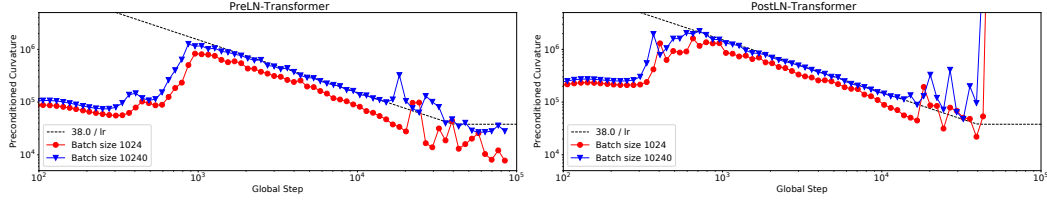


Figure 10: **Learning rate warmup gradually reduces the preconditioned sharpness during optimization.** *Left:* The evolution of the preconditioned sharpness for a 6L-6L encoder-decoder Pre-LN transformer trained on the WMT En=>De task at  $\eta = .001$  with a linear warmup period of 40000 steps. *Right:* Same as left, but with a PostLN Transformer. In both cases the preconditioned curvature closely tracks the  $38/\eta$  bound during warmup, however there is a noticeable gap at the smaller batch size. The PostLN Transformer training fails late in the warmup period.

## C Adam finds sharper solutions than momentum

Adam and other adaptive gradient methods are able to train stably in arbitrarily sharp regions of parameter space, whereas momentum gradient descent is repelled from regions of parameter space where the sharpness exceeds the algorithm’s stability threshold. We now demonstrate that as a consequence, Adam sometimes converges to far sharper solutions than momentum GD. In Figure 11, we train a WRN on CIFAR-100 (left two panes) and a fully-connected network on CIFAR-10 (right two panes) using both Adam and momentum GD. It is impossible to make a blanket claim that Adam finds sharper solutions than momentum GD, since momentum at a small learning rate finds sharper solutions than Adam at a large one; however, we observe that *for a given training speed*, Adam converges to a sharper solution.

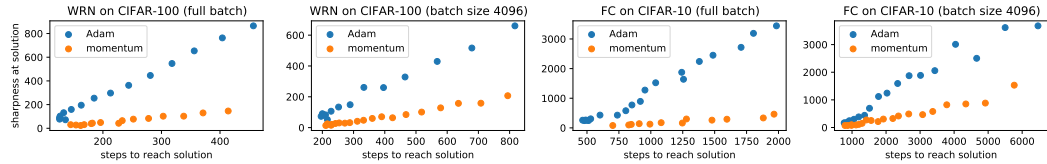


Figure 11: **At comparable training speeds, Adam finds sharper solutions than momentum gradient descent.** We train networks using both Adam (blue) and momentum (orange) at a range of learning rates. For each algorithm and learning rate, we plot the maximum Hessian eigenvalue at the solution as a function of the number of steps to reach the solution. For both algorithms, smaller learning rates (= longer training speeds) find sharper solutions; but in general, Adam finds a much sharper solution than momentum GD for the same training time.

## D A closer look at the AEoS

In this section, we take a closer look at Adam’s behavior at the AEoS. We will see that this behavior sometimes differs substantially from that of non-adaptive optimizers. In particular, whereas non-adaptive optimizers at the EoS are blocked from entering high-curvature regions of parameter space, adaptive gradient methods at the AEoS can and do enter high-curvature regions via their ability to adapt the preconditioner.

In Figure 6, we train a fully-connected network on CIFAR-10 using full-batch Adam at four learning rates. Observe that even though the preconditioned sharpness  $\lambda_1(\mathbf{P}_t^{-1}\mathbf{H}_t)$  flatlines, the raw sharpness  $\lambda_1(\mathbf{H}_t)$  continues to rise while the network is training at the AEoS. This behavior stands in stark contrast to the behavior of frozen Adam at the EoS in Figure 3; there, the raw sharpness ceased to appreciably increase after training had entered the EoS.

Why might this be the case? A crucial distinction between adaptive and non-adaptive optimizers is that the latter are necessarily repelled from high-curvature regions of parameter space. For example, EMAHB( $\eta, \beta_1$ ) is repelled from the set  $\{\mathbf{x} : \lambda_1(\mathbf{H}(\mathbf{x})) > \frac{2+2\beta_1}{\eta(1-\beta_1)}\}$ , and PRECONEMAHB( $\eta, \beta_1, \mathbf{P}$ ) is repelled from the set  $\{\mathbf{x} : \lambda_1(\mathbf{P}^{-1}\mathbf{H}(\mathbf{x})) > \frac{2+2\beta_1}{\eta(1-\beta_1)}\}$ . By contrast, while Adam’s trajectory  $\{(\mathbf{P}_t, \mathbf{x}_t)\}_t$  must approximately respect the constraint  $\lambda_1(\mathbf{P}_t^{-1}\mathbf{H}(\mathbf{x}_t)) \leq \frac{2+2\beta_1}{\eta(1-\beta_1)}$ , Adam is allowed to change the preconditioner, and hence can train stably in regions of arbitrarily high curvature.

To gain more insight into the behavior of Adam at the AEoS, in Figure 8 we “zoom in” on 200 steps of Adam training, for both a small learning rate (5e-5, the blue line in Figure 6) and a large learning rate (4e-4, the red line in Figure 6). In the top row, which corresponds to the small learning rate, we observe the following cycle. Whenever the preconditioned sharpness is below the stability threshold, both the sharpness and preconditioned sharpness rise. Eventually, the preconditioned sharpness crosses the stability threshold — an event which we mark with a vertical dotted red line. This causes Adam to oscillate explosively along a certain direction, as discussed in §2 (see Appendix I for more evidence). The large gradients associated with this explosive growth increase the second-moment accumulator  $\nu_t$  and hence Adam’s preconditioner  $\mathbf{P}_t$ . This increase in  $\mathbf{P}_t$  in turn causes the preconditioned sharpness  $\lambda_1(\mathbf{P}_t^{-1}\mathbf{H}_t)$  to decrease until it drops below the stability threshold, at which point the large gradients go away. The cycle then repeats itself. By this mechanism, the preconditioned sharpness equilibrates at the stability threshold, but the raw sharpness continues to grow. In other words, instability generates large gradients, which in turn increase the preconditioner, allowing Adam to continue advancing into regions of higher sharpness.

