

PAPER: Interdisciplinary statistical mechanics

Catapult dynamics and phase transitions in quadratic nets

David Meltzer¹, Min Chen² and Junyu Liu^{2,3,*}¹ Department of Physics, Cornell University, Ithaca, NY 14850, United States of America² Department of Computer Science, The University of Pittsburgh, Pittsburgh, PA 15260, United States of America³ Pritzker School of Molecular Engineering, The University of Chicago, Chicago, IL 60637, United States of AmericaE-mail: junyuliu@pitt.edu

Received 6 January 2025

Accepted for publication 20 July 2025

Published 25 September 2025



CrossMark

Online at stacks.iop.org/JSTAT/2025/093406
<https://doi.org/10.1088/1742-5468/adf821>

Abstract. Neural networks trained with gradient descent can undergo non-trivial phase transitions as a function of the learning rate. In Lewkowycz *et al* (2020 arXiv:2003.02218) it was discovered that wide neural nets can exhibit a catapult phase for super-critical learning rates, where the training loss grows exponentially quickly at early times before rapidly decreasing to a small value. During this phase the top eigenvalue of the neural tangent kernel also undergoes significant evolution. In this work, we will prove that the catapult phase exists in a large class of models, including quadratic models and two-layer, homogeneous neural nets. To do this, we show that for a certain range of learning rates the weight norm decreases whenever the loss becomes large. We also empirically study learning rates beyond this theoretically derived range and show that the activation map of ReLU nets trained with super-critical learning rates becomes increasingly sparse as we increase the learning rate.

Keywords: machine learning, network dynamics, nonlinear dynamics

* Author to whom any correspondence should be addressed.



Original Content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

Contents

1. Introduction 3

2. Related works 4

3. Critical points in the learning rate 6

 3.1. Model definitions 6

 3.2. Summary of results 8

 3.3. Sketch of a proof 10

4. Experiments 11

 4.1. Linear meta-feature function 11

 4.2. Generic homogeneous net 12

 4.3. ReLU net 13

5. Conclusions 14

Acknowledgment 16

Appendix A. Early time dynamics 16

 A.1. Catapult phase in the quadratic model 16

 A.2. Weight norm in the quadratic model 20

 A.3. Homogeneous, two-layer MLP 22

Appendix B. Derivations for a single datapoint 23

 B.1. Pure quadratic model 23

 B.2. Quadratic model with bias 26

 B.3. Homogeneous, two-layer, MLPs 29

 B.4. ReLU net 31

Appendix C. Multiple datapoints 35

 C.1. Pure quadratic model 35

 C.2. Quadratic model with bias 40

 C.3. Homogeneous, two-layer, MLPs 42

 C.4. Justifications for the simplifications 43

Appendix D. More Experiments 44

 D.1. Quadratic model 45

 D.1.1. Definitions 45

 D.1.2. Pure quadratic model 46

 D.1.3. Quadratic model with bias 47

 D.2. Generic two-layer homogeneous MLPs 48

 D.3. ReLU MLPs for images 51

 D.3.1. Architectures and datasets 51

D.3.2. Two layers	51
D.3.3. Three layers	54
References.....	57

1. Introduction

In recent years, the field of deep learning has seen a marked revival and lead to remarkable progress in machine learning (Goodfellow *et al* 2016). However, there remain many open problems concerning its theoretical foundations. Recent progress has been driven in large part by the development of neural tangent kernel (NTK) theory (Lee *et al* 2017, 2019, Jacot *et al* 2018, Arora *et al* 2019, Yaida 2019, Sohl-Dickstein *et al* 2020, Roberts *et al* 2022). In particular, it has been proven, under certain conditions on the learning rate and weight initialization, that infinitely wide neural nets reduce to linear models. To study more realistic models, which can learn new representations of the data, one can go beyond the strict infinite width and study perturbative, $1/\text{width}$ corrections (Dyer and Gur-Ari 2019, Huang and Yau 2019, Yaida 2020).

In this work we are interested in understanding models that display non-trivial, non-perturbative dynamics when trained with full-batch gradient descent. Specifically, we will study the *catapult mechanism* of Lewkowycz *et al* (2020), which occurs when we train the model with a learning rate η that is larger than the maximal learning rate of the corresponding linearized, or infinite-width, model. We will refer to learning rates above the naive, linear stability threshold as ‘super-critical’.

The catapult phase refers to an intermediate dynamical regime of gradient descent that lies between the so-called ‘lazy phase’ and ‘divergent phase’. In this regime, the training loss initially increases before turning around and decreasing to a small value. The phase is characterized by a transient growth of the NTK or related metrics, and is enabled by large learning rates that exceed the linearized stability threshold termed as super-critical learning rate above. While in the lazy phase the learning rate is sufficiently small such that the model weights do not deviate significantly from their values at initialization. In this regime one can use large width perturbation theory to make predictions. In the ‘divergent phase’ the learning rate is taken to be large enough such that the loss diverges. In Lewkowycz *et al* (2020) it was demonstrated in a variety of examples that the catapult phase leads to a smaller generalization loss than the lazy phase.

Overall, our understanding of the catapult phase, e.g. why it exists and how the neural net evolves in this phase, pales in comparison to our knowledge about neural nets in the lazy phase. To gain a better understanding of the catapult phase, we will study two classes of toy models: quadratic models and homogeneous MLPs with one hidden layer. Quadratic models are minimal models for representation learning (Roberts *et al* 2022), as they are the simplest deviation from linear models whose NTK evolves under gradient descent. We will derive sufficient, but not necessary, conditions for quadratic models to converge at large learning rates in terms of the maximal eigenvalue of the meta-feature function, which we will define later. We will also derive sufficient conditions

for homogeneous MLPs to have a catapult phase, which will depend on the slope of the activation function in the positive/negative regions. For both classes of models, the existence of the catapult phase is proven by showing that there exists a range of super-critical learning rates such that the weight norm θ^2 decreases whenever the loss becomes large. We will provide evidence for all analytical results by numerically studying various quadratic models and two-layer MLPs.

More generally, the weight norm θ^2 is an interesting quantity to study because its behavior changes qualitatively as we change the learning rate *inside* the catapult phase. We observe that for sufficiently large learning rates, the weight norm in ReLU nets can increase significantly over the course of training. We empirically observe that this increase in the weight norm does not hurt the generalization performance of ReLU MLPs. We conjecture that these models still generalize well because the activation map of ReLU nets becomes increasingly sparse as we increase the learning rate. We believe this result points to interesting future directions for the study of ReLU nets at large learning rates.

Notations: We use bold letters, e.g. θ , ψ to refer to vectors, matrices, etc in weight space. We use the corresponding un-bolded letter when referring to individual indices, e.g. θ_μ and $\psi_{\mu\nu}$ refer to specific components. We also use un-bolded letters for one-dimensional quantities in weight space. We use Greek letters from the middle of the alphabet to refer to indices in weight space, as above, and Greek letters from the beginning of the alphabet to refer to points in sample space, e.g. $(\mathbf{x}_\alpha, \mathbf{y}_\alpha)$ refers to the α th data-point and its corresponding label. We will keep the indices in sample space explicit, except when we consider toy examples with only one data-point, in which case we drop this index for compactness. The symbols \mathcal{N} and \mathcal{U} refer to the normal and uniform distribution, respectively.

For vector-norms we use the notation $\theta^2 \equiv \|\theta\|_2^2$, where $\|\cdot\|_2$ is the L_2 norm. For matrices we use the norm induced by the vector L_2 norm, $\|\psi\| \equiv \|\psi\|_2$. In addition, we use $\lambda_i(\psi)$ to refer to the i th eigenvalue. We use the λ_i notation for matrices in both weight and sample space. For the NTK we use $\lambda_{\max}(H_{\alpha\beta})$ and $\|H_{\alpha\beta}\|$ interchangeably⁴.

Finally, we use the t -subscript to refer to quantities at step t of gradient descent. We will exclusively use the 0 subscript to denote quantities at initialization, e.g. θ_0 is the vector θ at initialization and not its 0th component.

2. Related works

The catapult phase was first analyzed by Lewkowycz *et al* (2020), who studied gradient descent dynamics in a variety of architectures, including linear models, deep fully-connected networks, convolutional nets, and residual networks. They showed that in the

⁴ Since we reserve index-free notation for objects in weight space, $\lambda_{\max}(H_{\alpha\beta}) = \|H_{\alpha\beta}\|$ is the top eigenvalue of the full NTK matrix, and does not refer to the norm of individual components.

catapult phase, models can access large learning rates stably, and often achieve better generalization performance compared to those trained in the lazy phase. Theoretically, they analyzed a finite-width linear MLP trained with mean-squared error (MSE) loss, and identified phase boundaries based on the learning rate and Hessian curvature. Empirically, they demonstrated that models trained in the catapult regime not only converge, but may also exhibit more favorable generalization ability. Besides, the catapult phase in models trained with logistic loss was studied in Huang *et al* (2020). The effect of large learning rates on matrix factorization was studied in Wang *et al* (2021).

The work we present here is similar to the interesting analysis of Zhu *et al* (2022). There they proved that the catapult phase exists in the quadratic model that approximates a two-layer ReLU MLP. They also demonstrated that the generalization loss of neural nets and their quadratic approximations are qualitatively similar in the catapult phase. Our work differs from Zhu *et al* (2022) in a few key areas. The first is, we will show that the catapult phase exists in a large class of quadratic models, while in Zhu *et al* (2022) they focused on the quadratic approximation of a two-layer ReLU MLP. We also show that the catapult phase exists in general two-layer, homogeneous MLPs, without making a quadratic approximation. We believe it would be interesting to understand how their analysis for the quadratic approximation to the ReLU net could be used to study the class of quadratic models considered in this work.

Quadratic models have also been studied more generally in the literature. In Roberts *et al* (2022) the quadratic model was proposed as a minimal model to study feature learning. In Agarwala *et al* (2022) they proved that in quadratic models the maximal eigenvalue of the NTK displays the edge of stability (EOS) phenomena (Cohen *et al* 2021)⁵. Taylor expansions of neural nets were also studied in Bai *et al* (2020). The pure quadratic model with $\psi(x)$ defined in section 3.1 with specific form being xx^\top is equivalent to the classical phase retrieval problem, which has been extensively studied in both signal processing and machine learning (Candes *et al* 2015, Sun *et al* 2018, Davis *et al* 2020, Sarao Mannelli *et al* 2020a, 2020b, Mignacco *et al* 2021, Arnaboldi *et al* 2023, Martin *et al* 2024). These works analyze the optimization landscape and the dynamics of GD/SGD under small and moderate learning rates with various regimes. In contrast, we aim to analyze the dynamics and implicit regularization in the large learning rate regime, particularly the catapult phase and associated NTK evolution.

Finally, we also empirically study the sparsity of ReLU MLPs in the catapult phase as a function of the learning rate. For a review on sparsity see Hoefler *et al* (2021). Two recent works (Andriushchenko *et al* 2022, Li *et al* 2022) show that increasing the learning rate promotes sparsity, but to our knowledge they do not consider training the model in the catapult phase. Large initial learning rates in ReLU nets have also been studied in Iyer *et al* (2022).

⁵ In Cohen *et al* (2021) the EOS was studied for the top eigenvalue of the Hessian. As explained in appendix A of Agarwala *et al* (2022), if the model is trained with MSE then the spectrum of the NTK approaches that of the Hessian as the model converges.

3. Critical points in the learning rate

3.1. Model definitions

Quadratic models. We start by defining the general quadratic model. We take the data-points to be d -dimensional, $\mathbf{x}_\alpha \in \mathbb{R}^d$, where $\alpha = 1, \dots, D$, and the labels to be one-dimensional $y_\alpha \in \mathbb{R}$.

The general quadratic model (for a $1d$ output) is defined by:

$$z(\mathbf{x}_\alpha) = \boldsymbol{\theta}^T \boldsymbol{\phi}(\mathbf{x}_\alpha) + \frac{\zeta}{2} \boldsymbol{\theta}^T \boldsymbol{\psi}(\mathbf{x}_\alpha) \boldsymbol{\theta}, \tag{1}$$

where $z(\mathbf{x}) \in \mathbb{R}$, $\boldsymbol{\theta} \in \mathbb{R}^n$, $\boldsymbol{\phi}(\mathbf{x}) \in \mathbb{R}^n$, $\boldsymbol{\psi}(\mathbf{x}) \in \mathbb{R}^{n \times n}$, and $\zeta \in \mathbb{R}$. The weights are drawn from a normal distribution with unit variance, $\boldsymbol{\theta} \sim \mathcal{N}(0, \mathbb{I}_{n \times n})$. The functions $\boldsymbol{\phi}$ and $\boldsymbol{\psi}$ are part of the definition of the model and are not trainable. We will refer to $\boldsymbol{\phi}$ and $\boldsymbol{\psi}$ as the feature and meta-feature functions, respectively.

The parameter ζ is in principle arbitrary and measures the deviation of (1) from linearity. When $\zeta = 0$ the quadratic model (1) reduces to a linear model, which has static features. When $\zeta \neq 0$ the above model can exhibit non-trivial representation learning. In this work we take $0 < \zeta \ll 1$, but we will still see large deviations from linearity when the learning rate η is taken to be sufficiently large.

We train the model to minimize the MSE:

$$L = \frac{1}{2D} \sum_{\alpha=1}^D \epsilon_\alpha^2 = \frac{1}{2D} \sum_{\alpha=1}^D (z_\alpha - y_\alpha)^2, \tag{2}$$

where the error term is $\epsilon_\alpha = z_\alpha - y_\alpha$. Throughout this work we will assume $y_\alpha = O(1)$. Under gradient descent the parameters evolve as:

$$\begin{aligned} \boldsymbol{\theta}_{t+1} &= \boldsymbol{\theta}_t - \frac{\eta}{D} \frac{\partial L_t}{\partial \boldsymbol{\theta}_t} \\ &= \boldsymbol{\theta}_t - \frac{\eta}{D} \sum_{\alpha=1}^D \epsilon_{\alpha,t} (\boldsymbol{\phi}(\mathbf{x}_\alpha) + \zeta \boldsymbol{\psi}(\mathbf{x}_\alpha) \boldsymbol{\theta}_t), \end{aligned} \tag{3}$$

where η is the learning rate and we explicitly introduced the time-dependence.

One quantity of particular interest is the NTK, $H_{\alpha\beta}$, which is defined by:

$$H_{\alpha\beta} \equiv H(\mathbf{x}_\alpha, \mathbf{x}_\beta) = \frac{1}{D} \sum_{\mu=1}^n \frac{\partial z(\mathbf{x}_\alpha)}{\partial \theta_\mu} \frac{\partial z(\mathbf{x}_\beta)}{\partial \theta_\mu}. \tag{4}$$

In vector notation $H_{\alpha\beta} = D^{-1} (\partial_{\boldsymbol{\theta}} z(\mathbf{x}_\alpha))^T \partial_{\boldsymbol{\theta}} z(\mathbf{x}_\beta)$. In the quadratic model the NTK is given by:

$$H_{\alpha\beta} = \frac{1}{D} \sum_{\mu=1}^n \phi_\mu^E(\mathbf{x}_\alpha, \boldsymbol{\theta}) \phi_\mu^E(\mathbf{x}_\beta, \boldsymbol{\theta}), \tag{5}$$

$$\phi_\mu^E(\mathbf{x}_\alpha, \boldsymbol{\theta}) = \phi_\mu(\mathbf{x}_\alpha) + \zeta \sum_{\nu=1}^n \psi_{\mu\nu}(\mathbf{x}_\alpha) \theta_\nu, \tag{6}$$

where $\phi^E(\mathbf{x}_\alpha, \boldsymbol{\theta})$ are the effective feature functions of the quadratic model. When $\zeta \neq 0$ they are θ -dependent and evolve under gradient descent, while when $\zeta = 0$ they reduce to the static feature functions ϕ . Correspondingly, the NTK only evolves when $\zeta \neq 0$.

We study two classes of quadratic models, the *pure quadratic model* which is defined to have $\phi = 0$ and the *quadratic model with bias* which is defined by the condition that $\boldsymbol{\psi}\phi = 0$, i.e. that ϕ is an eigenvector of $\boldsymbol{\psi}$ with eigenvalue 0.

Homogeneous nets. The quadratic models capture some aspects of representation learning in finite-width neural nets. Unfortunately, the quadratic approximation of neural nets, i.e. simply taking the output function $z(\mathbf{x}_\alpha, \boldsymbol{\theta})$ and Taylor expanding $\boldsymbol{\theta}$ around its value at initialization, breaks down when the weights $\boldsymbol{\theta}$ evolve significantly. When this happens we need to keep higher-order terms in the Taylor expansion to get a reasonable approximation to the full neural net⁶.

For this reason, we will also consider a more realistic set of toy models: neural nets with a single hidden layer and a scale-invariant activation function. These also scale quadratically in the weights, but will not generically be differentiable. They can therefore be thought of as *generalized quadratic models*. We will see how the methods which proved essential in studying quadratic models can be carried over to homogeneous nets with one hidden layer.

We study the following MLP:

$$z(x) = \frac{1}{\sqrt{n}} \mathbf{v}^T \sigma(\mathbf{u}x) = \frac{1}{\sqrt{n}} \sum_{\mu=1}^n v_\mu \sigma(u_\mu x), \tag{7}$$

where $x \in \mathbb{R}$, $\mathbf{u} \in \mathbb{R}^n$, $\mathbf{v} \in \mathbb{R}^n$. The activation function is:

$$\sigma(x) = a_+ x \mathbb{1}_{x \geq 0} + a_- x \mathbb{1}_{x < 0}, \tag{8}$$

where $\mathbb{1}$ is the indicator function. We assume that $0 \leq a_- \leq a_+$. The ReLU function corresponds to $(a_-, a_+) = (0, 1)$.

The NTK in this model is⁷:

$$\begin{aligned} H &= \frac{1}{n} \left(\sigma(\mathbf{u}x)^2 + (\mathbf{v} \circ \sigma'(\mathbf{u}x))^2 \right) \\ &= \frac{1}{n} \sum_{\mu=1}^n \left(\sigma^2(u_\mu x) + (v_\mu \sigma'(u_\mu x))^2 \right), \end{aligned} \tag{10}$$

where \circ is the Hadamard product. Finally, we need that the weight norm is simply given by $\boldsymbol{\theta}^2 = \mathbf{u}^2 + \mathbf{v}^2$.