The bottom row of Figure 8 depicts 200 steps of Adam at a larger learning rate. At this larger learning rate, Adam behaves somewhat differently. As above, we see that whenever the preconditioned sharpness is below the threshold, both the sharpness and preconditioned sharpness rise. And, as above, we see that after the preconditioned sharpness crosses the threshold (red vertical lines), the gradient norm begins to increase explosively. However at the larger learning rate, this instability is “resolved” differently than at the small learning rate. In particular, for this large learning rate, we observe that the instability is resolved when both the sharpness and the preconditioned sharpness drop, indicating that the optimizer has escaped to a region where the curvature in many different directions (not just the maximum eigenvector of the preconditioned Hessian) is lower. This behavior is similar to the “catapult” effect studied in [25] for non-adaptive optimizers.

## E Mathematical background

### E.1 Stability of gradient descent algorithms on quadratics

We now review the stability properties of gradient descent algorithms on quadratic objective functions. We first refer the reader to Appendix A of [7]; that appendix derives the stability thresholds for GD( $\eta$ ), STANDARDHB( $\eta, \beta_1$ ), and STANDARDNESTEROV( $\eta, \beta_1$ ). In this appendix, we derive the stability thresholds for EMAHB( $\eta, \beta_1$ ), and EMANESTEROV( $\eta, \beta_1$ ). These derivations are minor variants of the corresponding derivations in [7], since the EMA-style momentum algorithms yield the same recurrence relation as the standard-style momentum algorithms, except with a learning rate scaled by  $(1 - \beta_1)$ .

Algorithm	Update rule	Stability threshold	Reference for threshold
GD( $\eta$ )	$\mathbf{x}_{t+1} = \mathbf{x}_t - \eta \mathbf{g}_{t+1}$	$2/\eta$	common knowledge, reproduced in [7] for completeness
STANDARDHB( $\eta, \beta_1$ )	$\mathbf{m}_{t+1} = \beta_1 \mathbf{m}_t + \mathbf{g}_{t+1}$ $\mathbf{x}_{t+1} = \mathbf{x}_t - \eta \mathbf{m}_{t+1}$	$\frac{2+2\beta_1}{\eta}$	[16], reproduced in [7] for completeness
EMAHB( $\eta, \beta_1$ )	$\mathbf{m}_{t+1} = \beta_1 \mathbf{m}_t + (1 - \beta_1) \mathbf{g}_{t+1}$ $\mathbf{x}_{t+1} = \mathbf{x}_t - \eta \mathbf{m}_{t+1}$	$\frac{2+2\beta_1}{\eta(1-\beta_1)}$	this paper
STANDARDNESTEROV( $\eta, \beta_1$ )	$\mathbf{m}_{t+1} = \beta_1 \mathbf{m}_t + \mathbf{g}_{t+1}$ $\hat{\mathbf{m}}_{t+1} = \beta_1 \mathbf{m}_{t+1} + \mathbf{g}_{t+1}$ $\mathbf{x}_{t+1} = \mathbf{x}_t - \eta \hat{\mathbf{m}}_{t+1}$	$\frac{2+2\beta_1}{\eta(1+2\beta_1)}$	[7]
EMANESTEROV( $\eta, \beta_1$ )	$\mathbf{m}_{t+1} = \beta_1 \mathbf{m}_t + (1 - \beta_1) \mathbf{g}_{t+1}$ $\hat{\mathbf{m}}_{t+1} = \beta_1 \mathbf{m}_{t+1} + (1 - \beta_1) \mathbf{g}_{t+1}$ $\mathbf{x}_{t+1} = \mathbf{x}_t - \eta \hat{\mathbf{m}}_{t+1}$	$\frac{2+2\beta_1}{\eta(1-\beta_1)(1+2\beta_1)}$	this paper

Table 1: Gradient descent algorithms and their stability thresholds.

**Lemma 1.** Consider running EMANESTEROV( $\eta, \beta_1$ ) on a quadratic objective  $f(\mathbf{x}) = \frac{1}{2} \mathbf{x}^T \mathbf{A} \mathbf{x} + \mathbf{b}^T \mathbf{x} + c$  starting from any initialization. Let  $(\mathbf{q}, a)$  be an eigenvector/eigenvalue pair of  $\mathbf{A}$ . If  $a > \frac{1}{\eta} \left( \frac{2+2\beta_1}{(1-\beta_1)(1+2\beta_1)} \right)$ , then the sequence  $\{\mathbf{q}^T \mathbf{x}\}$  will diverge.

*Proof.* First, we re-write EMANESTEROV( $\eta, \beta_1$ ) as a recursion in  $\mathbf{x}_t$  alone:

$$\mathbf{x}_{t+1} = (1 + \beta_1) \mathbf{x}_t - \beta_1 \mathbf{x}_{t-1} - \eta(1 - \beta_1)(1 + \beta_1) \nabla f(\mathbf{x}_t) + \eta(1 - \beta_1) \nabla f(\mathbf{x}_{t-1})$$

For the quadratic objective, we have  $\nabla f(\mathbf{x}_t) = \mathbf{A} \mathbf{x}_t + \mathbf{b}$ , so this update becomes:

$$\mathbf{x}_{t+1} = (1 + \beta_1)(\mathbf{I} - \eta \mathbf{A}) \mathbf{x}_t - \beta_1(\mathbf{I} - \eta(1 - \beta_1) \mathbf{A}) \mathbf{x}_{t-1} - \eta(1 - \beta_1) \mathbf{b}.$$