⁶ The breakdown of the quadratic approximation is discussed in Roberts *et al* (2022), Agarwala *et al* (2022), Zhu *et al* (2022).

⁷ The derivative of the activation function is not defined at the origin, but for definiteness we can define:

$$\sigma'(0) \equiv \frac{1}{2} (a_+ + a_-). \tag{9}$$

3.2. Summary of results

Here we will summarize the main results. Table 1⁸ gives the primary analytic results of our paper, where we summarize our sufficiency conditions for the existence of the catapult phase using a toy dataset. The toy dataset is defined to have a single, one-dimensional datapoint, $x = 1$, with label $y = 0$. The proofs are given in appendix B. Although this dataset is too simple to model realistic datasets, it does give qualitatively good insight into the dynamics of these models at large learning rates. The generalization of these analytic results to more generic datasets is given in appendix C. Below, we will briefly comment on the results derived.

- Our results in table 1 are rigorously derived and significantly extend the results of Lewkowycz *et al* (2020), Zhu *et al* (2022). These conditions are in general sufficient, but not necessary, for the existence of the catapult phase. Our conditions are sufficient because they are derived by imposing that the weight norm θ_t^2 decays monotonically during training. In quadratic models and homogeneous MLPs, the loss can only diverge if some of the weights θ_μ diverge. Therefore, if θ_t^2 is bounded the loss must be finite. On the other hand, our conditions are not necessary because imposing that θ_t^2 monotonically decreases is stronger than only requiring that θ_t^2 is finite. We find numerically that the catapult phase generically does exist beyond the range we have derived.
- Given that our conditions were derived by requiring that the weight norm decreases monotonically under gradient descent, we can restate our results as follows: large learning rates which satisfy the conditions listed in table 1 lead to an implicit L_2 regularization of the weight norm θ_t^2 , i.e. these conditions imply that the weight norm θ_t^2 decreases monotonically during training, acting as an implicit L_2 -type regularization, even though no explicit regularization term is included in the loss function. This result, along with the fact that the catapult phase lands on flatter minima Keskar *et al* (2016), Lewkowycz *et al* (2020)⁹, may help explain why models in the catapult phase generalize better than models in the lazy phase¹⁰.
- The analysis of ReLU nets is more subtle than the analysis of general homogeneous nets. If we take our bounds for homogeneous nets and set $a_- = 0$ and $a_+ = 1$, the allowed range for η shrinks to zero size¹¹. However, by using special properties of the ReLU function, we can use a different argument to show that a two-layer ReLU net with one, $1d$ datapoint converges for $\eta H_0 \leq 4$. This agrees with the bound found in Zhu *et al* (2022) for the quadratic approximation of the same ReLU net with one datapoint of generic dimension.

On the other hand, in practice we know that the catapult phase for ReLU nets exists for larger learning rates, up to $\eta H_0 \lesssim 12$ Lewkowycz *et al* (2020). This presents

⁸ In particular, the derived learning rate bounds ensure that θ_t^2 decays during training, which can be viewed as an implicit L_2 regularization induced by the dynamics of gradient descent at large learning rates.

⁹ In the two-layer linear net with $1d$ input, the NTK is proportional to the weight norm (Lewkowycz *et al* 2020), so a decrease in the weight norm implies that the minima is flatter.

¹⁰ In general however, flatness is not necessarily correlated with better generalization behavior (Dinh *et al* 2017, Kaur *et al* 2022).

¹¹ To see this from the third row of table 1, we use that in the ReLU net with one datapoint on average only half of the weights at $t = 0$ contribute to the NTK. Therefore, at $t = 0$ the weight norm is twice the NTK on average and the allowed range disappears.

Table 1. Rigorous guarantees for the existence of the catapult phase in different models trained on the toy dataset $(x, y) = (1, 0)$.

Model	z	H	Derived catapult phase range for η
Pure quadratic	$\frac{\zeta}{2} \boldsymbol{\theta}^T \boldsymbol{\psi} \boldsymbol{\theta}$	$\zeta^2 \boldsymbol{\theta}^T \boldsymbol{\psi}^2 \boldsymbol{\theta}$	$(\frac{2}{H_0}, \frac{4}{\zeta^2 \boldsymbol{\theta}_0^2 \lambda_{\max}(\boldsymbol{\psi}^2)})$
Quadratic with biases	$\boldsymbol{\phi}^T \boldsymbol{\theta} + \frac{\zeta}{2} \boldsymbol{\theta}^T \boldsymbol{\psi} \boldsymbol{\theta}, \boldsymbol{\psi} \boldsymbol{\phi} = 0$	$\boldsymbol{\phi}^2 + \zeta^2 \boldsymbol{\theta}^T \boldsymbol{\psi}^2 \boldsymbol{\theta}$	$(\frac{2}{H_0}, \frac{4}{2\boldsymbol{\phi}^2 + \zeta^2 \lambda_{\max}(\boldsymbol{\psi}^2) (\boldsymbol{\theta}_0^2 + \frac{(\boldsymbol{\phi}^T \boldsymbol{\theta}_0)^2}{\boldsymbol{\phi}^2})})$
Homogeneous nets	$\frac{1}{\sqrt{n}} \mathbf{v}^T \sigma(\mathbf{u}x)$ $\sigma = a_+ x \mathbb{1}_{x \geq 0} + a_- x \mathbb{1}_{x \leq 0}$ $a_- > 0$	$\frac{1}{n} (\sigma(\mathbf{u}x)^2 + (\mathbf{v} \circ \sigma'(\mathbf{u}x))^2)$	$(\frac{2}{H_0}, \frac{4n}{a_+^2 (\mathbf{u}_0^2 + \mathbf{v}_0^2)})$
ReLU nets	$\frac{1}{\sqrt{n}} \mathbf{v}^T \sigma(\mathbf{u}x)$ $\sigma = a_+ x \mathbb{1}_{x \geq 0} + a_- x \mathbb{1}_{x \leq 0}$ $(a_-, a_+) = (0, 1)$	$\frac{1}{n} (\sigma(\mathbf{u}x)^2 + (\mathbf{v} \circ \sigma'(\mathbf{u}x))^2)$	$(\frac{2}{H_0}, \frac{4}{H_0})$

an interesting tension and suggests that the two-layer ReLU net may have characteristically different behaviors depending on the sign of $4 - \eta H_0$. We observe empirically that there is a qualitative difference: for $\eta H_0 \lesssim 4$ the weight norm θ^2 decreases during training, while for $\eta H_0 \gtrsim 4$ we find that θ^2 receives positive updates during training. In particular, if we take ηH_0 large enough, then the final weight norm can be greater than its value at initialization, $\theta_\infty^2 > \theta_0^2$. Despite this increase in θ^2 , models trained with $\eta H_0 \gtrsim 4$ still generalize well (Lewkowycz *et al* 2020). We conjecture that this is the case because the activation map of ReLU nets becomes sparse in the catapult phase.

3.3. Sketch of a proof

Here we will sketch the proof for the results in table 1 and leave the details for appendix B. In both two-layer, homogeneous nets and the pure quadratic model, the output z is a homogeneous function of θ with scaling weight two¹². Restoring the explicit weight dependence we have:

$$z(\mathbf{x}_\alpha, \lambda \theta) = \lambda^2 z(\mathbf{x}_\alpha, \theta). \tag{11}$$

This homogeneity property simplifies the update equation for θ_t^2 and implies that the NTK is also a homogeneous function of θ_t with weight two. Both of these properties are important in deriving the results of table 1.

If we train a net with the property (11) on the toy dataset $(x, y) = (1, 0)$ we find:

$$\theta_{t+1}^2 = \theta_t^2 + \eta z_t^2 (\eta H_t - 4). \tag{12}$$

From (12) we observe that the value $\eta H_t = 4$ is special: if $\eta H_t < 4$ then the weight norm decreases, $\theta_{t+1}^2 < \theta_t^2$, while for $\eta H_t > 4$ it increases. Therefore, if we can guarantee that $\eta H_t < 4$ for all t , then (12) implies that the weight norm monotonically decreases. This condition is sufficient, but not necessary, to ensure the loss does not diverge. It is also clear that imposing $\eta H_0 < 4$ does not guarantee convergence: it ensures that the first step of gradient descent decreases θ_t^2 , but H_t can increase during training and cause the model to diverge.

Therefore, to ensure convergence, we want to bound the NTK H_t throughout training. Specifically, we will bound H_t in terms of the weight norm θ_t^2 :

$$H_t \leq C \theta_t^2, \tag{13}$$

where C is some constant, positive number¹³. We will show that the bound (13) holds in the pure quadratic model and for two-layer, homogeneous MLPs in appendix B. The bound (13) allows us to ensure that H_t cannot become too large during training and cause the weight norm to get positive updates. Here it is important that the homogeneity property (11) implies H_t also has scaling weight two in θ_t in order that both sides of (13) have the same scaling in θ_t .

¹² A slightly modified argument is needed for the quadratic model with bias, see appendix B.

¹³ In principle, one can also consider bounding the NTK in terms of a positive, monotonic function of θ_t^2 , but we will find the above bound sufficient for the models we study.

Finally, if the inequality (13) holds, then θ_t^2 is a monotonically decreasing quantity when:

$$\eta < \frac{4}{C\theta_{t=0}^2}. \tag{14}$$

The argument is simple, if (14) holds then the first step of gradient descent decreases the weight norm:

$$\begin{aligned} \theta_{t=1}^2 - \theta_{t=0}^2 &= \eta z_{t=0}^2 (\eta H_{t=0} - 4) \\ &\leq \eta z_{t=0}^2 (\eta C \theta_{t=0}^2 - 4) < 0. \end{aligned} \tag{15}$$

To get the first inequality we plugged in (13) and to get the second inequality we used (14). Equation (14) guarantees that the weight norm decreases monotonically during training, which acts as an effective or implicit L_2 -type regularization. We emphasize that although no explicit regularization term is present, the dynamical effect of large learning rates in this regime suppresses weight growth similarly to L_2 regularization.

Now after a second step of gradient descent we have:

$$\begin{aligned} \theta_{t=2}^2 - \theta_{t=1}^2 &= \eta z_{t=1}^2 (\eta H_{t=1} - 4) \\ &\leq \eta z_{t=1}^2 (\eta C \theta_{t=1}^2 - 4) \\ &< \eta z_{t=1}^2 (\eta C \theta_{t=0}^2 - 4) < 0. \end{aligned} \tag{16}$$

Once again, to get the first inequality we used (13). To get the second inequality we used (15), or that the first step decreased θ_t^2 . The final inequality follows from condition (14). It is now clear that we can extend this argument to all t using an inductive proof, which is given in appendix B.

4. Experiments

Here we will present experimental results for quadratic models and two-layer, homogeneous nets. Notably, the experimental choices are for clarity and numerical stability. They are not essential to observe catapult dynamics. In the following we find qualitatively similar behavior under a wide range of settings. Further details and experiments, including extensions to more generic quadratic models and datasets, can be found in appendix D.

4.1. Linear meta-feature function

In general, the meta-feature function ψ is an arbitrary function of \mathbf{x} . In this section we study the pure quadratic model with a linear, meta-feature function:

$$\psi_{\mu\nu}^{\text{lin}}(\mathbf{x}_\alpha) = \sum_{i=1}^d x_{\alpha,i} W_{\mu\nu}^i. \tag{17}$$

This is arguably the simplest class of meta-feature functions to consider. The tensor $W_{\mu\nu}^i$ is symmetric under $\mu \leftrightarrow \nu$ and can be diagonalized for each i :

$$W_{\mu\nu}^i = \sum_{\sigma=1}^n \lambda_{\sigma}^i q_{\sigma\mu}^i q_{\sigma\nu}^i. \quad (18)$$

Here for each, fixed i the matrix $q_{\sigma\mu}^i$ is an orthogonal matrix.

We will train the pure quadratic model on the toy dataset $(x, y) = (1, 0)$ ¹⁴. We take the number of weights to be $n = 1000$ and set $\zeta^2 = 2/n$. This setting serves as a form of NTK normalization. It ensures that the effective feature functions and NTK remain well-behaved in the large-width limit, similar to the $1/\sqrt{d}$ scaling used in Transformer attention mechanisms (Vaswani *et al* 2017). See also the discussions in Roberts *et al* (2022) for related justifications. For each i we take $q_{\mu\nu}^i$ to be a random orthogonal matrix. In addition, for each i we split the eigenvalues λ_{σ}^i into two sets of identical size, corresponding to the positive and negative eigenvalues. We draw the positive eigenvalues from $\mathcal{U}([1, 2])$ and take the negative eigenvalues to be exactly -1 times the positive eigenvalues. We impose that the eigenvalues come in positive/negative pairs so that $\mathbb{E}[z_0] = 0$. We also choose the range $[1, 2]$ because the two-layer, linear MLP with one datapoint, $x = 1$, corresponds to a quadratic model with eigenvalues ± 1 Agarwala *et al* (2022), Zhu *et al* (2022). We can then think of this quadratic model as a simple deformation of the linear MLP. The results for this set-up are shown in figure 1. We see that this quadratic model undergoes catapult dynamics and our bounds correctly predict the model does converge for a finite window above $\eta H_0 = 2$.

4.2. Generic homogeneous net

We also perform a similar experiment for homogeneous nets with one hidden layer of width 1024. To be concrete, we consider a net with $a_+ = 1$ and $a_- = 1/2$. We train this model on the toy dataset $(x, y) = (1, 0)$. The results are shown in figure 2. When $\eta H_0 \lesssim 4$, the results are qualitatively the same as for the pure quadratic model. However, one new feature is this model converges for larger learning rates. In particular, for $\eta H_0 = 4.5$ we see that θ_t^2 initially grows during the catapult phase, before decreasing to a small value. This growth in θ_t^2 occurs because the product $\eta H_t > 4$ when $z_t = O(\sqrt{n})$. From the update equation for θ_t^2 (12), this means that θ_t^2 receives a large positive correction during training. However, eventually ηH_t is pushed below 4 and the model converges.

This result is consistent with our analytic results summarized in table 1. When η obeys the bounds summarized there, the weight norm θ_t^2 must decrease monotonically. Since $\eta H_0 = 4.5$ exceeds our bounds, there is no guarantee θ_t^2 must decrease. We can also observe in figure 2(e) that there is a transition in $\theta_{\infty}^2/\theta_0^2$ when $\eta H_0 \approx 4$. When $2 \lesssim \eta H_0 \lesssim 4$ the final weight norm is smaller than its value at initialization. As we

¹⁴ The toy setting $(x, y) = (1, 0)$ is a simplification that offers qualitatively meaningful insights into the dynamics of the model under large learning rates. This allows us to analytically demonstrate the existence of the catapult phase in a clean and transparent manner. Though not conducting traditional representation learning, it reveals that the model can still exhibit nontrivial training dynamics. In appendix D.1 we will consider experiments for both the pure quadratic model and the quadratic model with bias trained on random datasets and teacher-student set-ups.

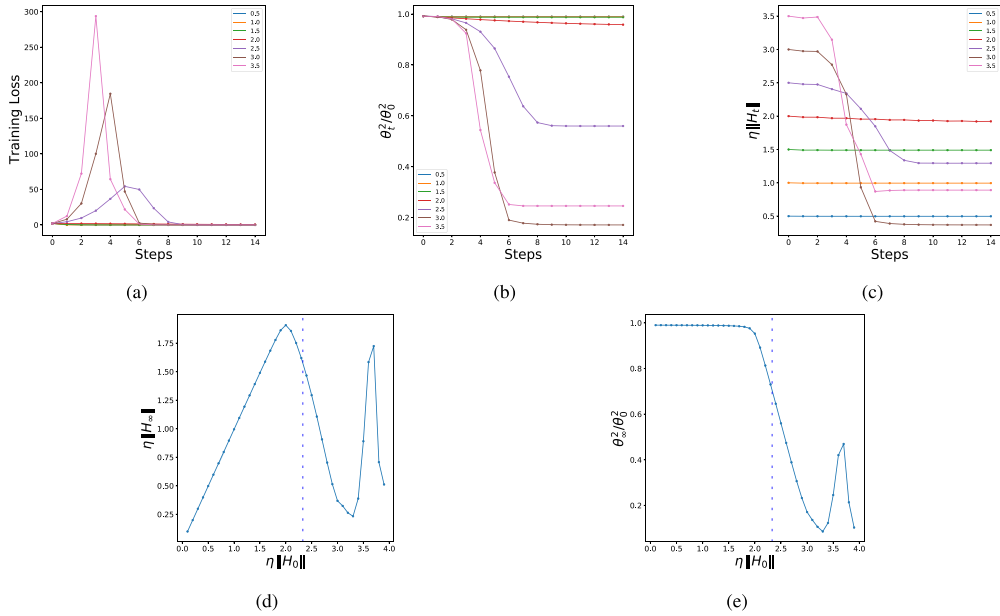


Figure 1. Results for the pure quadratic model with a linear meta-feature function trained on the toy dataset $(x, y) = (1, 0)$. (a)–(c) give the evolution of the loss L_t , the weight norm θ_t^2 , and ηH_t as a function of time, respectively. The different colors in (a)–(c) correspond to different choices of ηH_0 . In (a) we see that the loss experiences a transient phase of exponential growth before decreasing down to 0. In (b)–(c) we see that in the catapult phase both θ_t^2 and ηH_t decrease significantly. In (d)–(e) we show the final value of ηH_t and θ_t^2 as a function of ηH_0 . The dashed vertical lines correspond to the upper bound in the first row of table 1. We correctly predict that the model does have a catapult phase and empirically observe the model converges up to $\eta H_0 = 4$.

approach $\eta H_0 \rightarrow 4$ this trend reverses and there is a peak at $\eta H_0 = 4$, before θ_t^2 decreases again. To make this change in the evolution of θ_t^2 as a function of η more pronounced, we will study ReLU nets.

4.3. ReLU net

In this section we will study the behavior of fully-connected ReLU nets trained on a two-class version of MNIST. We will train the model to distinguish images of ‘0’ and ‘1’, which are assigned labels -1 and 1 , respectively. The training set has size 128 and the test set has size 2115. Our analytic bounds for ReLU nets is restricted to $1d$ input data, so here the results will be purely empirical¹⁵.

We train a two-layer net with width 1024 using full-batch gradient descent. The model is trained for each learning rate until the change in the training loss is $< 10^{-8}$. The results are shown in figure 3. The new feature, in comparison to the previous experiments, is that the final weight norm θ_∞^2 can be greater than the value at initialization,

¹⁵ A comparison between analytic and numeric for one-dimensional data and ReLU nets will be given in appendix D.2.