This is exactly the recurrence in Theorem 1 of [7], except with  $(1 - \beta_1)\eta$  in place of  $\eta$ . Thus, by the same logic as in Theorem 1 of [7],  $\mathbf{q}^T \mathbf{x}_t$  diverges if  $a > \frac{1}{\eta} \left( \frac{2+2\beta_1}{(1-\beta_1)(1+2\beta_1)} \right)$ .  $\square$

**Lemma 2.** Consider running EMAHB( $\eta, \beta_1$ ) on a quadratic objective  $f(\mathbf{x}) = \frac{1}{2} \mathbf{x}^T \mathbf{A} \mathbf{x} + \mathbf{b}^T \mathbf{x} + c$  starting from any initialization. Let  $(\mathbf{q}, a)$  be an eigenvector/eigenvalue pair of  $\mathbf{A}$ . If  $a > \frac{1}{\eta} \left( \frac{2+2\beta_1}{1-\beta_1} \right)$ , then the sequence  $\{\mathbf{q}^T \mathbf{x}\}$  will diverge.

*Proof.* First, we re-write EMAHB( $\eta, \beta_1$ ) as a recursion in  $\mathbf{x}_t$  alone:

$$\mathbf{x}_{t+1} = (1 + \beta_1) \mathbf{x}_t - \beta_1 \mathbf{x}_{t-1} - \eta(1 - \beta_1) \nabla f(\mathbf{x}_t).$$

For the quadratic objective, we have  $\nabla f(\mathbf{x}_t) = \mathbf{A} \mathbf{x}_t + \mathbf{b}$ , so this update becomes:

$$\mathbf{x}_{t+1} = [(1 + \beta_1) \mathbf{I} - \eta(1 - \beta_1) \mathbf{A}] \mathbf{x}_t - \beta_1 \mathbf{x}_{t-1} - \eta(1 - \beta_1) \mathbf{b}.$$

This is exactly the recurrence in Theorem 2 of [7], except with  $(1 - \beta_1)\eta$  in place of  $\eta$ . Thus, by the same logic as in Theorem 2 of [7],  $\mathbf{q}^T \mathbf{x}_t$  diverges if  $a > \frac{1}{\eta} \left( \frac{2+2\beta_1}{1-\beta_1} \right)$ .  $\square$

## E.2 Stability of preconditioned gradient descent

As discussed in §2, gradient descent can be preconditioned using a preconditioner  $\mathbf{P}$ . For example, the preconditioned version of EMAHB( $\eta, \beta_1$ ), which we denote PRECONEMAHB( $\eta, \beta_1, \mathbf{P}$ ) is:

$$\begin{aligned} \mathbf{m}_{t+1} &= \beta_1 \mathbf{m}_t + (1 - \beta_1) \mathbf{g}_{t+1}, \\ \mathbf{x}_{t+1} &= \mathbf{x}_t - \eta \mathbf{P}^{-1} \mathbf{m}_{t+1}. \end{aligned}$$

On quadratic objective functions, the stability behavior of preconditioned gradient descent algorithms differs slightly from that of their non-preconditioned counterparts. Whereas non-preconditioned gradient descent algorithms become unstable when the maximum eigenvalue of the Hessian  $\mathbf{H}$  exceeds a certain stability threshold, preconditioned gradient descent algorithms become unstable when the maximum eigenvalue of the *preconditioned Hessian*  $\mathbf{P}^{-1/2} \mathbf{H} \mathbf{P}^{-1/2}$  exceeds the threshold of their non-preconditioned counterpart.

The simplest way to see this is to note that preconditioned gradient descent algorithms are isomorphic to their non-preconditioned counterparts up to a certain reparameterization of the training objective. In particular, preconditioned gradient descent on an objective  $f$  using preconditioner  $\mathbf{P}$  is isomorphic to non-preconditioned gradient descent on the reparameterized objective  $\tilde{f}(\mathbf{x}) = f(\mathbf{P}^{-1/2}\mathbf{x})$ , in the sense that the iterates of the two algorithms are related by a linear map. Below, we make this rigorous for EMAHB( $\eta, \beta_1$ ) and PRECONEMAHB( $\eta, \beta_1, \mathbf{P}$ ), but note that one can carry out the same argument for the other momentum algorithms and their preconditioned counterparts:

Let  $\{(\mathbf{x}_t, \mathbf{m}_t)\}$  denote the iterates of PRECONEMAHB( $\eta, \beta_1, \mathbf{P}$ ) on the objective function  $f(\mathbf{x})$ , and let  $\{(\tilde{\mathbf{x}}_t, \tilde{\mathbf{m}}_t)\}$  denote the iterates of EMAHB( $\eta, \beta_1$ ) on the reparameterized objective function  $\tilde{f}(\mathbf{x}) = f(\mathbf{P}^{-1/2}\mathbf{x})$  starting from the initialization  $(\tilde{\mathbf{x}}_0, \tilde{\mathbf{m}}_0) = (\mathbf{P}^{1/2}\mathbf{x}_0, \mathbf{P}^{-1/2}\mathbf{m}_0)$ . Then we have the equivalence  $\tilde{\mathbf{x}}_t = \mathbf{P}^{1/2}\mathbf{x}_t$  and  $\tilde{\mathbf{m}}_t = \mathbf{P}^{-1/2}\mathbf{m}_t$  for all steps  $t$ .

*Proof.* The equivalence is true at initialization, by definition. Thus, it remains to show that if the equivalence is true at step  $t$ , it is true at step  $t + 1$ .

The updates for  $\{(\mathbf{x}_{t+1}, \mathbf{m}_{t+1})\}$  are:

$$\begin{aligned}\mathbf{m}_{t+1} &= \beta_1 \mathbf{m}_t + (1 - \beta_1) \nabla f(\mathbf{x}_t), \\ \mathbf{x}_{t+1} &= \mathbf{x}_t - \eta \mathbf{P}^{-1} \mathbf{m}_{t+1}.\end{aligned}$$

Meanwhile, the updates for  $\{(\tilde{\mathbf{x}}_{t+1}, \tilde{\mathbf{m}}_{t+1})\}$  are:

$$\begin{aligned}\tilde{\mathbf{m}}_{t+1} &= \beta_1 \tilde{\mathbf{m}}_t + (1 - \beta_1) \nabla \tilde{f}(\tilde{\mathbf{x}}_t) \\ &= \beta_1 \tilde{\mathbf{m}}_t + (1 - \beta_1) \mathbf{P}^{-1/2} \nabla f(\mathbf{P}^{-1/2} \tilde{\mathbf{x}}_t), \\ \tilde{\mathbf{x}}_{t+1} &= \tilde{\mathbf{x}}_t - \eta \tilde{\mathbf{m}}_{t+1}.\end{aligned}$$

To verify the equivalence  $\tilde{\mathbf{m}}_{t+1} = \mathbf{P}^{-1/2} \mathbf{m}_{t+1}$ , observe that

$$\begin{aligned}\mathbf{P}^{-1/2} \mathbf{m}_{t+1} &= \beta_1 \mathbf{P}^{-1/2} \mathbf{m}_t + (1 - \beta_1) \mathbf{P}^{-1/2} \nabla f(\mathbf{x}_t) \\ &= \beta_1 \tilde{\mathbf{m}}_t + (1 - \beta_1) \mathbf{P}^{-1/2} \nabla f(\mathbf{P}^{-1/2} \tilde{\mathbf{x}}_t) \\ &= \tilde{\mathbf{m}}_{t+1}.\end{aligned}$$

where the two replacements in the middle step are due to the inductive hypothesis holding for step  $t$ .

To verify the equivalence  $\tilde{\mathbf{x}}_{t+1} = \mathbf{P}^{1/2} \mathbf{x}_{t+1}$ , observe that

$$\begin{aligned}\mathbf{P}^{1/2} \mathbf{x}_{t+1} &= \mathbf{P}^{1/2} \mathbf{x}_t - \eta \mathbf{P}^{1/2} \mathbf{P}^{-1} \mathbf{m}_{t+1} \\ &= \tilde{\mathbf{x}}_t - \eta \tilde{\mathbf{m}}_{t+1} \\ &= \tilde{\mathbf{x}}_{t+1}.\end{aligned}$$

where the two replacements in the middle step are due to the inductive hypothesis holding for step  $t$ .  $\square$

Due to Proposition E.2, we know that PRECONEMAHB( $\eta, \beta_1, \mathbf{P}$ ) on a quadratic function

$$f(\mathbf{x}) = \frac{1}{2} \mathbf{x}^T \mathbf{A} \mathbf{x} + \mathbf{b}^T \mathbf{x} + c$$

has the same stability behavior as EMAHB( $\eta, \beta_1$ ) on the reparameterized quadratic

$$\tilde{f}(\mathbf{x}) = \frac{1}{2} \mathbf{x}^T (\mathbf{P}^{-1/2} \mathbf{A} \mathbf{P}^{-1/2}) \mathbf{x} + \mathbf{b}^T \mathbf{P}^{-1/2} \mathbf{x} + c.$$

Thus, PRECONEMAHB( $\eta, \beta_1, \mathbf{P}$ ) diverges on  $f$  whenever the maximum eigenvalue of  $(\mathbf{P}^{-1/2} \mathbf{A} \mathbf{P}^{-1/2})$  exceeds the stability threshold of EMAHB( $\eta, \beta_1$ ). When this occurs, the component of  $\mathbf{P}^{1/2} \mathbf{x}_t$  that is aligned with  $\mathbf{v}_1$ , the top eigenvector of the preconditioned Hessian, will oscillate with exponentially increasing magnitude. Equivalently, we can say that the component of  $\mathbf{x}_t$  that is aligned with  $(\mathbf{P}^{1/2})^T \mathbf{v}_1$  will oscillate with exponentially increasing magnitude.

## F Adaptive preconditioning

In Figure F, we optimize the objective function  $f(x) = \frac{1}{2}x^2$  using rmsprop with learning rates that would be divergent if the preconditioner  $\nu_t$  were held fixed at its initial value of  $\nu_0 = 1.0$ . Yet since the preconditioner is adaptive rather than fixed, rmsprop only oscillates explosively for a finite number of steps before the large gradients associated with this explosive movement cause  $\nu_t$  to increase, which decreases rmsprop's effective learning rate and halts divergence.

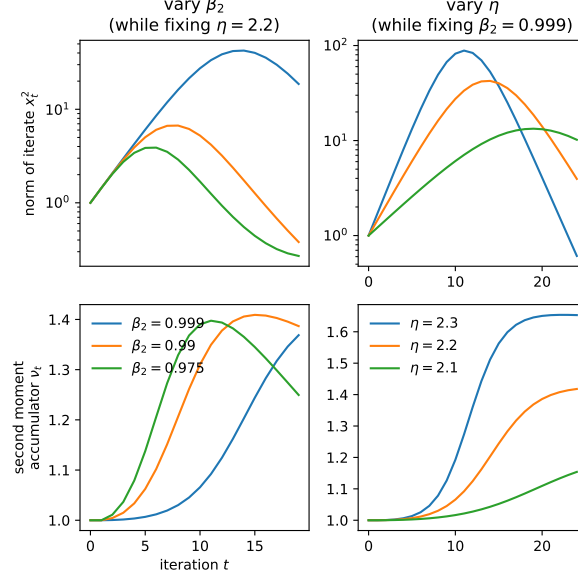


Figure 12: **Adaptive preconditioning often prevents rmsprop from diverging.** We optimize the objective function  $f(x) = \frac{1}{2}x^2$  using rmsprop. In the left column, we hold  $\eta$  fixed at  $\eta = 2.2$  and vary  $\beta_2$ ; in the right column, we hold  $\beta_2$  fixed at  $\beta_2 = 0.999$  and vary  $\eta$ . Observe that rmsprop only oscillates explosively for a finite number of steps before the large gradients associated with this explosive movement cause  $\nu_t$  to increase, which decreases rmsprop's effective learning rate and halts divergence.