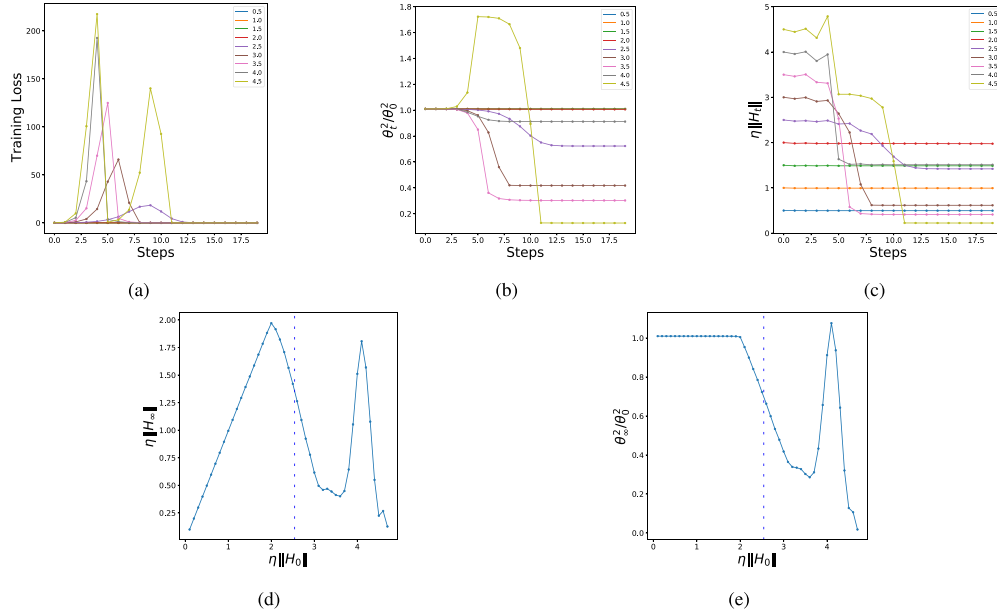


Figure 2. Results for the two-layer, homogeneous net with $a_+ = 1$ and $a_- = 1/2$. (a)–(c) give the evolution of the loss L_t , the weight norm θ_t^2 , and ηH_t as a function of time, respectively. The different colors in (a)–(c) correspond to different choices of ηH_0 . (d) and (e) give the final values for ηH_t and θ_t^2 after the model has converged as a function of ηH_0 . The hidden layer has width 1024. The vertical dashed line corresponds to the upper bound in the third row of table 1. We again correctly predict that the model converges for super-critical learning rates to the left of this line.

see figures 3(b) and (e). Despite the increase in θ_t^2 the model still generalizes well, see figure 3(g). We conjecture that the model is able to generalize well, despite the increase in θ_t^2 , because the activation map becomes sparse in the catapult phase, see figure 3(f). The sparsity of a given layer is the fraction of nodes in that layer which are zero after acting with the ReLU function for a single input. To get figure 3(f) we averaged the sparsity of the activation map over all inputs.

Finally, we can note that the final value of $\eta \lambda_{\max}(H_{\alpha\beta,t})$, which measures the sharpness of the final minima, first decreases in the catapult phase before increasing again to the EOS, $\eta \lambda_{\max}(H_t) = 2$, see figure 3(d). This is consistent with existing work on the EOS phenomena for the top eigenvalue of the NTK (Cohen *et al* 2021, Agarwala *et al* 2022).

5. Conclusions

In this work we have studied the properties of simple machine learning models, the quadratic model and two-layer, homogeneous nets, for super-critical learning rates. We have proven that these models can converge beyond the stability threshold of the corresponding linearized model. In the process, we have shown that training these models for




Figure 3. Results for the two-layer ReLU net trained on a two-class version of MNIST. The hidden layer has width 1024. (a)–(c) give the early-time evolution of the loss, the weight norm, and $\eta\lambda_{\max}(H_{\alpha\beta,t})$, respectively. The different lines in (a)–(c) correspond to different choices of $\eta\lambda_{\max}(H_{\alpha\beta,0})$. (d)–(g) give the final values for $\eta\lambda_{\max}(H_{\alpha\beta,t})$, θ_t^2 , the sparsity of the activation map, and the generalization gap, respectively, for the converged model as a function of $\eta\lambda_{\max}(H_{\alpha\beta,0})$. From (f) we observe that increasing the learning rate in the catapult phase promotes sparsity.

super-critical learning rates can lead to an implicit L_2 regularization where the weight norm θ^2 decreases significantly during training. While our theoretical analysis mainly rely on weight norm decay, we emphasize that the catapult phase exhibits richer behavior beyond this. For instance, we have also empirically studied ReLU nets for even larger learning rates, $\eta H_0 \gtrsim 4$, and showed that training models in this range leads to an increase in both the sparsity and the weight norm. The ReLU nets can generalize well even when the weight norm increases during training, likely due to emergent sparsity in the activation map. These results, both analytical and empirical, shed new light on the implicit biases of gradient descent at large learning rates.

There are many interesting open questions to consider. One is to understand how to weaken our bounds, which in general can miss large portions of the catapult regime.

The reason our bounds can be too strong is because they are designed to hold for all possible configurations of θ . In particular, they hold when inequalities of the form (13) are saturated. To derive weaker bounds we need a better understanding on how the weights evolve on average for super-critical learning rates. In addition, proving that ReLU nets can converge for $4 \leq \eta H_0 \lesssim 12$ remains an open question.

Another question is: why do models trained in the catapult phase generalize well? Our empirical and analytic results align with the expectation that large learning rates can lead to implicit regularization (Li *et al* 2019) which improves model performance. It would be interesting to understand what features of the catapult mechanism lead to a small generalization error for more generic, wide neural nets.

Acknowledgment

We thank Yasaman Bahri, Daniel Roberts, Yi-Hsien Du, and Han Zheng for discussions. We additionally thank Daniel Roberts for comments on the draft. D M is supported by the NSF Grant PHY-2014071. JL is supported in part by the University of Pittsburgh, School of Computing and Information, Department of Computer Science, Pitt Cyber, PQI Community Collaboration Awards, John C. Mascaró Faculty Scholar in Sustainability, NASA under award number 80NSSC25M7057, and Fluor Marine Propulsion LLC (U.S. Naval Nuclear Laboratory) under award number 140449-R08. This research used resources of the Oak Ridge Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725. JL is also supported in part by International Business Machines (IBM) Quantum through the Chicago Quantum Exchange, and the Pritzker School of Molecular Engineering at the University of Chicago through AFOSR MURI (FA9550-21-1-0209).

Appendix A. Early time dynamics

In this appendix, we study the early time dynamics of quadratic models and two-layer homogeneous MLPs. The goal of this appendix is to show that when $\eta \lambda_{\max}(H_{\alpha\beta,0}) > 2$, the loss grows exponentially quickly at early times. We will then estimate the scale at which perturbation theory breaks down. This appendix builds off of the theoretical analysis done in Lewkowycz *et al* (2020). We will also study the evolution of the weight norm and show that it only receives large corrections near the peak of the catapult phase.

The analysis in this appendix will be important in proving that models trained on generic datasets can converge for super-critical learning rates but will not be necessary for proving convergence on the toy dataset $(x, y) = (1, 0)$.

A.1. Catapult phase in the quadratic model

Here we will study the pure quadratic model and the quadratic model with bias. The general quadratic model is:

$$z(\mathbf{x}_\alpha) = \boldsymbol{\theta}^\top \boldsymbol{\phi}_\alpha + \frac{\zeta}{2} \boldsymbol{\theta}^\top \boldsymbol{\psi}_\alpha \boldsymbol{\theta}, \quad (19)$$

where $\boldsymbol{\theta} \in \mathbb{R}^n$, $\zeta^2 = O(1/n)$ and $n \gg 1$. We also use the notation, $\boldsymbol{\phi}_\alpha \equiv \boldsymbol{\phi}(\mathbf{x}_\alpha)$ and $\boldsymbol{\psi}_\alpha \equiv \boldsymbol{\psi}(\mathbf{x}_\alpha)$. We train the model with MSE:

$$L = \frac{1}{2D} \sum_{\alpha=1}^D \epsilon_\alpha^2 = \frac{1}{2D} \sum_{\alpha=1}^D (z_\alpha - y_\alpha)^2. \tag{20}$$

Here $y_\alpha \in \mathbb{R}$ and $\mathbf{x}_\alpha \in \mathbb{R}^d$. The weights evolve as:

$$\boldsymbol{\theta}_{t+1} = \boldsymbol{\theta}_t - \frac{\eta}{D} \sum_{\alpha=1}^D \epsilon_{\alpha,t} (\boldsymbol{\phi}_\alpha + \zeta \boldsymbol{\psi}_\alpha \boldsymbol{\theta}_t). \tag{21}$$

To study the quadratic model with bias we impose the condition,

$$\boldsymbol{\psi}_\alpha \boldsymbol{\phi}_\beta = 0, \quad \forall \alpha, \beta. \tag{22}$$

To study the pure quadratic model we can set $\boldsymbol{\phi} = 0$. The NTK in this model is

$$H_{\alpha\beta} = \frac{1}{D} (\boldsymbol{\phi}_\alpha^T \boldsymbol{\phi}_\beta + \zeta^2 \boldsymbol{\theta}^T \boldsymbol{\psi}_\alpha \boldsymbol{\psi}_\beta \boldsymbol{\theta}). \tag{23}$$

We draw the initial weights from a normal distribution with zero mean and unit variance, $\boldsymbol{\theta} \sim \mathcal{N}(0, \mathbb{I}_{n \times n})$. Then at initialization we have:

$$\mathbb{E}[z_{\alpha,0}] = \frac{\zeta}{2} \text{Tr}(\boldsymbol{\psi}_\alpha), \tag{24}$$

$$\mathbb{E}[z_{\alpha,0} z_{\beta,0}] = \boldsymbol{\phi}_\alpha^T \boldsymbol{\phi}_\beta + \frac{\zeta^2}{2} \text{Tr}(\boldsymbol{\psi}_\alpha \boldsymbol{\psi}_\beta) + \mathbb{E}[z_\alpha] \mathbb{E}[z_\beta], \tag{25}$$

$$\mathbb{E}[H_{\alpha\beta,0}] = \frac{1}{D} (\boldsymbol{\phi}_\alpha^T \boldsymbol{\phi}_\beta + \zeta^2 \text{Tr}(\boldsymbol{\psi}_\alpha \boldsymbol{\psi}_\beta)), \tag{26}$$

$$\mathbb{E}[\boldsymbol{\theta}_0^2] = n. \tag{27}$$

We will choose the (meta-)feature functions such that $z_{\alpha,0}, H_{\alpha\beta,0} = O(\zeta^0)$. To impose this, we will assume that the eigenvalues of the meta-feature functions are order-one, $\lambda_i(\boldsymbol{\psi}_\alpha) = O(\zeta^0)$ and come in approximately positive/negative pairs. In this case:

$$\zeta \text{Tr}(\boldsymbol{\psi}_\alpha) = \zeta \sum_{a=1}^n \lambda_a(\boldsymbol{\psi}_\alpha) \ll 1, \tag{28}$$

$$\zeta^2 \text{Tr}(\boldsymbol{\psi}_\alpha^2) = \zeta^2 \sum_{a=1}^n \lambda_a^2(\boldsymbol{\psi}_\alpha) = O(\zeta^0). \tag{29}$$

Finally, we impose $\boldsymbol{\phi}_\alpha^2 = O(\zeta^0)$.

The update equations for ϵ_α and $H_{\alpha\beta}$ are:

$$\epsilon_{\alpha,t+1} = \sum_{\beta=1}^D (\delta_{\alpha\beta} - \eta H_{\alpha\beta,t}) \epsilon_{\beta,t} + \frac{\eta^2 \zeta^3}{2D^2} \sum_{\beta,\gamma=1}^D \epsilon_{\beta,t} \epsilon_{\gamma,t} \boldsymbol{\theta}_t^T \boldsymbol{\psi}_\beta \boldsymbol{\psi}_\alpha \boldsymbol{\psi}_\gamma \boldsymbol{\theta}_t, \tag{30}$$

$$\begin{aligned}
 H_{\alpha\beta,t+1} &= H_{\alpha\beta,t} - \frac{\eta\zeta^3}{D} \sum_{\gamma=1}^D [\epsilon_{\gamma,t} \boldsymbol{\theta}_t^T \boldsymbol{\psi}_\gamma \boldsymbol{\psi}_\alpha \boldsymbol{\psi}_\beta \boldsymbol{\theta}_t + (\alpha \leftrightarrow \beta)] \\
 &\quad + \frac{\eta^2\zeta^4}{D^2} \sum_{\gamma,\rho=1}^D \epsilon_{\gamma,t} \epsilon_{\rho,t} \boldsymbol{\theta}_t^T \boldsymbol{\psi}_\gamma \boldsymbol{\psi}_\alpha \boldsymbol{\psi}_\beta \boldsymbol{\psi}_\rho \boldsymbol{\theta}_t,
 \end{aligned} \tag{31}$$

where $\delta_{\alpha\beta}$ is the Kronecker delta function. Next we will show that when $\eta\lambda_{\max}(H_{\alpha\beta,0}) > 2$ and $\zeta \ll 1$ the loss increases exponentially quickly at early times and that small ζ -perturbation theory breaks down when $t = O(\log(\zeta^{-1}))$. First we will show that when $t \ll \log(\zeta^{-1})$ we can approximate the update equations by:

$$\epsilon_{\alpha,t+1} \approx \sum_{\beta=1}^D (\delta_{\alpha\beta} - \eta H_{\alpha\beta,t}) \epsilon_{\beta,t}, \tag{32}$$

$$H_{\alpha\beta,t+1} \approx H_{\alpha\beta,t}. \tag{33}$$

That is, we are dropping terms in (30) and (31) which are explicitly suppressed in ζ . When this approximation holds the NTK $H_{\alpha\beta,t}$ is constant and the error $\epsilon_{\alpha,t}$ increases fastest in the direction parallel to the top eigenvector of the NTK. To be more explicit, we write

$$\epsilon_{\alpha,t} = \sum_{i=1}^D c_t^i e_{\alpha,t}^i, \tag{34}$$

where $e_{\alpha,t}^i$ is the i th eigenvector of the NTK at time t . Assuming (32) and (33) hold, the eigenvectors are approximately constant and the coefficients c_t^i are given by:

$$c_t^i = (1 - \eta\lambda_i(H_{\alpha\beta,0}))^T c_{t=0}^i, \tag{35}$$

so the error grows fastest in the direction of the top eigenvector of the NTK.

To show that (32) and (33) are valid at early times we will bound the sub-leading terms in (30) and (31). We will first bound the size of the ζ^3 term in (30) using Cauchy-Schwarz inequalities:

$$\begin{aligned}
 \left| \frac{\eta^2\zeta^3}{2D^2} \sum_{\beta,\gamma=1}^D \epsilon_{\beta,t} \epsilon_{\gamma,t} \boldsymbol{\theta}_t^T \boldsymbol{\psi}_\beta \boldsymbol{\psi}_\alpha \boldsymbol{\psi}_\gamma \boldsymbol{\theta}_t \right| &\leq \frac{\eta^2\zeta^3}{2D^2} \sum_{\beta,\gamma=1}^D |\epsilon_{\beta,t} \epsilon_{\gamma,t} \boldsymbol{\theta}_t^T \boldsymbol{\psi}_\beta \boldsymbol{\psi}_\alpha \boldsymbol{\psi}_\gamma \boldsymbol{\theta}_t| \\
 &\leq \frac{\eta^2\zeta^3}{2D^2} \boldsymbol{\theta}_t^2 \sqrt{\lambda_{\max}(\boldsymbol{\psi}_\alpha^2)} \sum_{\beta,\gamma=1}^D |\epsilon_{\beta,t} \epsilon_{\gamma,t}| \sqrt{\lambda_{\max}(\boldsymbol{\psi}_\beta^2) \lambda_{\max}(\boldsymbol{\psi}_\gamma^2)} \\
 &\leq \frac{\eta^2\zeta^3}{2D^2} \boldsymbol{\theta}_t^2 \sqrt{\lambda_{\max}(\boldsymbol{\psi}_\alpha^2)} \sum_{\beta=1}^D \epsilon_{\beta,t}^2 \sum_{\gamma=1}^D \lambda_{\max}(\boldsymbol{\psi}_\gamma^2) \\
 &= O\left(\zeta^3 \boldsymbol{\theta}_t^2 D^{-1} \sum_{\beta=1}^D \epsilon_{\beta,t}^2\right).
 \end{aligned} \tag{36}$$

To obtain the last line we used the assumption that the eigenvalues of the meta-feature function are all $O(1)$.