## G Qualitative aspects of the AEOs

Looking closely at Figure 4, one can discern that the preconditioned sharpness behaves in qualitatively different ways depending on the type of momentum. For rmsprop and Adagrad, which do not employ any momentum, the preconditioned sharpness hovers above the stability threshold. For Nadam, too, which employs Nesterov-style momentum, the preconditioned sharpness also hovers above the stability threshold. In contrast, for the other algorithms, which employ heavy-ball style momentum, the preconditioned sharpness does not hover above the stability threshold — instead, the optimizer is rapidly thrown into flatter regions whenever the preconditioned sharpness crosses the threshold; as a result, the preconditioned sharpness oscillates around the stability threshold instead of hovering above it.

We now (1) confirm that these behaviors also occur in the simpler setting of non-adaptive gradient descent, and (2) suggest a speculative explanation.

**Non-adaptive gradient descent** The pattern described above also holds for non-adaptive gradient descent. For example, Figures 29, 95, and 96 in [7] depict the evolution of the sharpness as a network is trained using vanilla gradient descent, heavy ball momentum, and Nesterov momentum, respectively. In both Figure 29 (vanilla) and Figure 96 (Nesterov), the sharpness hovers above the stability threshold, whereas in Figure 95 (heavy ball), the sharpness oscillates around the threshold.

We now reproduce this observation using the same network from Figure 4. In Figure 13, we train the fully-connected network on CIFAR-10 using vanilla gradient descent  $\text{GD}(\eta)$ , heavy ball gradient descent  $\text{STANDARDHB}(\eta, 0.9)$ , and Nesterov gradient descent  $\text{STANDARDNESTEROV}(\eta, 0.9)$ , each at a range of learning rates. Observe that for vanilla and Nesterov gradient descent, the sharpness hovers above the stability threshold (except at the smallest learning rate considered), whereas for heavy ball gradient descent, the sharpness oscillates around the stability threshold.

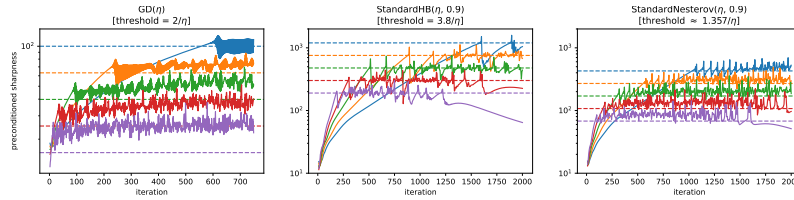
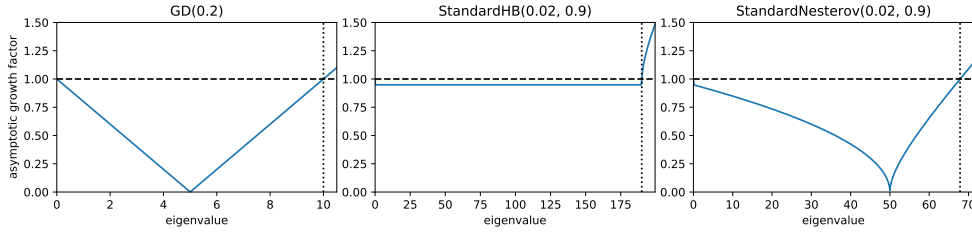


Figure 13: **Similar to Figure 4, the behavior of the sharpness during (non-adaptive) gradient descent depends on the type of momentum.** We train a fully-connected network on CIFAR-10 using three non-adaptive algorithms: vanilla gradient descent ( $\eta$  in logspace between 0.02 and 0.2), heavy ball momentum ( $\beta_1 = 0.9$ , and  $\eta$  in logspace between 0.002 and 0.02), and Nesterov momentum ( $\beta_1 = 0.9$ , and  $\eta$  in logspace between 0.002 and 0.02). Observe that for vanilla GD and Nesterov, the sharpness hovers above the stability threshold, whereas for heavy ball, it oscillates around the stability threshold. This pattern parallels that in Figure 4 for Adam.

**A speculative explanation** We now give a speculative explanation for the above observations. One difference between heavy ball gradient descent, on the one hand, and Nesterov and vanilla gradient descent, on the other, is that heavy ball momentum can be much less robust to increases in the sharpness beyond the optimizer’s stability threshold. By way of background, recall from [7] that on quadratic functions, all three of these algorithms act independently along each Hessian eigenvector; in particular, the component of the iterate that is aligned with each Hessian eigenvector evolves according to a linear recurrence relation whose coefficients are determined by the corresponding Hessian eigenvalue  $\lambda$  [7]. For example, the recurrence relation for  $\text{GD}(0.2)$  is  $x_{t+1} = (1 - 0.2\lambda)x_t$ , the recurrence relation for  $\text{STANDARDHB}(0.02, 0.9)$  is  $x_{t+1} = (1.9 - 0.02\lambda)x_t - 0.9x_{t-1}$ , and the recurrence relation for  $\text{STANDARDNESTEROV}(0.02, 0.9)$  is  $x_{t+1} = 1.9(1 - 0.02\lambda)x_t - 0.9(1 - 0.02\lambda)x_{t-1}$ . Asymptotically (i.e. over many optimizer steps), the growth rate of a linear recurrence relation is determined by the magnitude of the largest root of the characteristic polynomial. (Note that for a recurrence of the form  $x_{t+1} = ax_t + bx_{t-1}$ , the characteristic polynomial is  $x^2 - ax - b$ .) The stability threshold  $\lambda^*$  of an optimizer is the value of  $\lambda$  for which the magnitude of the largest-magnitude root is exactly 1. The optimizer’s robustness to barely-unstable eigenvalues ( $\lambda$  just above  $\lambda^*$ ) can be quantified by the asymptotic growth rate at those eigenvalues.

For example, consider the three optimization algorithms  $\text{GD}(0.2)$ ,  $\text{STANDARDHB}(0.02, 0.9)$ , and  $\text{STANDARDNESTEROV}(0.02, 0.9)$ . These algorithms all share the same effective learning rate of 0.2, and they correspond to the purple lines in Figure 13. In Figure 14, we visualize the asymptotic growth factor of these algorithms as a function of the input eigenvalue  $\lambda$ . The stability threshold  $\lambda^*$  is marked with a vertical dotted black line (of course, the asymptotic growth factor at  $\lambda^*$  is always 1), and we plot the asymptotic growth factor for eigenvalues up to 1.1 times the stability threshold.

Crucially, observe that the asymptotic growth rate for  $\text{STANDARDHB}(0.02, 0.9)$  at an eigenvalue 1.1 times its stability threshold is much higher than for  $\text{GD}(0.02)$  and  $\text{STANDARDNESTEROV}(0.02, 0.9)$ . This means that whenever  $\text{STANDARDHB}$  finds itself in a region of parameter space where the sharpness is 1.1 times its stability threshold, it diverges much faster than  $\text{GD}(0.02)$  or  $\text{STANDARDNESTEROV}(0.02, 0.9)$  do in an analogous situation. One can easily imagine that this may be related to our observation that the latter two algorithms can train at the AEoS with the sharpness hovering a bit above the stability threshold, whereas heavy ball momentum is flung into a flatter region whenever the sharpness crosses the stability threshold even by a small amount.



**Figure 14: Vanilla gradient descent and Nesterov momentum are both more robust than heavy ball momentum to eigenvalues shortly above the stability threshold.** We plot the asymptotic growth factors of three algorithms: vanilla GD, heavy ball GD, and Nesterov GD. That is, for each possible Hessian eigenvalue, we plot the multiplicative constant which the corresponding component of the iterate will be asymptotically multiplied by, over many steps of gradient descent. The vertical dotted line marks each algorithm’s stability threshold, and we plot the asymptotic growth factor (AGF) for eigenvalues up to 1.1 times that stability threshold. Observe that for heavy ball momentum, the AGF at 1.1 times the stability threshold is much higher than for Nesterov momentum or vanilla gradient descent.

## H Corner case

While we generally observe that during full-batch Adam, the preconditioned sharpness equilibrates at the stability threshold of frozen Adam, we have also consistently observed one notable corner case where this does not occur. Namely, when running Adam at *extremely small* learning rates, we often find that the preconditioned sharpness either flatlines short of the stability threshold, or increases at a very slow rate while remaining below the threshold. These behaviors are evident in Figure 15, in which we train a CNN on CIFAR-10 using full-batch Adam at a range of very small learning rates. (This task is the same one as in Figure 5, left panel, just at smaller learning rates.) Observe that at these small learning rates, the preconditioned sharpness fails to rise to the stability threshold.

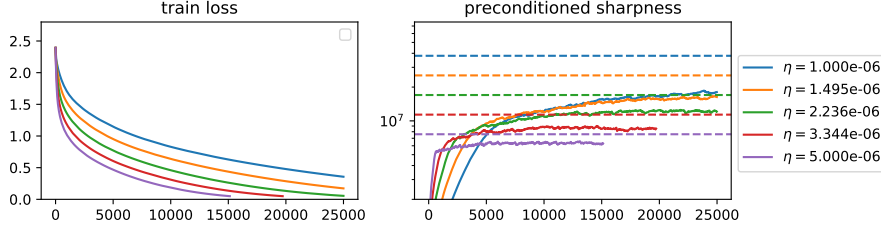


Figure 15: **When running Adam at *extremely small* learning rates, the sharpness can fail to rise all the way to the stability threshold.** We train a CNN on CIFAR-10 using full-batch Adam at a range of extremely small learning rates. Observe that the preconditioned sharpness fails to rise all the way to the stability threshold.

Why does the preconditioned sharpness not rise to the stability threshold? One clue is that this behavior seems to be related to adaptivity. In Figure 16, we train using full-batch Adam at the green learning rate from Figure 15, and at step 6000 of training, we suddenly freeze the preconditioner (i.e. we stop updating  $\nu_t$ ). The train loss and preconditioned sharpness of frozen Adam are plotted in pink. Observe that upon freezing the preconditioner, the preconditioned sharpness rapidly rises all the way to the stability threshold. Thus, we can infer that it was due to Adam’s adaptivity that the preconditioned sharpness was located below the stability threshold in the first place.

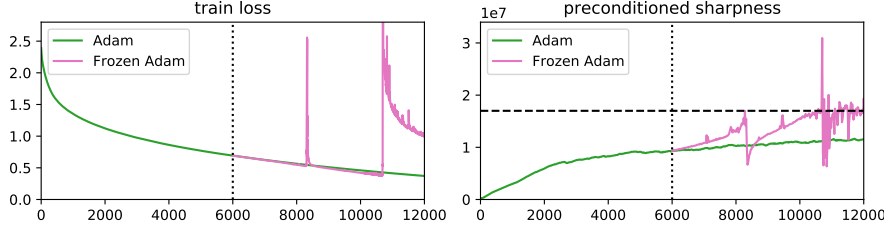
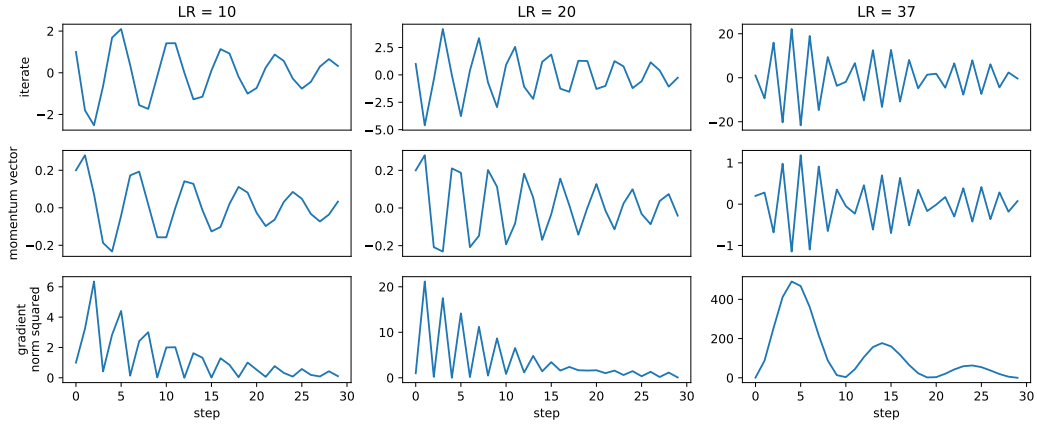