We can similarly bound the size of the ζ^3 term in (31) using Cauchy–Schwarz inequalities:

$$\begin{aligned} \left| \frac{\eta\zeta^3}{D} \sum_{\gamma=1}^D \epsilon_{\gamma,t} \boldsymbol{\theta}_t^T \boldsymbol{\psi}_\gamma \boldsymbol{\psi}_\alpha \boldsymbol{\psi}_\beta \boldsymbol{\theta}_t \right| &\leq \frac{\eta\zeta^3}{D} \boldsymbol{\theta}_t^2 \sqrt{\lambda_{\max}(\boldsymbol{\psi}_\alpha^2) \lambda_{\max}(\boldsymbol{\psi}_\beta^2)} \sum_{\gamma=1}^D |\epsilon_{\gamma,t}| \sqrt{\lambda_{\max}(\boldsymbol{\psi}_\gamma^2)} \\ &\leq \frac{\eta\zeta^3}{D} \boldsymbol{\theta}_t^2 \sqrt{\lambda_{\max}(\boldsymbol{\psi}_\alpha^2) \lambda_{\max}(\boldsymbol{\psi}_\beta^2)} \sqrt{\sum_{\gamma=1}^D \epsilon_{\gamma,t}^2} \sqrt{\sum_{\rho=1}^D \lambda_{\max}(\boldsymbol{\psi}_\rho^2)} \\ &= O\left(\zeta^3 \boldsymbol{\theta}_t^2 D^{-1/2} \sqrt{\sum_{\gamma=1}^D \epsilon_{\gamma,t}^2}\right). \end{aligned} \tag{37}$$

Finally, we can bound the ζ^4 term of (31):

$$\begin{aligned} \left| \frac{\eta^2\zeta^4}{D^2} \sum_{\gamma,\rho=1}^D \epsilon_{\gamma,t} \epsilon_{\rho,t} \boldsymbol{\theta}_t^T \boldsymbol{\psi}_\gamma \boldsymbol{\psi}_\alpha \boldsymbol{\psi}_\beta \boldsymbol{\psi}_\rho \boldsymbol{\theta}_t \right| &\leq \frac{\eta^2\zeta^4}{D^2} \boldsymbol{\theta}_t^2 \sqrt{\lambda_{\max}(\boldsymbol{\psi}_\alpha^2) \lambda_{\max}(\boldsymbol{\psi}_\beta^2)} \sum_{\gamma,\rho=1}^D |\epsilon_{\gamma,t} \epsilon_{\rho,t}| \\ &\quad \times \sqrt{\lambda_{\max}(\boldsymbol{\psi}_\gamma^2) \lambda_{\max}(\boldsymbol{\psi}_\rho^2)} \\ &\leq \frac{\eta^2\zeta^4}{D^2} \boldsymbol{\theta}_t^2 \sqrt{\lambda_{\max}(\boldsymbol{\psi}_\alpha^2) \lambda_{\max}(\boldsymbol{\psi}_\beta^2)} \sum_{\gamma=1}^D \epsilon_{\gamma,t}^2 \sum_{\rho=1}^D \lambda_{\max}(\boldsymbol{\psi}_\rho^2) \\ &= O\left(\zeta^4 \boldsymbol{\theta}_t^2 D^{-1} \sum_{\gamma} \epsilon_{\gamma,t}^2\right). \end{aligned} \tag{38}$$

Next, we can argue that (36)–(38) are all small at initialization when $\zeta \ll 1$. We will use that at $t = 0$:

$$\boldsymbol{\theta}_0^2 = O(\zeta^{-2} D^0), \tag{39}$$

$$|\epsilon_{\beta,0}^2| = O(1), \tag{40}$$

$$H_{\alpha\beta,0} = O(\zeta^0 D^{-1}). \tag{41}$$

If we use these estimates in (36)–(38), we find that at $t = 0$ these terms scale like:

$$\left| \frac{\eta^2\zeta^3}{2D^2} \sum_{\beta,\gamma=1}^D \epsilon_{\beta,0} \epsilon_{\gamma,0} \boldsymbol{\theta}_0^T \boldsymbol{\psi}_\beta \boldsymbol{\psi}_\alpha \boldsymbol{\psi}_\gamma \boldsymbol{\theta}_0 \right| = O(\zeta), \tag{42}$$

$$\left| \frac{\eta\zeta^3}{D} \sum_{\gamma=1}^D \epsilon_{\gamma,0} \boldsymbol{\theta}_0^T \boldsymbol{\psi}_\gamma \boldsymbol{\psi}_\alpha \boldsymbol{\psi}_\beta \boldsymbol{\theta}_0 \right| = O(\zeta), \tag{43}$$

$$\left| \frac{\eta^2 \zeta^4}{D^2} \sum_{\gamma, \rho=1}^D \epsilon_{\gamma,0} \epsilon_{\rho,0} \boldsymbol{\theta}_0^T \boldsymbol{\psi}_\gamma \boldsymbol{\psi}_\alpha \boldsymbol{\psi}_\beta \boldsymbol{\psi}_\rho \boldsymbol{\theta}_0 \right| = O(\zeta^2). \tag{44}$$

At initialization $\epsilon_{\alpha,0} = O(1)$, so the sub-leading term (42) is negligible if $\zeta \ll 1$. Similarly, at $t=0$ we have $H_{\alpha\beta,0} = O(D^{-1})$, so we can drop the sub-leading terms (43) and (44) if $\zeta D \ll 1$. Both conditions are satisfied if assume $\zeta \ll 1$ and D is generic.

This proves that, at $t=0$, we can ignore the sub-leading terms in the small ζ expansion. We can now determine at what scale perturbation theory breaks down by finding at what scale the sub-leading terms (36)–(38) are of the same order as the leading order terms. If we plug in the 0th-order solutions, (33)–(35) into the sub-leading terms (36)–(38), we find that perturbation theory breaks down when $c_t^{\max} = O(\zeta^{-1})$, which happens when $t = O(\log(\zeta^{-1}))$.

A.2. Weight norm in the quadratic model

In this section we will study the evolution of $\boldsymbol{\theta}^2$ and $\boldsymbol{\phi}_\alpha^T \boldsymbol{\theta}$ and show that small ζ -perturbation theory is valid when $t \ll \log(\zeta^{-1})$. Specifically, we will show:

$$\boldsymbol{\theta}_t^2 = \boldsymbol{\theta}_0^2 + O(1) \quad \text{for } t \ll \log(\zeta^{-1}). \tag{45}$$

The update equation for $\boldsymbol{\theta}_t^2$ is:

$$\begin{aligned} \boldsymbol{\theta}_{t+1}^2 &= \boldsymbol{\theta}_t^2 - \frac{2\eta}{D} \sum_{\alpha=1}^D (\epsilon_{\alpha,t} \boldsymbol{\phi}_\alpha^T \boldsymbol{\theta}_t + \zeta \epsilon_{\alpha,t} \boldsymbol{\theta}_t^T \boldsymbol{\psi}_\alpha \boldsymbol{\theta}_t) + \frac{\eta^2}{D} \sum_{\alpha, \beta=1}^D \epsilon_{\alpha,t} \epsilon_{\beta,t} H_{\alpha\beta,t} \\ &= \boldsymbol{\theta}_t^2 - \frac{4\eta}{D} \sum_{\alpha=1}^D \epsilon_{\alpha,t} z_{\alpha,t} + \frac{\eta^2}{D} \sum_{\alpha, \beta=1}^D \epsilon_{\alpha,t} \epsilon_{\beta,t} H_{\alpha\beta,t} + \frac{2\eta}{D} \sum_{\alpha=1}^D \epsilon_{\alpha,t} \boldsymbol{\phi}_\alpha^T \boldsymbol{\theta}_t. \end{aligned} \tag{46}$$

With the exception of the last term, we already know the time evolution of the right-hand side of (46). For early times, $t \ll \log(\zeta^{-1})$, the terms $\epsilon_{\alpha,t}$ and $z_{\alpha,t}$ grow exponentially quickly and the NTK $H_{\alpha\beta,t}$ is approximately constant, see (32)–(35). In addition, at initialization we have $\boldsymbol{\theta}_0^2 = O(\zeta^{-2})$. Therefore, these sub-leading terms in ζ become comparable to $\boldsymbol{\theta}_0^2$ when $t = O(\log(\zeta^{-1}))$. This is the same scale at which perturbation theory breaks down for the $z_{\alpha,t}$ and $H_{\alpha\beta,t}$ update equations, see the previous section.

In the pure quadratic model $\boldsymbol{\phi}_\alpha = 0$ and we could stop here. The above analysis proves that, in the pure quadratic model, we have:

$$\boldsymbol{\theta}_t^2 = \boldsymbol{\theta}_0^2 + O(1), \quad \text{for } t \ll \log(\zeta^{-1}). \tag{47}$$

To complete the proof for the quadratic model with bias we need to study the last term in (46). Using our assumption that $\boldsymbol{\psi}_\alpha \boldsymbol{\phi}_\beta = 0$, we find its update equation is:

$$\boldsymbol{\phi}_\alpha^T \boldsymbol{\theta}_{t+1} = \boldsymbol{\phi}_\alpha^T \boldsymbol{\theta}_t - \frac{\eta}{D} \sum_{\beta=1}^D \boldsymbol{\phi}_\alpha^T \boldsymbol{\phi}_\beta \epsilon_{\beta,t}, \tag{48}$$

whose solution we can find in closed form:

$$\phi_\alpha^T \theta_t = \phi_\alpha^T \theta_0 - \frac{\eta}{D} \sum_{\beta=1}^D \sum_{i=1}^{t-1} \phi_\alpha^T \phi_\beta \epsilon_{\beta,i}. \tag{49}$$

The time evolution of the last term in (46) is then:

$$\frac{2\eta}{D} \sum_{\alpha=1}^D \epsilon_{\alpha,t} \phi_\alpha^T \theta_t = \frac{2\eta}{D} \sum_{\alpha=1}^D \epsilon_{\alpha,t} \phi_\alpha^T \left(\theta_0 - \frac{\eta}{D} \sum_{\beta=1}^D \sum_{i=1}^{t-1} \phi_\beta \epsilon_{\beta,i} \right). \tag{50}$$

Note that this is an exact solution and does not require any assumptions on t .

We can now bound the size of the last term in (46):

$$\begin{aligned} \left| \frac{2\eta}{D} \sum_{\alpha=1}^D \epsilon_{\alpha,t} \phi_\alpha^T \theta_t \right| &= \frac{2\eta}{D} \left| \sum_{\alpha=1}^D \epsilon_{\alpha,t} \phi_\alpha^T \left(\theta_0 - \frac{\eta}{D} \sum_{\beta=1}^D \sum_{i=1}^{t-1} \phi_\beta \epsilon_{\beta,i} \right) \right| \\ &\leq \frac{2\eta}{D} \left(\left| \sum_{\alpha=1}^D \epsilon_{\alpha,t} \phi_\alpha^T \theta_0 \right| + \frac{\eta}{D} \sum_{\alpha,\beta=1}^D \sum_{i=1}^{t-1} |\epsilon_{\alpha,t} \phi_\alpha^T \phi_\beta \epsilon_{\beta,i}| \right) \\ &\leq \frac{2\eta}{D} \left(\sqrt{\theta_0^2 \lambda_{\max}(\phi_\alpha^T \phi_\beta^T) \sum_{\alpha=1}^D \epsilon_\alpha^2} + \frac{\eta}{D} (t-1) \lambda_{\max}(\phi_\alpha^T \phi_\beta) \sum_{\alpha=1}^D \epsilon_{\alpha,t}^2 \right). \end{aligned} \tag{51}$$

To go from the second to the third line we used Cauchy–Schwarz inequalities and also replaced $\epsilon_{\beta,i} \rightarrow \epsilon_{\beta,t}$ for each i . This gives a very weak, but valid, upper bound because the errors $\epsilon_{\beta,t}$ grow exponentially quickly when $\eta \lambda_{\max}(H_{\alpha\beta,0}) > 2$.

If we define our model such that $D^{-1} \lambda_{\max}(\phi_\alpha^T \phi_\beta) = O(1)$, then the first term in (51) becomes order $O(\zeta^{-2}) = O(n)$ when $t = O(\log(\zeta^{-2}))$. The second term is more non-trivial to analyze due to the explicit factor of t and we find it becomes order $O(\zeta^{-2})$ when $t = O(\log(\zeta^{-2}) - \log \log(\zeta^{-2}))$. If in (51) we had instead assumed:

$$\frac{\eta}{D} \sum_{\alpha,\beta=1}^D \sum_{i=1}^{t-1} |\epsilon_{\alpha,t} \phi_\alpha^T \phi_\beta \epsilon_{\beta,i}| \approx \frac{\eta}{D} \sum_{\alpha,\beta=1}^D |\epsilon_{\alpha,t} \phi_\alpha^T \phi_\beta \epsilon_{\beta,t}|, \tag{52}$$

then we would have found this term becomes of order $O(\zeta^{-2})$ for $t = O(\log(\zeta^{-1}))$. We think (52) is a more reasonable assumption because of the exponential growth of $\epsilon_{\beta,t}$. With either assumption, we find that the final term in (46) becomes $O(\zeta^{-2})$ in a time-scale that scales logarithmically in ζ^{-1} and that we can use perturbation theory if $t \ll \log(\zeta^{-1})$.

A.3. Homogeneous, two-layer MLP

Here we study the early-time dynamics of the homogeneous, two-layer MLP:

$$z_\alpha = \frac{1}{\sqrt{n}} \mathbf{v}^\top \sigma(\mathbf{U} \mathbf{x}_\alpha), \tag{53}$$

$$H_{\alpha\beta} = \frac{1}{nD} \left(\sigma(\mathbf{U} \mathbf{x}_\alpha)^\top \sigma(\mathbf{U} \mathbf{x}_\beta) + \mathbf{x}_\alpha^\top \mathbf{x}_\beta (\mathbf{v} \circ \sigma'(\mathbf{U} \mathbf{x}_\alpha))^\top (\mathbf{v} \circ \sigma'(\mathbf{U} \mathbf{x}_\beta)) \right), \tag{54}$$

where $\mathbf{v} \in \mathbb{R}^n$, $\mathbf{U} \in \mathbb{R}^{n \times d}$, $\mathbf{x}_\alpha \in \mathbb{R}^d$, $\alpha, \beta \in \{1, \dots, D\}$, and \circ is the Hadamard product. The scale-invariant activation function σ is defined in (8). We will choose an activation function with $0 \leq a_- \leq a_+$. The weights \mathbf{U} and \mathbf{v} evolve as:

$$\mathbf{v}_{t+1} = \mathbf{v}_t - \frac{\eta}{\sqrt{nD}} \sum_{\alpha=1}^D \epsilon_{\alpha,t} \sigma(\mathbf{U}_t \mathbf{x}_\alpha), \tag{55}$$

$$\mathbf{U}_{t+1} = \mathbf{U}_t - \frac{\eta}{\sqrt{nD}} \sum_{\alpha=1}^D \epsilon_{\alpha,t} (\mathbf{v}_t \circ \sigma'(\mathbf{U}_t \mathbf{x}_\alpha)) \mathbf{x}_\alpha^\top. \tag{56}$$

Then the update equation for $z_{\alpha,t}$ is:

$$z_{\alpha,t+1} = \frac{1}{\sqrt{n}} \left(\mathbf{v}_t - \frac{\eta}{\sqrt{nD}} \sum_{\beta=1}^D \epsilon_{\beta,t} \sigma(\mathbf{U}_t \mathbf{x}_\beta) \right)^\top \sigma \left(\mathbf{U}_t \mathbf{x}_\alpha - \frac{\eta}{\sqrt{nD}} \sum_{\gamma=1}^D \epsilon_{\gamma,t} (\mathbf{v}_t \circ \sigma'(\mathbf{U}_t \mathbf{x}_\gamma)) \mathbf{x}_\gamma^\top \mathbf{x}_\alpha \right). \tag{57}$$

In the infinite width limit the update equation for the error ϵ is:

$$\epsilon_{\alpha,t+1} = \sum_{\beta=1}^D (\delta_{\alpha\beta} - \eta H_{\alpha\beta,0}) \epsilon_{\beta,t} + O(1/n), \tag{58}$$

where the NTK is frozen at its initial value, see e.g. chapter 10 of Roberts *et al* (2022) for a review. When $\eta \lambda_{\max}(H_{\alpha\beta,0}) < 2$ the errors converge to zero exponentially fast while when $\eta \lambda_{\max}(H_{\alpha\beta,0}) > 2$ the errors grow exponentially quickly.

Here we run into an important subtlety, the activation function σ is not differentiable around $x = 0$. More generally, the series expansion of $\sigma(x)$ around a point $x = x_0$ has a radius of convergence $r = |x_0|$. This means that we cannot necessarily expand (57) around $n = \infty$ if ϵ_α is large enough to flip the sign of $\mathbf{U} \mathbf{x}_\alpha$ ¹⁶.

This breakdown of perturbation theory is of a different nature than what we saw for the quadratic model in the catapult phase. In the quadratic model the Taylor expansion

¹⁶ To be more precise, the large n expansion breaks down at finite width because Taylor expanding the activation function produces increasingly singular δ -functions. These distributions are not integrable when we compute expectation values over the initial weights. Therefore, we cannot use the large n expansion to study the evolution of the NTK. For more details see the discussion around equation ($\infty.39$) of Roberts *et al* (2022).

of the *update equations*, (30) and (31), truncates at order ζ^2 for all t . Perturbation theory breaks down for super-critical learning rates because we cannot similarly truncate z_t and $H_{\alpha\beta,t}$ at a low order in ζ once $z_t = O(\zeta^{-1})$. Instead, we are forced to work to all orders in the small ζ expansion. On the other hand, for the homogeneous MLPs studied here, the large- n expansion of the update equations is simply not valid when $z_{\alpha,t}$ is large because σ is not differentiable at the origin. This breakdown of perturbation theory also occurs when η is small and the model is in the lazy phase.

For this reason, to estimate at what time-scale the large- n expansion breaks down we will look at the update equation for \mathbf{U}_t directly. At initialization we have $v_{i,0}, U_{ij,0} \sim \mathcal{N}(0,1)$ for all i and j . Assuming \mathbf{x}_α and y_α do not scale with n , we expect an order-one fraction of the components of $\mathbf{U}_t \mathbf{x}_\alpha$ to flip signs when $z_t = O(\sqrt{n})$, which occurs after $O(\log(\sqrt{n}))$ time-steps. This argument is less rigorous than the one given above for the quadratic model, but it agrees with the numerical results of Lewkowycz *et al* (2020) and what we find numerically.

Appendix B. Derivations for a single datapoint

In this section we will prove that the catapult phase exists for super-critical learning rates when training on a single datapoint. With the exception of the ReLU MLP, we will work with the toy dataset $(x, y) = (1, 0)$. For the ReLU MLP we will take the label $y > 0$ to avoid the trivial solution where all the first layer weights are negative. We will drop the sample index everywhere since we only have one data-point.

B.1. Pure quadratic model

We will start with the pure quadratic model:

$$z = \frac{\zeta}{2} \boldsymbol{\theta}^T \boldsymbol{\psi} \boldsymbol{\theta}, \quad H = \zeta^2 \boldsymbol{\theta}^T \boldsymbol{\psi}^2 \boldsymbol{\theta}. \tag{59}$$

We will prove that the weight norm $\boldsymbol{\theta}_t^2$ decreases monotonically if:

$$\eta < \frac{4}{\zeta^2 \boldsymbol{\theta}_0^2 \lambda_{\max}(\boldsymbol{\psi}^2)} \tag{60}$$

and therefore that the loss will remain finite. The proof was sketched in section 3.3 and here we will fill in the details.

First, we use that under gradient descent the weight norm evolves as:

$$\boldsymbol{\theta}_{t+1}^2 = \boldsymbol{\theta}_t^2 + \eta z_t^2 (\eta H_t - 4). \tag{61}$$

Next, we use that the meta-feature function $\boldsymbol{\psi}$ is a real symmetric matrix. Therefore $\boldsymbol{\psi}^2$ is a positive semi-definite matrix and the NTK in this model obeys the bound:

$$H_t \leq \zeta^2 \lambda_{\max}(\boldsymbol{\psi}^2) \boldsymbol{\theta}_t^2. \tag{62}$$

This has the same form as the bound (13). We will now assume that at some time-step t^* we have:

$$\eta < \frac{4}{\zeta^2 \theta_{t^*}^2 \lambda_{\max}(\psi^2)}. \tag{63}$$

If (63) holds at some time t^* , then we can prove that the weight norm θ_t^2 is a monotonically decreasing function for all subsequent steps:

$$\theta_{t+1}^2 - \theta_t^2 < 0 \quad \forall t \geq t^*. \tag{64}$$

This statement follows directly from the update equation (61):

$$\begin{aligned} \theta_{t^*+1}^2 - \theta_{t^*}^2 &= \eta z_{t^*}^2 (\eta H_{t^*} - 4) \\ &\leq \eta z_{t^*}^2 (\eta \zeta^2 \theta_{t^*}^2 \lambda_{\max}(\psi^2) - 4) \\ &< 0. \end{aligned} \tag{65}$$

To get the second line we used (62) and to get the third line we used (63). This shows that the weight norm decreases under one step of gradient descent. Moreover, since $\theta_{t^*+1}^2 < \theta_{t^*}^2$ we also have that (63) continues to hold at time $t = t^* + 1$:

$$\eta < \frac{4}{\zeta^2 \theta_{t^*}^2 \lambda_{\max}(\psi^2)} < \frac{4}{\zeta^2 \theta_{t^*+1}^2 \lambda_{\max}(\psi^2)}. \tag{66}$$

Therefore, once (63) holds at some time t^* , it holds for all subsequent times. Then from (61) this implies that the weight norm θ_t^2 also decreases for all subsequent times. To complete the proof we then just need to assume the bound on η (63) holds at $t = 0$ to guarantee that the weight norm decreases monotonically for all time. This gives the original condition (60) and completes the proof.

Note that for this argument to work we did not need to use that the model undergoes any catapult dynamics. We then have two separate cases:

1. If $2H_0 > \zeta^2 \theta_0^2 \lambda_{\max}(\psi^2)$, then there is a finite window above the linear stability threshold $\eta H_0 = 2$ where the model converges and the loss exhibits the catapult mechanism.
2. If $2H_0 < \zeta^2 \theta_0^2 \lambda_{\max}(\psi^2)$, then this argument does not guarantee that the model has a catapult phase, but it does imply there exists a region in the lazy phase where the weight norm decays monotonically.

These conditions are dependent on the values of the weights at initialization. We can also ask when these conditions hold in expectation by averaging over the weights. For example, the inequality $2H_0 > \zeta^2 \theta_0^2 \lambda_{\max}(\psi^2)$ holds in expectation if:

$$\begin{aligned} \mathbb{E}_{\theta_0} [2H_0 - \zeta^2 \theta_0^2 \lambda_{\max}(\psi^2)] &> 0 \\ \implies \lambda_{\max}(\psi^2) &< \frac{2}{n} \text{Tr}(\psi^2). \end{aligned} \tag{67}$$

That is, the maximum eigenvalue of ψ^2 cannot differ too significantly from the mean. This inequality is closest to being saturated when all the eigenvalues of ψ^2 are the same, which is what happens in a two-layer MLP with a linear activation function (Agarwala *et al* 2022, Zhu *et al* 2022).

We can use the same type of argument to prove that if:

$$\eta > \frac{4}{\zeta^2 \theta_0^2 \lambda_{\min}(\psi^2)}, \tag{68}$$

then the loss diverges. The proof is almost the same as before, except we need to use the following lower bound on the NTK in terms of the minimal eigenvalue of ψ^2 :

$$H_t \geq \zeta^2 \lambda_{\min}(\psi^2) \theta_t^2. \tag{69}$$

Following the same ideas as before, we start by assuming that at some time-step t_* we have:

$$\eta > \frac{4}{\zeta^2 \theta_{t_*}^2 \lambda_{\min}(\psi^2)}. \tag{70}$$

Then from the update equation for θ^2 (61) we have:

$$\begin{aligned} \theta_{t_*+1}^2 - \theta_{t_*}^2 &= \eta z_{t_*}^2 (\eta H_{t_*} - 4) \\ &\geq \eta z_{t_*}^2 (\eta \zeta^2 \lambda_{\min}(\psi^2) \theta_{t_*}^2 - 4) \\ &> 0. \end{aligned} \tag{71}$$

To get the second line we used (69) and to get the third line we used our assumption (70). Since here the weight norm has increased after one step of gradient descent, this implies (70) continues to hold at time-step $t = t_* + 1$:

$$\eta > \frac{4}{\zeta^2 \theta_{t_*}^2 \lambda_{\min}(\psi^2)} > \frac{4}{\zeta^2 \theta_{t_*+1}^2 \lambda_{\min}(\psi^2)}. \tag{72}$$

Therefore, by induction, the weight norm monotonically increases for all future time-steps. Finally, this means if the lower bound holds at initialization, see (68), then the weight norm increases monotonically for all time.

The fact the weight norm monotonically increases for all times does not, by itself, imply the loss diverges. However, we expect the loss will diverge since if the lower bound (68) holds, then the learning rate η is necessarily super-critical:

$$\eta H_0 > \eta \lambda_{\min}(\psi^2) \zeta^2 \theta_0^2 > 4, \tag{73}$$

where to find the first inequality we used (69) and to find the second inequality we used (68). Moreover, since θ_t^2 is a monotonically increasing function, this implies $\eta H_t > 4$ for all time t and the model can never re-enter the lazy phase. Altogether, this implies the point $z_t = 0$ is an unstable fixed point of the update equations when the learning rate satisfies the lower bound (68). If z_t vanishes exactly, then the model stays at $z_t = 0$ for all time. However, if z_t is small, but non-zero, then we can use the small

ζ -perturbation theory arguments of appendix A to show that the loss will increase exponentially quickly. Therefore, barring fine-tuning where the update equations set $z_t = 0$ exactly, the output will be exponentially large and θ_t^2 receives, large, positive updates for all t and will diverge¹⁷.

To summarize, we have shown that if η satisfies the bound (60) then the weights decay monotonically and if η satisfies the bound (68) then the weights increase monotonically. This leaves the in-between region:

$$\frac{4}{\zeta^2 \theta_0^2 \lambda_{\max}(\psi^2)} < \eta < \frac{4}{\zeta^2 \theta_0^2 \lambda_{\min}(\psi^2)}, \tag{74}$$

where we cannot say anything definite with our methods. To say something definite about this region would require understanding how the details of the meta-feature eigensystem affects the evolution of the quadratic model. Here we can note that this region shrinks to zero size when:

$$\lambda_{\min}(\psi^2) = \lambda_{\max}(\psi^2). \tag{75}$$

In this case there is a very sharp delineation between the catapult and divergent phases. The condition (75) is satisfied in the two-layer MLP with linear activation functions. The phase boundary between the catapult and divergent phase for this model was first found in Lewkowycz *et al* (2020).

B.2. Quadratic model with bias

In this appendix we study the quadratic model with bias,

$$z = \theta^T \phi + \frac{\zeta}{2} \theta^T \psi \theta, \quad H = \phi^2 + \zeta^2 \theta^T \psi^2 \theta, \tag{76}$$

where $\psi \phi = 0$. We train this model on the toy dataset $(x, y) = (1, 0)$. We will prove this model converges if:

$$\eta < \frac{4}{2\phi^2 + \zeta^2 \lambda_{\max}(\psi^2) \left(\theta_0^2 + \frac{(\phi^T \theta_0)^2}{\phi^2} \right)}. \tag{77}$$

In the quadratic model with bias, the update equation for θ_t^2 is:

$$\theta_{t+1}^2 = \theta_t^2 - 2\eta z_t [\phi^T \theta_t + \zeta (\theta_t^T \psi \theta_t)] + \eta^2 z_t^2 H_t. \tag{78}$$

Unlike the pure quadratic model, here the update equation for θ_t^2 does not have a nice sign-definiteness property. In both the pure quadratic model and the quadratic model with bias, the η^2 term in the θ_t^2 update equation is manifestly positive semi-definite since z_t is real and the NTK is positive semi-definite. On the other hand, for the pure

¹⁷ The weight norm θ^2 can receive small updates if $\eta H_t \rightarrow 4$ as $t \rightarrow \infty$, but given our assumption (68) we have that $\eta H_t - 4$ is strictly bounded away from 0 for all t .

quadratic model the order η term was manifestly negative semi-definite, see (61), while in (78) it is not clear if the order η term has a definite sign. The fact the order η term in (61) is negative semi-definite was important in proving that the pure quadratic model can have a catapult phase. To remedy the lack of sign-definiteness in (78) we will study a slightly different update equation.

Specifically, we will study the evolution of:

$$\boldsymbol{\theta}_t^2 + \frac{(\boldsymbol{\phi}^\top \boldsymbol{\theta}_t)^2}{\phi^2}. \tag{79}$$

Under gradient descent the second term of (79) evolves as:

$$\frac{(\boldsymbol{\phi}^\top \boldsymbol{\theta}_{t+1})^2}{\phi^2} = \frac{(\boldsymbol{\phi}^\top \boldsymbol{\theta}_t)^2}{\phi^2} - 2\eta z_t \boldsymbol{\phi}^\top \boldsymbol{\theta}_t + \eta^2 z_t^2 \phi^2, \tag{80}$$

where we used that $\boldsymbol{\psi}\boldsymbol{\phi} = 0$. Using the above update equation we find the quantity (79) evolves as:

$$\begin{aligned} \boldsymbol{\theta}_{t+1}^2 + \frac{(\boldsymbol{\theta}_{t+1}^\top \boldsymbol{\phi})^2}{\phi^2} &= \boldsymbol{\theta}_t^2 + \frac{(\boldsymbol{\theta}_t^\top \boldsymbol{\phi})^2}{\phi^2} - 2\eta z_t (2(\boldsymbol{\theta}_t \boldsymbol{\phi}) + \zeta(\boldsymbol{\theta}_t^\top \boldsymbol{\psi} \boldsymbol{\theta}_t)) + \eta^2 z_t^2 (H_t + \phi^2) \\ &= \boldsymbol{\theta}_t^2 + \frac{(\boldsymbol{\theta}_t^\top \boldsymbol{\phi})^2}{\phi^2} + \eta z_t^2 (\eta(H_t + \phi^2) - 4). \end{aligned} \tag{81}$$

Now the order η and order η^2 terms in (81) are sign-definite and we can prove that $\boldsymbol{\theta}_t^2 + \frac{(\boldsymbol{\theta}_t^\top \boldsymbol{\phi})^2}{\phi^2}$ decays monotonically if (77) holds.

The proof will be the same as for the pure quadratic model. To start, we use that the NTK obeys the following upper bound:

$$\begin{aligned} H_t &= \phi^2 + \zeta^2 \boldsymbol{\theta}_t^\top \boldsymbol{\psi}^2 \boldsymbol{\theta}_t \leq \phi^2 + \zeta^2 \lambda_{\max}(\boldsymbol{\psi}^2) \boldsymbol{\theta}_t^2 \\ &\leq \phi^2 + \zeta^2 \lambda_{\max}(\boldsymbol{\psi}^2) \left(\boldsymbol{\theta}_t^2 + \frac{(\boldsymbol{\phi}^\top \boldsymbol{\theta}_t)^2}{\phi^2} \right). \end{aligned} \tag{82}$$

The first inequality follows from the fact $\boldsymbol{\psi}^2$ is a positive semi-definite matrix. The second inequality is trivial as we just added a manifestly positive quantity to the previous expression. Therefore the $\eta(H_t + \phi^2) - 4$ term in (81) will be negative at time-step $t = t^*$ if we impose the inequality:

$$\eta < \frac{4}{2\phi^2 + \zeta^2 \lambda_{\max}(\boldsymbol{\psi}^2) \left(\boldsymbol{\theta}_{t^*}^2 + \frac{(\boldsymbol{\phi}^\top \boldsymbol{\theta}_{t^*})^2}{\phi^2} \right)}. \tag{83}$$

This condition is sufficient to guarantee $\eta(H_{t^*} + \phi^2) - 4 < 0$, but may not be necessary. If the condition (83) holds at time t^* , then from (81) we have that at time $t = t^* + 1$:

$$\boldsymbol{\theta}_{t^*+1}^2 + \frac{(\boldsymbol{\theta}_{t^*+1}^\top \boldsymbol{\phi})^2}{\phi^2} < \boldsymbol{\theta}_{t^*}^2 + \frac{(\boldsymbol{\theta}_{t^*}^\top \boldsymbol{\phi})^2}{\phi^2}. \tag{84}$$

The denominator in (83) is a sum of manifestly positive quantities, so if $\theta^2 + \frac{\theta^T \phi}{\phi^2}$ is decreasing the right-hand side must increase:

$$\eta < \frac{4}{2\phi^2 + \zeta^2 \lambda_{\max}(\psi^2) \left(\theta_{t^*}^2 + \frac{(\phi^T \theta_{t^*})^2}{\phi^2} \right)} < \frac{4}{2\phi^2 + \zeta^2 \lambda_{\max}(\psi^2) \left(\theta_{t^*+1}^2 + \frac{(\phi^T \theta_{t^*+1})^2}{\phi^2} \right)}. \tag{85}$$

This shows that if (83) holds at time t^* , then it also holds for all subsequent steps and therefore $\theta_t^2 + \frac{(\phi^T \theta_t)^2}{\phi^2}$ will decrease monotonically for all subsequent steps as well. Finally, we can guarantee that the weight norm θ_t^2 does not diverge, and therefore the loss remains finite, by imposing this condition at initialization. This gives the original inequality (77)¹⁸.

Example: Linear Net with Bias

Here we will make the previous analysis more concrete by considering a simple example of a quadratic model with bias, the two-layer, linear net with bias on the final layer:

$$z = \frac{1}{\sqrt{n}} \mathbf{v}^T \mathbf{u} + b, \tag{86}$$

where $\mathbf{u}, \mathbf{v} \in \mathbb{R}^n$ and $b \in \mathbb{R}$ ¹⁹. The NTK is given by:

$$H = \frac{1}{n} (\mathbf{u}^2 + \mathbf{v}^2) + 1. \tag{87}$$

We train this model to minimize the MSE loss:

$$L_t = \frac{1}{2} z_t^2. \tag{88}$$

To explain the connection between the linear net with bias and the quadratic model with bias, it is simplest to work in index notation. First, we define the abstract weights $\theta \in \mathbb{R}^{2n+1}$ as the concatenation of \mathbf{u} , \mathbf{v} and b :

$$\theta = (\mathbf{u}, \mathbf{v}, b)^T. \tag{89}$$

Here it is understood that the first n indices of θ_μ correspond to \mathbf{u} , the second n indices to \mathbf{v} and last index to b . Then to make the map between the quadratic model with bias and the linear net with bias more explicit, we will label the indices of ϕ and ψ by the corresponding weights, \mathbf{u} , \mathbf{v} and b . For example, we will write the output z as:

$$\begin{aligned} z &= \phi^T \theta + \frac{\zeta}{2} \theta^T \psi \theta \\ &= \phi_b b + \zeta \sum_{i,j=1}^n u_i \psi_{u_i} v_j v_j, \end{aligned} \tag{90}$$

¹⁸ Here we are using the trivial inequality $\theta^2 \leq \theta^2 + \frac{(\theta^T \phi)^2}{\phi^2}$, so if the right-hand side is finite, the left-hand side will be finite as well.
¹⁹ We could also allow for bias in the hidden layer, but this will not change anything qualitatively.

where $\zeta^2 = 1/n$. Here the only non-trivial components of the (meta-)feature functions are given by:

$$\phi_b = 1, \quad \psi_{u_i v_j} = \delta_{ij}, \tag{91}$$

where δ_{ij} is the Kronecker delta function. The other components, e.g. ϕ_{u_i} , ϕ_{v_i} , $\psi_{u_i u_j}$, $\psi_{v_i v_j}$, etc all vanish in the linear net with bias.

Then the sufficiency condition that this model converges (77) becomes:

$$\eta < \frac{4}{2 + 1/n(\mathbf{u}_0^2 + \mathbf{v}_0^2 + 2b_0^2)}. \tag{92}$$

If we assume that at initialization the bias is set to zero, $b_0 = 0$, then this reduces to:

$$\eta < \frac{4}{2 + 1/n(\mathbf{u}_0^2 + \mathbf{v}_0^2)} = \frac{4}{H_0 + 1}. \tag{93}$$

B.3. Homogeneous, two-layer, MLPs

Here we will study the homogeneous, two-layer net,

$$z = \frac{1}{\sqrt{n}} \sum_{i=1}^n v_i \sigma(u_i), \quad H = \frac{1}{n} (\sigma^2(u_i) + v_i^2 \sigma'^2(u_i)), \tag{94}$$

trained on the toy dataset $(x, y) = (1, 0)$. The activation function σ is given in (8) and we assume $0 < a_- \leq a_+ \leq 1$. We will prove this model converges if:

$$\eta < \frac{4n}{a_+^2(\mathbf{u}_0^2 + \mathbf{v}_0^2)} = \frac{4n}{a_+^2 \theta_0^2}. \tag{95}$$

This analysis will be almost identical to the analysis of the pure quadratic model given in appendix B.1, so we will be brief.

First, the weight norm in this model is $\theta^2 = \mathbf{u}^2 + \mathbf{v}^2$. Under gradient descent θ^2 evolves as:

$$\theta_{t+1}^2 = \theta_t^2 + \eta z_t^2 (\eta H_t - 4). \tag{96}$$

Note that this is identical to the weight update equation for the pure quadratic model (61). This is not a coincidence and instead comes from the fact that in both models z_t is a homogeneous, quadratic function of θ and that we are training the models with MSE.

Next, we use the following upper bound for the NTK in terms of the weight norm:

$$H_t \leq a_+^2 \frac{\mathbf{u}_t^2 + \mathbf{v}_t^2}{n} = \frac{a_+^2}{n} \theta_t^2. \tag{97}$$

This bound is identical in form to the bound used when studying the pure quadratic model, see (62), except with the replacement:

$$\zeta^2 \lambda_{\max}(\psi^2) \implies \frac{a_+^2}{n}. \tag{98}$$

Since the weight update equation (96) and NTK bound (97) have the same functional form as in the pure quadratic model, we can repeat the steps used in appendix B.1 to derive the new bound (95) for the homogeneous MLP. In particular, if we take the bound derived for the pure quadratic model (60) and make the replacement (98), we land exactly on (95).