Figure 16: **If we freeze the preconditioner, the preconditioned sharpness immediately rises to the stability threshold.** In green, we train a CNN on CIFAR-10 using full-batch Adam at step size  $\eta = 2.236e-6$  (the green learning rate from Figure 15). In pink, we freeze the preconditioner at step 6,000. Observe that once the preconditioner is frozen, the preconditioned sharpness rises immediately to the stability threshold, suggesting that adaptive preconditioning was preventing it from rising to the stability threshold in the first place.

We hypothesize that this behavior is related to the following property of momentum algorithms: momentum algorithms can generate large gradients even when they are stable. For example, in Figure 17, we optimize the quadratic objective  $f(x) = \frac{1}{2}x^2$  using EMA heavy ball momentum at  $\beta_1 = 0.9$  and three different stable learning rates (columns). We set the initial iterate  $x_0$  to 1, and the initial momentum vector  $m_0$  to 0.2. We run EMAHB( $\eta, 0.9$ ) for thirty steps, plotting the iterate (top row), the momentum vector (middle row), and the gradient squared norm (bottom row). Observe that the gradient squared norm increases dramatically from its initial value.

We therefore hypothesize that the preconditioned sharpness  $\lambda_1(\mathbf{P}_t^{-1}\mathbf{H}_t)$  fails to rise all the way to the stability threshold because the large gradients generated by the momentum algorithm (even while

it is stable) are causing the preconditioner  $\mathbf{P}_t$  to constantly increase. This, in turn “cancels out” the effect of progressive sharpening, and causes the preconditioned sharpness  $\lambda_1(\mathbf{P}_t^{-1}\mathbf{H}_t)$  to remain short of the stability threshold.



**Figure 17: Momentum can generate large gradients even when stable.** We optimize the function  $f(x) = \frac{1}{2}x^2$  using 30 steps of EMAHB( $\eta, 0.9$ ) at various stable learning rates  $\eta$ . The top row plots the iterate, the second row the momentum vector (really a scalar), and the third row the squared gradient norm. Observe that the gradient norm increases from its initial value, sometimes dramatically. This does not happen with vanilla gradient descent.

## I Further plots of zoom in

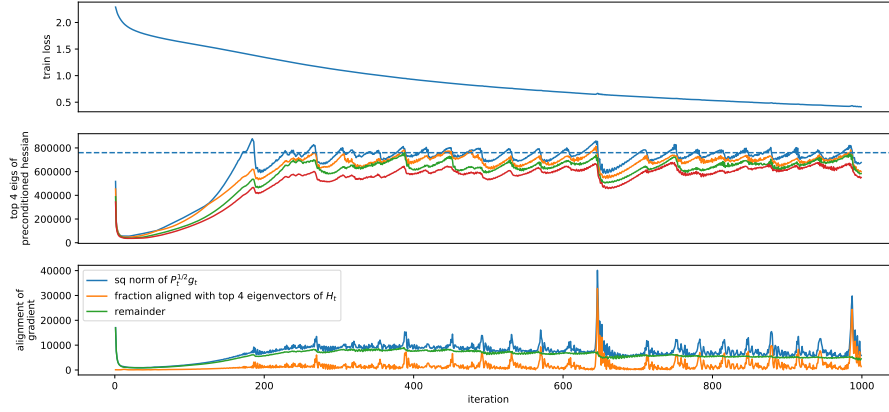


Figure 18: **Explosive growth in the (preconditioned) gradient is limited to the top eigenvectors of the preconditioned Hessian.** We train a fully-connected network on CIFAR-10 using full-batch Adam. In the top row, we plot the train loss. In the middle row, we plot the evolution of the top four eigenvalues of the preconditioned Hessian  $\mathbf{P}_t^{-1}\mathbf{H}_t$ . In the bottom row, we plot (1) in blue, the squared norm of the “semi-preconditioned gradient”  $\mathbf{P}^{1/2}\mathbf{g}_t$ , which is the quantity that is expected to grow explosively when the algorithm is unstable; in orange, the squared norm of the portion of  $\mathbf{P}^{1/2}\mathbf{g}_t$  that is aligned with the top four eigenvectors of  $\mathbf{P}_t^{-1}\mathbf{H}_t$ ; and in green, the squared norm of the remainder of  $\mathbf{P}^{1/2}\mathbf{g}_t$ . Observe that almost all of the spikes in the gradient norm can be attributed to the top eigenvectors of the preconditioned Hessian.

## J Test accuracy of Wide ResNet

In §D, we saw that Adam tended to find solutions with higher sharpness (maximum Hessian eigenvalue) whenever the learning rate  $\eta$  was smaller or the decay factor  $\beta_2$  was smaller. We now demonstrate that the test error tends to be higher in these same situations. In Figure 19, we plot the test error of the Wide ResNet from Figure 7 as a function of  $\eta$  and  $\beta_2$ . (For each training run, we select the test error at the moment the train loss first drops below the threshold of 0.05.) As in Figure 7, we experiment with both the full-batch and large batch (batch size 4096) settings. Observe that across both these settings, the test error tends to be higher whenever  $\eta$  is small (subplots (a) or (b)) or whenever  $\beta_2$  is small (subplots (c) or (d)). We emphasize that we do not intend to claim that there is necessarily a causal connection between the test error and the sharpness — we only mean to point out that the Adam’s implicit bias towards low-curvature regions when  $\eta$  is high or  $\beta_2$  is high coincides with superior generalization performance.