As with the pure quadratic model trained on $(x, y) = (1, 0)$, we now have two separate cases to consider:

1. If $n^{-1}a_+^2\theta_0^2 < 2H_0$, then there exists a finite window above the linear stability threshold, $\eta H_0 = 2$, where the model still converges. In this case the loss exhibits the catapult mechanism.
2. If $n^{-1}a_+^2\theta_0^2 > 2H_0$, then our bound only guarantees there is a choice of learning rate in the lazy phase such that θ_t^2 decreases monotonically.

To determine which case happens on average, we need the expectation values of the NTK and the weight norm at initialization:

$$\mathbb{E}[H_0] = (a_-^2 + a_+^2), \tag{99}$$

$$\mathbb{E}[\theta_0^2] = 2n. \tag{100}$$

Therefore:

$$\mathbb{E}\left[2H_0 - \frac{a_+^2}{n}\theta_0^2\right] > 0 \implies a_- \neq 0, \tag{101}$$

which includes all scale-invariant activation functions with the exception of the ReLU net. Of course, this does not mean that the ReLU net cannot obey the bound (95) for some choice of initial weights, but on average it obeys this condition only half of the time.

Finally, we can also prove that above a certain learning rate the model will diverge. Specifically, if

$$\eta > \frac{4n}{a_-^2\theta_0^2}, \tag{102}$$

then the weight norm grows monotonically and the loss diverges. To prove this, we use that the NTK obeys the following lower bound:

$$H_t \geq \frac{a_-^2}{n}\theta_t^2. \tag{103}$$

Note that here we are using our assumption that $0 < a_- \leq a_+$. If $a_- = 0$ this lower bound is trivial. The lower bound (103) also has the same form as the lower bound (68) for the pure quadratic model. To go from the pure quadratic model to the homogeneous net we make the replacement:

$$\zeta^2 \lambda_{\min}(\psi^2) \implies \frac{a_-^2}{n}. \tag{104}$$

Then, since the update equation for θ_t^2 and the lower bound on the NTK have the same form as in the pure quadratic model, we can again borrow from the analysis of appendix B.1. We will not repeat the argument here, but can instead use the replacement (104) in (68) to give exactly the lower bound (102) for the homogeneous MLP.

B.4. ReLU net

In the previous section we derived sufficiency conditions for the catapult phase to exist in homogeneous, two-layer nets. However, these conditions were too strong for ReLU nets: the allowed range for the learning rate η vanished on average. In this section we will derive weaker bounds on η by using some simplifying features of the ReLU function.

We again study a two-layer net:

$$z = \frac{1}{\sqrt{n}} \sum_{i=1}^n v_i \sigma_{\text{ReLU}}(u_i), \quad H = \frac{1}{n} (\sigma_{\text{ReLU}}^2(u_i) + v_i^2 \sigma_{\text{ReLU}}^{\prime 2}(u_i)). \quad (105)$$

Here we will train the model using MSE loss:

$$L = \frac{1}{2} (z - y)^2, \quad (106)$$

where $y \geq 0$. We will consider two cases, $y = 0$ and $y > 0$ with $y = O(1)$. The first case is trivial to solve, we can simply take $u_i < 0$ for all i . However, it is a useful case to consider to set up notation.

$y = 0$:

The reason our bounds for the generic homogeneous net are too strong for the ReLU net is that, on average, only half of the weights will evolve in the ReLU MLP when we have one datapoint. It is then overkill to study the evolution of the full weight norm $\theta_t^2 = \mathbf{u}_t^2 + \mathbf{v}_t^2$ and we should instead only focus on the weights which evolve for at least one step of gradient descent. We can do this by decomposing the weights at initialization as follows:

$$\mathbf{u}_0 = \mathbf{u}_{+,0} + \mathbf{u}_{-,0}, \quad (107)$$

$$\mathbf{u}_{+,0} = \sigma_{\text{ReLU}}(\mathbf{u}_0), \quad (108)$$

$$\mathbf{u}_{-,0} = \sigma_{\text{ReLU}}(-\mathbf{u}_0). \quad (109)$$

That is, $\mathbf{u}_{\pm,0}$ are two n -dimensional vectors whose non-zero components correspond to the components of \mathbf{u}_0 which are positive/negative at $t=0$, respectively. By definition they are orthogonal $\mathbf{u}_{+,0}^T \mathbf{u}_{-,0} = 0$. We can then define matrices which project onto the \mathbf{u}_{\pm} directions:

$$\mathbf{P}^+ = \mathbb{I}_{n \times n} \circ \text{diag}(\mathbf{1}_{\sigma(\mathbf{u}) \geq 0}), \quad (110)$$

$$\mathbf{P}^- = \mathbb{I}_{n \times n} \circ \text{diag}(\mathbf{1}_{\sigma(\mathbf{u}) < 0}). \quad (111)$$

Here ‘diag’ produces a diagonal matrix from the given vector. These projectors may be clearer in index notation, where we can write:

$$P_{ij}^+ = \delta_{ij} \mathbf{1}_{u_{i,0} \geq 0}, \quad (112)$$

$$P_{ij}^- = \delta_{ij} \mathbb{1}_{u_{i,0} < 0}, \tag{113}$$

where δ_{ij} is the Kronecker delta function. We can then define $\mathbf{v}_{\pm,t}$ as:

$$\mathbf{v}_{+,t} = \mathbf{P}^+ \mathbf{v}_t, \tag{114}$$

$$\mathbf{v}_{-,t} = \mathbf{P}^- \mathbf{v}_t. \tag{115}$$

That is, $\mathbf{v}_{\pm,t}$ is the projection of \mathbf{v}_t onto the directions where $\mathbf{u}_{\pm,0}$ are non-zero.

Under gradient descent the weights evolve as:

$$\mathbf{u}_{+,t+1} = \mathbf{u}_{+,t} - \frac{1}{\sqrt{n}} \eta z_t \mathbf{v}_{+,t} \circ \sigma'_{\text{ReLU}}(\mathbf{u}_{+,t}), \tag{116}$$

$$\mathbf{v}_{+,t+1} = \mathbf{v}_{+,t} - \frac{1}{\sqrt{n}} \eta z_t \sigma_{\text{ReLU}}(\mathbf{u}_{+,t}), \tag{117}$$

$$\mathbf{u}_{-,t+1} = \mathbf{u}_{-,t}, \tag{118}$$

$$\mathbf{v}_{-,t+1} = \mathbf{v}_{-,t}. \tag{119}$$

We can rewrite the output z and NTK as:

$$z_t = \frac{1}{\sqrt{n}} \mathbf{v}_{+,t}^T \sigma_{\text{ReLU}}(\mathbf{u}_{+,t}), \tag{120}$$

$$H_t = \frac{1}{n} \left((\sigma_{\text{ReLU}}(\mathbf{u}_{+,t}))^2 + (\mathbf{v}_{+,t} \circ \sigma'_{\text{ReLU}}(\mathbf{u}_{+,t}))^2 \right). \tag{121}$$

Then, instead of studying the full weight norm $\mathbf{u}_t^2 + \mathbf{v}_t^2$, we can study the reduced weight norm $\boldsymbol{\theta}_{+,t}^2 \equiv \mathbf{u}_{+,t}^2 + \mathbf{v}_{+,t}^2$,

$$\boldsymbol{\theta}_{+,t+1}^2 = \boldsymbol{\theta}_{+,t}^2 + \eta z_t^2 (\eta H_t - 4). \tag{122}$$

Note that the update equation for the reduced weight norm is identical in form to the update equation for the full weight norm in the pure quadratic model and two-layer, homogeneous net, see (61) and (96). Therefore, we can now recycle our previous proofs for the existence of the catapult phase for this model. As a first step, we use that H_t is bounded from above by this reduced weight norm:

$$H_t \leq \frac{1}{n} (\mathbf{u}_{+,t}^2 + \mathbf{v}_{+,t}^2) = \frac{\boldsymbol{\theta}_{+,t}^2}{n}. \tag{123}$$

At $t = 0$ this inequality is saturated but when $t > 0$ it is in general a strict inequality. This is because at $t = 0$ each component of $\mathbf{u}_{+,t}$ is positive and we have $\sigma_{\text{ReLU}}(\mathbf{u}_{+,0}) = \mathbf{u}_{+,0}$ by definition. However, under gradient descent some of the components of $\mathbf{u}_{+,t}$ can become negative, in which case they are set to zero by the ReLU activation function and only contribute to the right hand side of (123).

Finally, we can use the inequality (123) to show that if

$$\eta < \frac{4n}{(\mathbf{u}_{+,0}^2 + \mathbf{v}_{+,0}^2)} = \frac{4}{H_0}, \tag{124}$$

then the weight norm monotonically decreases under gradient descent. We will not write out the proof here since it is identical in form to the proof used for the pure quadratic model, see appendix B.1. The only differences are that here we are studying θ_+^2 instead of θ^2 and to go from the bound (62) to (124) we make the replacement $\zeta^2 \lambda_{\max}(\psi^2) \theta^2 \rightarrow \theta_+^2/n$. If we make this replacement in (60) we find the bound (124).

$y > 0$:

We can now generalize the previous argument to the case where $y > 0$, but is still $O(1)$. The main difference is now the update equations for \mathbf{u}_+ and \mathbf{v}_+ are:

$$\mathbf{u}_{+,t+1} = \mathbf{u}_{+,t} - \frac{1}{\sqrt{n}} \eta \epsilon_t \mathbf{v}_{+,t} \circ \sigma'_{\text{ReLU}}(\mathbf{u}_{+,t}), \tag{125}$$

$$\mathbf{v}_{+,t+1} = \mathbf{v}_{+,t} - \frac{1}{\sqrt{n}} \eta \epsilon_t \sigma_{\text{ReLU}}(\mathbf{u}_{+,t}). \tag{126}$$

The reduced weight norm θ_+^2 then evolves as:

$$\theta_{+,t+1}^2 = \theta_{+,t}^2 - 4\eta z_t \epsilon_t + \eta^2 \epsilon_t^2 H_t. \tag{127}$$

Unlike the previous case, where $y = 0$, now the order η term does not generically have a definite sign. However, this is a minor issue since when $\eta H_0 > 2$ the output z_t grows exponentially quickly. When $|z_t| \gg 1$ we have:

$$\epsilon_t = z_t - y \approx z_t. \tag{128}$$

Therefore, if $\eta H_0 > 2$ and we study the model at a time t^* such that $1 \ll t^* \ll \log(n)$, then $|y| \ll |z_{t^*}| \ll 1/\sqrt{n}$ and the update equation simplifies to²⁰:

$$\theta_{+,t^*+1}^2 \approx \theta_{+,t^*}^2 + \eta z_{t^*}^2 (\eta H_{t^*} - 4), \quad \text{for } 1 \ll t^* \ll \log(n). \tag{129}$$

In this case, the weight norm receives large, $O(n)$, negative corrections when $z_t = O(\sqrt{n})$ if we impose:

$$\eta < \frac{4n}{a_+^2 \theta_{+,t^*}^2} = \frac{4n}{a_+^2 \theta_{+,0}^2} + O(1/n). \tag{130}$$

Here we used that when $t^* \ll \log(n)$ the weight norm θ_t^2 changes at most by an $O(1)$ amount (recall θ_0^2 is $O(n)$), see appendix A. To conclude, if η is smaller than $\frac{4n}{a_+^2 \theta_{+,0}^2}$ by more than a parametrically order $1/n$ amount, then the reduced weight norm $\theta_{+,t}^2$ receives $O(n)$ negative corrections in the catapult phase. Since the weight norm is a positive-definite quantity, it must eventually stop receiving large, additive, negative updates. This will happen when $|z_t| = O(1)$, since at this point the order η term in (127) does not have a definite sign. When this happens, the model has either re-entered the lazy phase ($\eta H_t < 2$) and converges or the model still has a super-critical learning rate ($\eta H_t > 2$).

²⁰ In practice, because z_t increases exponentially fast and $|y|$ is an order-one number, we can take t^* also to be an order-one number. The condition $t^* \gg 1$ ensures that $\epsilon_{t^*} \approx z_{t^*}$ up to exponentially suppressed corrections.

To complete our argument, we need to show that if the model has a super-critical learning rate after the first catapult mechanism has finished, then subsequent catapult mechanisms will eventually drive the model to the lazy phase²¹. To do this, we will use that during the first catapult mechanism θ^2 decreased by an $O(n)$ amount, but that between the end of the first catapult mechanism and the beginning of the second one, the weight norm θ^2 can change at most by an $O(1)$ amount. It then follows that (130) will hold during the second catapult mechanism and the weight norm will again receive $O(\zeta^{-2})$ negative updates.

To make this concrete, we will denote the time at which the first catapult mechanism ends by t_1 . At t_1 we have:

$$|z_{t_1}| = O(1), \tag{131}$$

$$\eta H_{t_1} > 2, \tag{132}$$

$$|\theta_{+,t^*}^2 - \theta_{+,t_1}^2| = O(n), \tag{133}$$

$$\theta_{+,t_1}^2 < \theta_{+,t^*}^2. \tag{134}$$

Since the model still has a super-critical learning rate at time t_1 , the output z_t will again grow exponentially fast. Therefore, at time t_2 , such that $1 \ll t_2 - t_1 \ll \log(n)$, we again have $\epsilon_{t_2} \approx z_{t_2}$. We then want to show the analog of (130) holds at the beginning of the second catapult mechanism:

$$\eta < \frac{4n}{a_+^2 \theta_{+,t_2}^2} \quad \text{for } 1 \ll t_2 - t_1 \ll \log(n). \tag{135}$$

If (135) is true, then θ_t^2 also decreases by an $O(n)$ amount during the second catapult mechanism. To prove (135) holds, we use that under gradient descent $\theta_{+,t}^2$ changes by:

$$|\theta_{+,t+1}^2 - \theta_{+,t}^2| = O(\eta z_t^2). \tag{136}$$

We can now track how $\theta_{+,t}^2$ evolves from time t_1 , when the first catapult mechanism ends and $|z_{t_1}| = O(1)$, to time t_2 , when the output again becomes large, or $|z_t| \gg 1$. Since $y = O(1)$ and z_t grows exponentially fast, we can take $t_2 - t_1 = O(1)$. In this case, the weight norm cannot change significantly from time t_1 to time t_2 :

$$|\theta_{+,t_2}^2 - \theta_{+,t_1}^2| \ll O(n). \tag{137}$$

Here we used that $t_2 - t_1 = O(1)$ and that if $t_2 - t_1 \ll \log(n)$ then $|z_{t_2}| \ll \sqrt{n}$. Combining (133), (134) and (137), we have $\theta_{+,t_2}^2 < \theta_{+,t^*}^2$ and the inequality (135) follows from the inequality (130). Finally, this means $\theta_{+,t}^2$ receives large, negative updates during the second catapult mechanism. One can extend this argument by induction to all possible subsequent catapult mechanisms.

²¹ Here by ‘catapult mechanism’ we mean that the loss curve exhibits the characteristic catapult shape where it rapidly increases and then decreases. This should be contrasted with the ‘catapult phase’ which is the choice of super-critical learning rates where the model converges. For a sufficiently large learning rate the model can in principle undergo many catapult mechanisms before the loss converges, i.e. the loss curve has multiple peaks.

We can note that the above subtlety, that once we include non-trivial labels y_α we can only argue that the updates to θ_t^2 are negative when the loss is large, is not unique to the ReLU MLP. This is a generic feature of the θ_t^2 update equation with non-trivial labels and will arise repeatedly in appendix C. We should also note that the same type of argument was used in Lewkowycz *et al* (2020) to argue that the NTK of the two-layer, linear MLP trained on a generic dataset receives large negative updates whenever the loss is large. For the models studied in this work, we generally find it is the weight norm θ_t^2 , and not the NTK, which decreases monotonically when the loss is large.

Appendix C. Multiple datapoints

In this section we will explain how to derive bounds for generic datasets $(\mathbf{x}_\alpha, y_\alpha)$ with $\alpha = 1, \dots, D$. We will always take the labels to be one-dimensional, although the generalization to multi-dimensional output is straightforward. We will also assume that the labels y_α do not scale with ζ or D , $y_\alpha = O(\zeta^0 D^0)$. The main results for the pure quadratic model are given in (138) and (154), for the quadratic model with bias in (163), and for the two-layer, homogeneous MLP in (180).

C.1. Pure quadratic model

We will consider two different methods to prove that pure quadratic models with generic data can converge at super-critical learning rates. The first method will be more rigorous but will give upper bounds that can be too strong. The second method gives weaker upper bounds, but will involve approximations on how the NTK eigenvectors evolve in the catapult phase.

Method 1

For the first method we claim that if the learning rate η lies in the range:

$$\frac{2}{\lambda_{\max}(H_{\alpha\beta,0})} < \eta \lesssim \frac{4}{\lambda_{\max}(\Omega_{\alpha\mu,\beta\nu})\theta_0^2}, \tag{138}$$

then the model is in the catapult phase. We will define Ω momentarily and will also explain why we use the approximate inequality ‘ \lesssim ’ above.

The update equation for θ_t^2 is:

$$\theta_{t+1}^2 = \theta_t^2 - \frac{4\eta}{D} \sum_{\alpha=1}^D z_{\alpha,t} \epsilon_{\alpha,t} + \sum_{\alpha,\beta=1}^D \frac{\eta^2}{D} \epsilon_{\alpha,t} \epsilon_{\beta,t} H_{\alpha\beta,t}. \tag{139}$$

To simplify this update equation we use that when $\eta\lambda_{\max}(H_{\alpha\beta,0}) > 2$ the output $z_{\alpha,t}$ grows exponentially quickly at early times in the direction of the top eigenvector of $H_{\alpha\beta,0}$, see (35). Here we will only need to know that $z_{\alpha,t}$ is large and will not need to know that it is aligned with the top eigenvector. We will write²²:

$$z_{\alpha,t} = \|z_{\alpha,t}\| \hat{z}_{\alpha,t}, \tag{140}$$

²² Here $\|z_{\alpha,t}\|^2 = \sum_{\alpha=1}^D z_{\alpha,t}^2$, i.e. it is the L_2 norm in sample-space.

where $\hat{z}_{\alpha,t}$ is a unit vector in sample space, i.e. $\sum_{\alpha=1}^D \hat{z}_{\alpha,t}^2 = 1$.

At time t_* such that $1 \ll t_* \ll \log(\zeta^{-1})$, the output becomes much larger than the labels, $\|z_{\alpha,t_*}\| \gg \|y_\alpha\|$, and the update equation for θ_t^2 simplifies to:

$$\theta_{t_*+1}^2 \approx \theta_{t_*}^2 + \frac{\eta}{D} \|z_{\alpha,t_*}\|^2 \left(\sum_{\alpha,\beta=1}^D \eta \hat{z}_{\alpha,t_*} \hat{z}_{\beta,t_*} H_{\alpha\beta,t_*} - 4 \right). \tag{141}$$

By taking $t_* \gg 1$ we are ensuring that z_{α,t_*} has a large norm and the approximation $\epsilon_{\alpha,t_*} \approx z_{\alpha,t_*}$ is valid. Moreover, since $t_* \ll \log(n)$ we also have $\|z_{\alpha,t_*}\| \ll \zeta^{-1}$ and can still use small ζ -perturbation theory. This means that the weight norm at t_* has not deviated significantly from its value at initialization, $\theta_{t_*}^2 - \theta_0^2 = O(1)$ where $\theta_0^2 = O(\zeta^{-2})$.

Next, we want to derive conditions on η such that after $t = t_*$, the weight norm θ_t^2 decreases by an $O(\zeta^{-2})$ amount. As in our previous analysis, we start by bounding the NTK $H_{\alpha\beta,t}$ in terms of the weight norm θ_t^2 . We will do this by expressing the NTK in terms of the meta-feature functions:

$$\sum_{\alpha,\beta=1}^D \hat{z}^\alpha \hat{z}^\beta H_{\alpha\beta} = \frac{\zeta^2}{D} \sum_{\alpha,\beta=1}^D \sum_{\mu,\nu,\rho=1}^n \hat{z}^\alpha \theta^\mu \psi_{\alpha,\mu\nu} \psi_{\beta,\nu\rho} \theta^\rho \hat{z}^\beta = \sum_{\alpha,\beta=1}^D \sum_{\mu,\rho=1}^n \hat{z}^\alpha \theta^\mu \Omega_{\alpha\mu,\beta\rho} \hat{z}^\beta \theta^\rho, \tag{142}$$

where we defined the matrix:

$$\Omega_{\alpha\mu,\beta\rho} = \frac{\zeta^2}{D} \sum_{\nu=1}^n \psi_{\alpha,\mu\nu} \psi_{\beta,\nu\rho}. \tag{143}$$

We can think of Ω as a matrix acting in $z \otimes \theta$ space, i.e. $\Omega \in \mathbb{R}^{(n \times d) \times (n \times d)}$. The matrix $\Omega_{\alpha\mu,\beta\rho}$ is also positive semi-definite in $z \otimes \theta$ space, which implies:

$$\sum_{\alpha,\beta=1}^D \hat{z}_t^\alpha H_{\alpha\beta} \hat{z}_t^\beta \leq \lambda_{\max}(\Omega_{\alpha\mu,\beta\rho}) \theta_t^2, \tag{144}$$

where we are treating Ω as a $(n \times d) \times (n \times d)$ -dimensional matrix and used that $\hat{z}_{\alpha,t}$ is a unit vector.

Using the approximate update equation (141) and the bound (144) we claim that if:

$$\eta < \frac{4}{\lambda_{\max}(\Omega_{\alpha\mu,\beta\rho}) \theta_{t_*}^2}, \quad \text{for } 1 \ll t_* \ll \log(\zeta^{-1}), \tag{145}$$

then the weight norm θ_t^2 decreases by an order $O(\zeta^{-2})$ during the catapult phase. The argument follows the usual inductive proof we have used so far. If the inequality (145) holds at some time t_* such that $1 \ll t_* \ll \log(\zeta^{-1})$, then from (141) the weight norm decreases:

$$\theta_{t_*+1}^2 < \theta_{t_*}^2. \tag{146}$$

Then, since the matrix Ω is constant, this means the inequality (145) continues to hold at time $t = t_* + 1$:

$$\eta < \frac{4}{\lambda_{\max}(\Omega_{\alpha\mu,\beta\rho})\theta_{t_*}^2} < \frac{4}{\lambda_{\max}(\Omega_{\alpha\mu,\beta\rho})\theta_{t_*+1}^2}, \quad \text{for } 1 \ll t_* \ll \log(\zeta^{-1}). \quad (147)$$

Therefore, if the inequality (145) holds, then while $\|z_{\alpha,t}\| \gg \|y_{\alpha,t}\|$ the weight norm θ_t^2 decreases. When $t = O(\log(\zeta^{-1}))$ we have $\|z_{\alpha,t}\| = O(\zeta^{-1})$ and the squared weights θ_t^2 receive negative $O(\zeta^{-2})$ corrections. Since θ_t^2 is a manifestly positive semi-definite quantity, eventually this process will terminate. This happens when $\|z_{\alpha,t}\| = O(1)$, so the updates to θ_t^2 become small and no longer have a definite sign²³. In practice, the output decreases from $\|z_{\alpha,t}\| = O(\zeta^{-1})$ to $\|z_{\alpha,t}\| = O(1)$ in an order-one number of steps.

When $\|z_{\alpha,t}\| = O(1)$ we have two possibilities. The first is that $\eta\lambda_{\max}(H_{\alpha\beta,t}) < 2$, in which case the model is in the lazy phase, we can apply small ζ -perturbation theory, and the loss converges. The other possibility is that $\eta\lambda_{\max}(H_{\alpha\beta,t}) > 2$, in which case $z_{\alpha,t}$ grows exponentially fast again. Following the arguments used in appendix B.4 to study positive labels in the two-layer ReLU MLP, we can argue that in this case there is a second catapult mechanism where θ_t^2 again receives $O(n)$, negative updates. The arguments below will be identical in form to the ones used in appendix B.4, so the reader can safely skip the remainder of this section.

Following the conventions of appendix B.4, we will denote the time at which the first catapult mechanism ends by t_1 , at which time $\|z_{\alpha,t_1}\| = O(1)$. Since θ_t^2 received negative $O(\zeta^{-2})$ updates during the first catapult mechanism we have:

$$|\theta_{t_*}^2 - \theta_{t_1}^2| = O(\zeta^{-2}), \quad (148)$$

$$\theta_{t_1}^2 < \theta_{t_*}^2. \quad (149)$$

Next, because the learning rate is still super-critical at time $t = t_1$, the output $z_{\alpha,t}$ grows exponentially fast and after an $O(1)$ number of steps becomes exponentially large again. We will denote the time at which $z_{\alpha,t}$ becomes exponentially large by t_2 and take $1 \ll t_2 - t_1 \ll \log(\zeta^{-1})$. If we take $t_2 - t_1$ to be large, but $O(1)$ (i.e. it does not scale with ζ^{-1}), then we have:

$$|\theta_{t_1}^2 - \theta_{t_2}^2| = O(1). \quad (150)$$

Our assumption that $1 \ll t_2 - t_1 \ll \log(\zeta^{-1})$ is crucial, the lower bound allows us to make the approximation $\epsilon_{\alpha,t_2} \approx z_{\alpha,t_2}$ while the upper bound allows us to use small ζ -perturbation theory to compute $\theta_{t_2}^2$ and argue it can change at most by an $O(1)$ amount when we evolve for an $O(1)$ number of steps. Together, (148)–(150) yield:

$$\theta_{t_*}^2 - \theta_{t_2}^2 < 0. \quad (151)$$

²³ The other possibility one can think of is that $\eta\sum_{\alpha\beta}\hat{z}_\alpha\hat{z}_\beta H_{\alpha\beta} - 4 \rightarrow 0$ so the updates to θ_t^2 become small while $\|z_\alpha\|$ remains large. However, this is not compatible with the bound (144) and the inequality (145) which imply $\eta\sum_{\alpha\beta}\hat{z}_\alpha\hat{z}_\beta H_{\alpha\beta} - 4$ is strictly negative and bounded away from 0.

Finally, (151) implies that the inequality (145) continues to hold at time $t = t_2$:

$$\eta < \frac{4}{\lambda_{\max}(\Omega_{\alpha\mu,\beta\rho})\theta_{t_2}^2}, \quad \text{for } 1 \ll t_2 - t_1 \ll \log(\zeta^{-1}). \quad (152)$$

The bound (152) implies that θ_t^2 will again receive $O(\zeta^{-2})$ negative corrections when $z_{\alpha,t} = O(\zeta^{-1})$ for the second time. Given the fact the NTK is a positive semi-definite quantity which is bounded from above by θ_t^2 , see (144), eventually these large decreases in θ_t^2 will drive the model to the lazy phase, $\eta\lambda_{\max}(H_{\alpha\beta,t}) < 2$.

We can now restate the condition (145) in terms of the weights at initialization. Because we assumed $t_* \ll \log(\zeta^{-1})$ we have $\theta_{t_*}^2 = \theta_0^2 + O(1)$. Therefore, we can rewrite (145) as:

$$\eta < \frac{4}{\lambda_{\max}(\Omega_{\alpha\mu,\beta\rho})\theta_{t_*}^2} = \frac{4}{\lambda_{\max}(\Omega_{\alpha\mu,\beta\rho})\theta_0^2} + O(\zeta^2), \quad \text{for } 1 \ll t_* \ll \log(\zeta^{-1}), \quad (153)$$

where we have $O(\zeta^2)$ because of the implicit ζ^2 in the definition of Ω , see (143). The $O(\zeta^2)$ correction is computable using small ζ -perturbation theory, but we will not give it here. Instead (153) is our final result and the order ζ^2 term in (153) is the source of the ' \lesssim ' in (138).

Method 2

While the previous method works, the window (138) can be very small or simply not exist. For this reason we will consider a different argument for the existence of the catapult phase in the pure quadratic model. We conjecture that if the learning rate is in the range:

$$\frac{2}{\lambda_{\max}(H_{\alpha\beta,0})} < \eta \lesssim \frac{4}{\zeta^2\theta_0^2\lambda_{\max}(\psi_{\text{eff}}^2)}, \quad (154)$$

and $\lambda_{\max}(\psi_{\text{eff}}^2)$ is $O(1)$, then the model is in the catapult phase. The effective meta-feature function is defined by:

$$\psi_{\text{eff}} = \frac{1}{\sqrt{D}} \sum_{\alpha=1}^D e_{\max,0}^{\alpha} \psi_{\alpha}, \quad (155)$$

where $e_{\max,0}^{\alpha}$ is the top eigenvector of the NTK at initialization.

To argue that the model is in the catapult phase, we follow the set-up of the previous section. Assuming $\eta\lambda_{\max}(H_{\alpha\beta,0}) > 2$, the output $z_{\alpha,t}$ grows exponentially fast at early times. At a time t_* such that $1 \ll t_* \ll \log(\zeta^{-1})$, the output z_{α,t_*} is exponentially large but we can still use small ζ -perturbation theory. The leading order solutions take the same form as before:

$$z_{\alpha,t_*} = c_{t_*}^{\max} e_{\alpha,0}^{\max} + O(\zeta), \quad (156)$$

$$c_{t_*}^{\max} = e^{(1-\eta\lambda_{\max}(H_{\alpha\beta,0}))t_*} c_0^{\max} + O(\zeta), \quad (157)$$

$$H_{\alpha\beta,t_*} = H_{\alpha\beta,0} + O(\zeta), \quad (158)$$

$$\boldsymbol{\theta}_{t^*}^2 = \boldsymbol{\theta}_0^2 + O(\zeta). \tag{159}$$

In contrast to the previous method, at this point we will assume that the eigenvectors of the NTK remain approximately static until $t = O(\log(\zeta^{-1}))$. If we make this approximation, then the update equation (139) becomes:

$$\boldsymbol{\theta}_{t^*+1}^2 = \boldsymbol{\theta}_{t^*}^2 + \frac{\eta}{D} (c_{t^*}^{\max})^2 \left(\eta \sum_{\alpha, \beta=1}^D e_{\alpha,0}^{\max} e_{\beta,0}^{\max} H_{\alpha\beta,t^*} - 4 \right). \tag{160}$$

The above update equation for $\boldsymbol{\theta}_t^2$ is essentially the same as the update equation in the pure quadratic model with a single datapoint, see (61). The important difference is that when we had a single data-point the NTK was simply a number while here we are studying a particular matrix element of the NTK.

Following our analysis of the model with a single datapoint given in appendix B.1, we can bound the NTK in terms of the meta-feature functions:

$$\sum_{\alpha, \beta=1}^D e_{\alpha,0}^{\max} e_{\beta,0}^{\max} H_{\alpha\beta,t^*} \leq \zeta^2 \boldsymbol{\theta}_{t^*}^2 \lambda_{\max}(\boldsymbol{\psi}_{\text{eff}}^2). \tag{161}$$

Using the bound (161), we can show that if the learning rate η obeys the following bound:

$$\eta < \frac{4}{\zeta^2 \boldsymbol{\theta}_{t^*}^2 \lambda_{\max}(\boldsymbol{\psi}_{\text{eff}}^2)} = \frac{4}{\zeta^2 \boldsymbol{\theta}_0^2 \lambda_{\max}(\boldsymbol{\psi}_{\text{eff}}^2)} + O(\zeta^2) \quad \text{for } 1 \ll t^* \ll \log(\zeta), \tag{162}$$

then the weight norm $\boldsymbol{\theta}_t^2$ receives negative corrections while $z_{\alpha,t}$ is large and aligned with the eigenvector $e_{\alpha,0}^{\max}$. The second assumption is crucial and underlies the difference between this method and the previous method.

Ignoring this subtlety for a moment, we know from appendix A that perturbation theory breaks down when $t = O(\log(\zeta^{-1}))$ and $\|z_{\alpha,t}\| = O(\zeta^{-1})$. If the learning rate η obeys the bound (162) and $z_{\alpha,t}$ is aligned with the top eigenvector of the NTK, then $\boldsymbol{\theta}_t^2$ decreases by an order $O(\zeta^{-2})$ amount when $\|z_{\alpha,t}\| = O(\zeta^{-1})$. We conjecture that this large decrease in $\boldsymbol{\theta}_t^2$ is sufficient to ensure that the model does not diverge.

The above argument is not rigorous because we do not have analytic control over the evolution of the NTK eigenvectors when perturbation theory breaks down near $t = O(\log(\zeta^{-1}))$. Specifically, the issue is that the bound (162) depends on the top eigenvector of the NTK at initialization through our definition of $\boldsymbol{\psi}_{\text{eff}}$. If the top eigenvector evolves significantly during training, and $z_{\alpha,t}$ becomes primarily aligned with this new direction in sample space, there is not a guarantee that (162) continues to hold with the new top eigenvector. That being said, we find numerically that the bound (162) is always sufficient to ensure convergence and in fact is generically stronger than necessary. See appendix D for more details.

C.2. Quadratic model with bias

Here we will extend the above analysis to include quadratic models with bias. The arguments in this section will largely be the same as ‘method 2’ above for the pure quadratic model. We will argue that if the learning rate lies in the range:

$$\frac{2}{\lambda_{\max}(H_{\alpha\beta,0})} < \eta \lesssim \frac{4}{2\phi_{\text{eff}}^2 + \zeta^2 \lambda_{\max}(\psi_{\text{eff}}^2) \left(\theta_0^2 + \frac{(\phi_{\text{eff}}^T \theta_0)^2}{\phi_{\text{eff}}^2} \right)}, \tag{163}$$

then the model converges. Here the effective (meta-)feature functions are defined by:

$$\phi_{\text{eff}} = \frac{1}{\sqrt{D}} \sum_{\alpha=1}^D e_{\alpha,0}^{\max} \phi_{\alpha}, \tag{164}$$

$$\psi_{\text{eff}} = \frac{1}{\sqrt{D}} \sum_{\alpha=1}^D e_{\alpha,0}^{\max} \psi_{\alpha}, \tag{165}$$

where $e_{\alpha,0}^{\max}$ is the top eigenvector of the NTK at initialization.

We will study the quadratic model with bias at time t_* such that $1 \ll t_* \ll \log(\zeta^{-1})$. In this regime we can use the results from small ζ -perturbation theory given in (156)–(159). If we assume that the eigenvectors of the NTK remain approximately static until $t = O(\log(\zeta^{-1}))$, then we can approximate the update equation for θ_t^2 (46) as:

$$\theta_{t_*+1}^2 = \theta_{t_*}^2 + \frac{\eta}{D} (c_{t_*}^{\max})^2 \left(\eta \sum_{\alpha,\beta=1}^D e_{\alpha,0}^{\max} e_{\beta,0}^{\max} H_{\alpha\beta,t_*} - 4 \right) + \frac{2\eta}{D} c_{t_*}^{\max} \phi_{\text{eff}}^T \theta_{t_*}. \tag{166}$$

Following our analysis for the same model with a single data-point, see appendix B.2, we will study the evolution of the quantity:

$$\theta_t^2 + \frac{\left(\sum_{\alpha=1}^D \hat{z}_{\alpha,t} \phi_{\alpha}^T \theta_t \right)^2}{\left(\sum_{\alpha=1}^D \hat{z}_{\alpha,t} \phi_{\alpha} \right)^2}. \tag{167}$$

Here $\hat{z}_{\alpha,t}$ is the unit vector in sample space parallel to $z_{\alpha,t}$. We will prove that under certain conditions on η , the quantity (167) decreases whenever the loss is large.