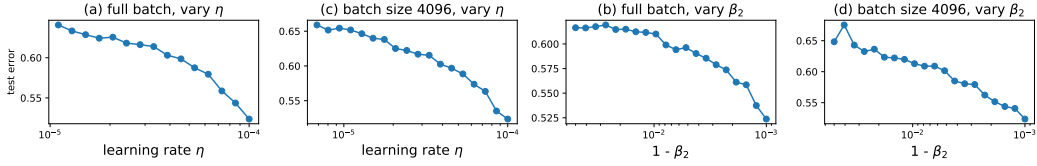


Figure 19: **Adam tends to generalize better when  $\eta$  or  $\beta_2$  is small.** We train a Wide ResNet on CIFAR-100 using full-batch (panels a, c) or large-batch (panels b, d) Adam at various  $\eta$  (panels a, b) or  $\beta_2$  (panels b, d). We plot the test error, measured when the training loss first drops below the threshold 0.05.

## K Experimental details

### K.1 Architectures

**Fully-connected network** This network has 5 hidden layers of width 200 each. The activation function is tanh, and the weights are initialized with variance  $1 / \text{fan\_in}$ . The fully-connected layers have bias parameters, which are initialized to zero.

**CNN** This network repeats (Conv  $\Rightarrow$  ReLU  $\Rightarrow$  maxpool) three times, then flattens the activations and does FC  $\Rightarrow$  ReLU  $\Rightarrow$  FC  $\Rightarrow$  ReLU  $\Rightarrow$  FC. The first conv layer has 64 filters, a kernel size of 5, “valid” padding, a 3x3 window size, and stride 2. The second conv layer has 96 filters, a kernel size of 3, “valid” padding, a 3x3 window size, and stride 2. The third conv layer has 128 filters, a kernel size of 3, “same” padding, a 3x3 window size, and stride 2. The first fully-connected layer has width 512, and the second has width 256. All biases are initialized to zero, and all weights are initialized with variance  $1 / \text{fan\_in}$ .

**Wide ResNet (normalized)** This Wide ResNet has four blocks per group, and a channel multiplier of 10.

**Wide ResNet (un-normalized)** When we remove batch normalization, we insert a learnable scalar multiplier (initialized to 1) after the residual branch of each block.

## L The role of Adam’s hyperparameters

The preceding discussion suggests that Adam’s behavior at the AEoS — and in particular, the extent to which the stability-based constraint on  $\lambda_1(\mathbf{P}_t^{-1}\mathbf{H}_t)$  is maintained by adapting  $\mathbf{P}$  vs. by constraining  $\mathbf{H}$  — may depend on Adam’s hyperparameters. To test this hypothesis, in Figure 7(a), we use full-batch Adam at a range of learning rates to train a Wide Resnet on CIFAR-100 until reaching a training loss value of 0.1. We plot the sharpness  $\lambda_1(\mathbf{H}(\mathbf{x}^*))$  at the solution  $\mathbf{x}^*$ , as a function of the learning rate. Observe that full-batch Adam at small learning rates finds sharper solutions. In Figure 7(b), we confirm that this trend extends to minibatch Adam at the large batch size of 4096. Meanwhile, in Figure 7(c), we sweep full-batch Adam at a range of values of  $\beta_2$ . We observe that when  $\beta_2$  is small, Adam tends to find sharper solutions. Figure 7(d) demonstrates that this trend extends to minibatch Adam at batch size 4096. Incidentally, we show in Appendix J that the models trained with higher  $\eta$  or higher  $\beta_2$  generalize better.

One speculative hypothesis is that Adam’s ability to resolve instabilities via preconditioner adaptation is contingent on how fast the preconditioner adapts (smaller  $\beta_2$  = faster adaptation) relative to the speed at which the optimizer is moving (larger  $\eta$  = faster movement).

### L.1 Figures

**Figure 3** We extracted the fixed preconditioner from step 1100 of a network trained using full-batch Adam at learning rate  $2\text{e-}5$ .

**Figure 5** In the left pane, we trained for 2000 steps or until reaching a threshold loss value of 0.05. In the middle pane, we trained for 5000 steps. In the right pane, we trained until reaching a threshold loss value of 0.05.

**Figure 4 Adam with bias correction:**  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ ,  $\epsilon = 1\text{e-}7$ , learning rates were:

`np.logspace(np.log10(3e-5), np.log10(1e-3), 5)`

**AdamW:**  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ ,  $\epsilon = 1\text{e-}7$ , learning rates were:

`np.logspace(np.log10(3e-5), np.log10(1e-3), 5)`

**Adafactor:** The decay rate was 0.8, the momentum was 0.9, epsilon was  $1\text{e-}7$ , and the learning rates were

`np.logspace(-5, -3, 5)`

**Amsgrad:**  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ ,  $\epsilon = 1\text{e-}7$ , bias correction was off, and learning rates were:

`np.logspace(-5, -3, 5)`

**Adam:**  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ ,  $\epsilon = 1\text{e-}7$ , bias correction was off, the Adam exponent was 0.25, and learning rates were:

`np.logspace(-3, -1, 5)`

**Nadam:**  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ ,  $\epsilon = 1\text{e-}7$ , bias correction was off, and learning rates were:

`np.logspace(-5, -3, 5)`

**rmsprop:** decay was 0.995, epsilon was  $1\text{e-}7$ , and learning rates were:

`np.logspace(-4, -3, 5)`

**Adagrad:** Learning rates were `np.logspace(-2, -1, 5)`.

**Figure 7** For the learning rate sweeps we fixed  $\beta_2 = 0.999$ , and for the  $\beta_2$  sweeps we fixed  $\eta = 1\text{e-}4$ .

**Figure 11** For Adam we used  $\beta_1 = 0.9$  and  $\beta_2 = 0.999$ ; for momentum we used  $\beta_1 = 0.9$ .