In the range $1 \ll t_* \ll \log(\zeta^{-1})$, the output z_{α,t_*} becomes aligned with the top eigenvector of the NTK and we have:

$$\theta_{t_*}^2 + \frac{\left(\sum_{\alpha=1}^D \hat{z}_{\alpha,t_*} \phi_{\alpha} \theta_{t_*} \right)^2}{\left(\sum_{\alpha=1}^D \hat{z}_{\alpha,t_*} \phi_{\alpha} \right)^2} = \theta_{t_*}^2 + \frac{(\phi_{\text{eff}}^T \theta_{t_*})^2}{\phi_{\text{eff}}^2} + O(1), \quad \text{for } 1 \ll t_* \ll \log(\zeta^{-1}). \tag{168}$$

J. Stat. Mech. (2025) 093406

The update equation for $\phi_{\text{eff}}^T \theta_t$ in the range $1 \ll t_* \ll \log(\zeta^{-1})$ is:

$$\phi_{\text{eff}}^T \theta_{t_*+1} = \phi_{\text{eff}}^T \theta_{t_*} - \frac{\eta}{\sqrt{D}} \sum_{\beta=1}^D \phi_{\text{eff}}^T \phi_{\beta} \epsilon_{t_*}^{\beta} = \phi_{\text{eff}}^T \theta_{t_*} - \eta c_{t_*}^{\max} \phi_{\text{eff}}^2 + O(\zeta), \quad (169)$$

and the update equation for $\frac{(\phi_{\text{eff}}^T \theta_t)^2}{\phi_{\text{eff}}^2}$ is:

$$\frac{(\phi_{\text{eff}}^T \theta_{t_*+1})^2}{\phi_{\text{eff}}^2} \approx \frac{(\phi_{\text{eff}}^T \theta_{t_*})^2}{\phi_{\text{eff}}^2} - 2\eta \phi_{\text{eff}}^T \theta_{t_*} + \eta^2 (c_{t_*}^{\max})^2 \phi_{\text{eff}}^2. \quad (170)$$

Together (166) and (170) give the update equation:

$$\begin{aligned} \theta_{t_*+1}^2 + \frac{(\phi_{\text{eff}}^T \theta_{t_*+1})^2}{\phi_{\text{eff}}^2} &\approx \theta_{t_*}^2 + \frac{(\phi_{\text{eff}}^T \theta_{t_*})^2}{\phi_{\text{eff}}^2} + \eta (c_{t_*}^{\max})^2 \\ &\times \left(\eta \left(\frac{1}{D} \sum_{\alpha, \beta=1}^D e_{\alpha,0}^{\max} e_{\beta,0}^{\max} H_{\alpha\beta,t_*} + \phi_{\text{eff}}^2 \right) - 4 \right). \end{aligned} \quad (171)$$

Finally, we can bound the size of the η^2 term in (171) in terms of the effective meta-feature function:

$$\begin{aligned} \frac{1}{D} \sum_{\alpha, \beta=1}^D e_{\alpha,0}^{\max} e_{\beta,0}^{\max} H_{\alpha\beta,t} + \phi_{\text{eff}}^2 &\leq 2\phi_{\text{eff}}^2 + \zeta^2 \lambda_{\max}(\psi_{\text{eff}}^2) \theta_t^2 \\ &\leq 2\phi_{\text{eff}}^2 + \zeta^2 \lambda_{\max}(\psi_{\text{eff}}^2) \left(\theta_t^2 + \frac{(\phi_{\text{eff}}^T \theta_t)^2}{\phi_{\text{eff}}^2} \right). \end{aligned} \quad (172)$$

To get the first inequality we used that ψ_{eff}^2 is a positive semi-definite matrix. The second inequality is trivial since we simply added a positive term. Finally, we find that if η obeys the following bound for $1 \ll t_* \ll \log(\zeta^{-1})$:

$$\begin{aligned} \eta &\leq \frac{4}{2\phi_{\text{eff}}^2 + \zeta^2 \lambda_{\max}(\psi_{\text{eff}}^2) \left(\theta_{t_*}^2 + \frac{(\phi_{\text{eff}}^T \theta_{t_*})^2}{\phi_{\text{eff}}^2} \right)} \\ &= \frac{4}{2\phi_{\text{eff}}^2 + \zeta^2 \lambda_{\max}(\psi_{\text{eff}}^2) \left(\theta_0^2 + \frac{(\phi_{\text{eff}}^T \theta_0)^2}{\phi_{\text{eff}}^2} \right)} + O(\zeta), \end{aligned} \quad (173)$$

then the quantity $\theta_t^2 + \frac{(\phi_{\text{eff}}^T \theta_t)^2}{\phi_{\text{eff}}^2}$ receives large $O(\zeta^{-2})$ negative corrections during the catapult phase. The argument is essentially the same as ‘method 2’ of appendix C.1. If the inequality (173) holds at $t = t_*$, then from the approximate update equation (171) we see that $\theta_t^2 + \frac{(\phi_{\text{eff}}^T \theta_t)^2}{\phi_{\text{eff}}^2}$ decreases after one step of gradient descent. This then implies the inequality (173) continues to hold at $t = t_* + 1$. This procedure continues until the small

ζ expansion breaks down when $z_t = O(\zeta^{-1})$, at which point $\theta_t^2 + \frac{(\phi_{\text{eff}}^T \theta_t)^2}{\phi_{\text{eff}}^2}$ decreases by an $O(\zeta^{-2})$ amount. Given that the NTK is bounded from above in terms of this quantity by an order-one amount, see (172), we expect the large decrease in $\theta_t^2 + \frac{(\phi_{\text{eff}}^T \theta_t)^2}{\phi_{\text{eff}}^2}$ will drive the model back to the lazy phase, $\eta \lambda_{\max}(H_{\alpha\beta,t}) < 2$.

As with our previous version of the above argument, method 2 of appendix C.1, this argument is not completely rigorous because we do not have analytic control over the evolution of the NTK's eigenvectors during training. To make the argument rigorous we need to understand the evolution of ϕ_{eff} and ψ_{eff} when perturbation theory breaks down near $t = O(\log(\zeta^{-1}))$. That said, we empirically observe in appendix D that the bound (173) is stronger than necessary to ensure convergence at super-critical learning rates.

C.3. Homogeneous, two-layer, MLPs

As our final example, we will study two-layer MLPs with a scale-invariant activation function and multiple datapoints, see (53) and (54) for $z_{\alpha,t}$ and $H_{\alpha\beta,t}$ in this model. In practice, we will find our bounds for this MLP will only be non-trivial for uni-dimensional data, $d = 1$, although we believe there should be a way to extend our type of arguments to general d .

The weight norm in this model is given by:

$$\theta_t^2 = \text{Tr}(\mathbf{U}_t^T \mathbf{U}_t) + \mathbf{v}_t^2. \tag{174}$$

Under gradient descent it evolves as:

$$\theta_{t+1}^2 = \theta_t^2 - \frac{4\eta}{D} \sum_{\alpha=1}^D \epsilon_{\alpha,t} z_{\alpha,t} + \frac{\eta^2}{D} \sum_{\alpha,\beta=1}^D \epsilon_{\alpha,t} \epsilon_{\beta,t} H_{\alpha\beta,t}. \tag{175}$$

If we assume $\eta \lambda_{\max}(H_{\alpha\beta,0}) > 2$ then the output $z_{\alpha,t}$ becomes exponentially large and in an order $O(n^0)$ number of steps we can make the approximation $\epsilon_{\alpha,t} = z_{\alpha,t} - y_\alpha \approx z_{\alpha,t}$. Then for $1 \ll t^* \ll \log(n)$, we again have $\|y_\alpha\| \ll \|z_{\alpha,t^*}\| \ll 1/\sqrt{n}$ and the θ^2 update equation can be approximated as:

$$\theta_{t+1}^2 \approx \theta_{t^*}^2 + \frac{\eta}{D} \|z_{\alpha,t^*}\|^2 \left(\eta \sum_{\alpha,\beta=1}^D \hat{z}_{\alpha,t^*} \hat{z}_{\beta,t^*} H_{\alpha\beta,t^*} - 4 \right), \tag{176}$$

where $\hat{z}_{\alpha,t}$ is a unit vector in sample space and $\|z_{\alpha,t^*}\|$ is the L_2 norm in sample space.

To derive a bound on η we need to bound the NTK in terms of the weight norm:

$$\begin{aligned} \sum_{\alpha,\beta=1}^D \hat{z}_{\alpha,t} \hat{z}_{\beta,t} H_{\alpha\beta,t} &\leq \frac{a_+^2}{nD} \sum_{\alpha\beta} \hat{z}_{\alpha,t} \hat{z}_{\beta,t} (\mathbf{x}_\alpha^T \mathbf{U}_t^T \mathbf{U}_t \mathbf{x}_\beta + \mathbf{x}_\alpha^T \mathbf{x}_\beta \mathbf{v}_t^2) \\ &\leq \frac{a_+^2}{nD} \sum_{\alpha,\beta=1}^D \hat{z}_{\alpha,t} \hat{z}_{\beta,t} \mathbf{x}_\alpha^T \mathbf{x}_\beta (\text{Tr}(\mathbf{U}_t^T \mathbf{U}_t) + \mathbf{v}_t^2) \end{aligned}$$

$$\begin{aligned} &\leq \frac{a_+^2}{nD} \sum_{\alpha, \beta=1}^D \lambda_{\max}(\mathbf{x}_\alpha^\top \mathbf{x}_\beta) (\text{Tr}(\mathbf{U}_t^\top \mathbf{U}_t) + \mathbf{v}_t^2) \\ &= \frac{a_+^2}{nD} \lambda_{\max}(\mathbf{x}_\alpha^\top \mathbf{x}_\beta) \boldsymbol{\theta}_t^2. \end{aligned} \tag{177}$$

To obtain the first inequality we used our assumption $0 < a_0 \leq a_+$. The second inequality then follows from a Cauchy–Schwarz inequality. The final inequality follows from the fact $\mathbf{x}_\alpha^\top \mathbf{x}_\beta$ is a positive semi-definite matrix and \hat{z}_α is a unit vector in sample space²⁴.

We can then show that if at time t^* , such that $1 \ll t^* \ll \log(n)$, the learning rate η obeys the bound:

$$\frac{2}{\lambda_{\max}(H_{\alpha\beta,0})} < \eta < \frac{4nD}{a_+^2 \lambda_{\max}(\mathbf{x}_\alpha^\top \mathbf{x}_\beta) \boldsymbol{\theta}_{t^*}^2} = \frac{4nD}{a_+^2 \lambda_{\max}(\mathbf{x}_\alpha^\top \mathbf{x}_\beta) \boldsymbol{\theta}_0^2} + O(1/n), \tag{180}$$

then the model is in the catapult phase. The argument for (180) is identical in form to the argument given in appendix C.1, so we will not repeat it here. Instead we will use that the update equation for $\boldsymbol{\theta}_t^2$ in the two-layer, homogeneous MLP (175) has the same form as the $\boldsymbol{\theta}_t^2$ update equation in the pure quadratic model (139) and also that the upper bound for the NTK (177) has the same form as the same bound in the pure quadratic model (144) with the replacement:

$$\lambda_{\max}(\Omega_{\alpha\mu,\beta\rho}) \implies \frac{a_+^2}{nD} \lambda_{\max}(\mathbf{x}_\alpha^\top \mathbf{x}_\beta). \tag{181}$$

Then if we make the replacement (181) in our result for the pure quadratic model, (153), we find the result (180).

C.4. Justifications for the simplifications

The above analysis of the training dynamics under the regime of small- ζ employs several controlled approximations to simplify the proofs. Although some of these approximations are heuristic, they remain justified through a combination of perturbation theory and assumptions consistently applied throughout our work, specifically the condition $\zeta^{-1} \gg |y_\alpha|$ for all α . We summarize the main approximations below:

²⁴ Note that the second inequality is generically very weak because we used an inequality of the form:

$$\mathbf{e}^\top \mathbf{U}^\top \mathbf{U} \mathbf{e} \leq \mathbf{e}^\top \mathbf{e} \text{Tr}(\mathbf{U}^\top \mathbf{U}), \tag{178}$$

when we could have used the stronger inequality:

$$\mathbf{e}^\top \mathbf{U}^\top \mathbf{U} \mathbf{e} \leq \mathbf{e}^\top \mathbf{e} \lambda_{\max}(\mathbf{U}^\top \mathbf{U}), \tag{179}$$

since $\mathbf{U}^\top \mathbf{U}$ is a positive semi-definite matrix. The fact the bound (178) is weak implies our final upper bounds on η may be stronger than necessary to ensure convergence. However, using the bound (178) is a necessary evil if we want to bound the NTK in terms of the weight norm $\boldsymbol{\theta}_t^2$. To derive weaker bounds on η it may be necessary to bound the NTK in a different manner.

- 1) To derive equation (141), we made the approximation $\epsilon_{\alpha,t^*} \approx z_{\alpha,t^*}$ when $1 \ll t^* \ll \log(\zeta^{-1})$. As shown in appendix A, $z_{\alpha,t}$ grows exponentially in this early phase, ensuring that the approximation $\epsilon_{\alpha,t^*} = z_{\alpha,t^*} - y_\alpha \approx z_{\alpha,t^*}$ becomes increasingly accurate as t^* and ζ^{-1} grow, provided the relationship $1 \ll t^* \ll \log(\zeta^{-1})$ is maintained.
- 2) In establishing the sharp inequality (145) within appendix C.1, we initially bounded the expression in terms of $\theta_{t^*}^2$ and subsequently converted it into a bound on θ_0^2 . This conversion relies on small- ζ perturbation theory to argue that the difference $\theta_{t^*}^2 - \theta_0^2$ is $O(1)$, as detailed in equation (47). The outcome is captured clearly in equation (153). Such approximations closely follow analogous procedures used in prior work (Lewkowycz *et al* 2020, Zhu *et al* 2022).
- 3) Lastly, we considered the case where gradient descent induces $z_{\alpha,t}$ to decrease from $\Theta(\zeta^{-1})$ to $o(\zeta^{-1})$, while maintaining a super-critical learning rate. Under these conditions, the model effectively reverts to its initial dynamic regime, validating the reuse of approximation (1) above. However, the key distinction here is a significant reduction in θ_t^2 by an amount of $O(\zeta^{-2})$, allowing the inequality analogous to equation (145) to remain applicable in analyzing the second ‘‘catapult mechanism.’’

Furthermore, We also realize the time-scale t_1 can be defined more clearly as the first time t after the catapult mechanism has started such that θ_t^2 receives a positive, additive correction. This is almost equivalent to the definition we gave since θ_t^2 can only receive a positive correction when $\|z_{\alpha,t}\| = O(1)$ but the converse is not necessarily true.

These approximations align with the spirit of previous theoretical analyses (Lewkowycz *et al* 2020, Zhu *et al* 2022), particularly concerning NTK approximations during early and peak phases. Besides, our application of the small- ζ expansion in a newly introduced time-scale ($1 \ll t^* \ll \log(\zeta^{-1})$), coupled with the identification and analysis of double-peaked loss curves (e.g. figures 2a and 9a), are our contributions beyond existing studies.

Appendix D. More Experiments

In this appendix we will explain our experimental set-up and present new numerical results. All experiments were carried out in PyTorch (Paszke *et al* 2019) and the models were trained using full-batch gradient descent. In figures 4–16 the different colored lines in plots (a)-(c), which give the time evolution of the training loss L_t , the weight norm θ_t^2 , and $\eta\lambda_{\max}(H_{\alpha\beta,t})$, correspond to different initial choices of $\eta\lambda_{\max}(H_{\alpha\beta,0})$. In addition, the vertical dashed lines (blue or red) in our plots of $\|\eta H_{\alpha\beta,\infty}\|$, the weight norm θ_∞^2 , and the generalization loss correspond to the theoretical predictions of appendices B and C. That is, the methods of appendices B and C imply the models should converge for super-critical learning rates to the left of those vertical lines. For each experiment we fix a random seed in order to compare different learning rates with a fixed weight initialization.

D.1. Quadratic model

D.1.1. Definitions. Here we will give more details about how we set up the quadratic model,

$$z_\alpha = \phi_\alpha^T \boldsymbol{\theta} + \frac{\zeta}{2} \boldsymbol{\theta}^T \boldsymbol{\psi}_\alpha \boldsymbol{\theta}. \tag{182}$$

We focus on the quadratic model with bias, but the pure quadratic model can be recovered by setting $\phi = 0$. In the quadratic model with bias we need to impose:

$$\boldsymbol{\psi}_\alpha \boldsymbol{\phi}_\beta = 0, \quad \forall \alpha, \beta. \tag{183}$$

We will impose this by assuming that the feature and meta-feature functions span different subspaces of \mathbb{R}^n . That is, we split the indices $\mu = \{1, \dots, n\}$ into two sets $A = \{1, \dots, m\}$ and $B = \{m + 1, \dots, n\}$ so that $\psi_{\alpha, \mu_i \mu_j} \neq 0$ only if $i, j \in A$ and $\phi_{\alpha, \mu_i} \neq 0$ only if $i \in B$. We will define the meta-feature hidden dimension to be $n_\psi = m$ and the feature hidden dimension to be $n_\phi = n - m$. In addition, we will assume $\phi_\alpha^2 = O(1)$ and $\zeta^2 = 1/m$ so that z_α is $O(1)$ at initialization when the eigenvalues of $\boldsymbol{\psi}_\alpha$ are $O(1)$. At initialization the weights $\theta_{\mu, t=0}$ are all drawn from $\mathcal{N}(0, 1)$.

We will consider the following feature and meta-feature functions:

$$\boldsymbol{\phi}(\mathbf{x}_\alpha) = \mathbf{U} \mathbf{x}_\alpha, \tag{184}$$

$$\boldsymbol{\psi}(\mathbf{x}_\alpha) = g \left(\sum_{i=1}^d \mathbf{W}^i x_{\alpha, i} \right), \tag{185}$$

where $\mathbf{x}_\alpha \in \mathbb{R}^d$, $\mathbf{U} \in \mathbb{R}^{n \times d}$, and $\mathbf{W}^i \in \mathbb{R}^{n \times n}$ for each i . Here g is a generic activation function. The matrices \mathbf{U} and \mathbf{W}^i are fixed throughout training. We will draw the components of \mathbf{U} from $\mathcal{N}(0, 1)$. For \mathbf{W}^i we will parameterize it so we can tune its eigenvalues by hand. Specifically, as in (18), we take:

$$W_{\mu\nu}^i = \sum_{\sigma=1}^n \lambda_\sigma^i q_{\sigma\mu}^i q_{\sigma\nu}^i, \tag{186}$$

where λ_σ^i are the eigenvalues of $W_{\mu\nu}^i$ for a fixed i and $q_{\sigma\mu}^i$ are the corresponding eigenvectors. We will define $\mathbf{q}^i \in \mathbb{R}^{n \times n}$ as:

$$\mathbf{q}^i = e^{\mathbf{B}^i}, \tag{187}$$

where \mathbf{B}^i is an anti-symmetric matrix whose independent components are drawn from $\mathcal{N}(0, 1)$. In (187) we are taking the matrix exponential, and not the component-wise exponential, and using that the exponential of an anti-symmetric matrix gives an orthogonal matrix (more precisely it gives a matrix in $\text{SO}(N)$).

We will train the quadratic model on three types of datasets:

1. The toy dataset $(x, y) = (1, 0)$.

2. Random datasets where we draw \mathbf{x}_α and y_α from cubes of the appropriate dimension, $\mathbf{x}_\alpha \in \mathcal{U}([-k, k]^d)$ and $y_\alpha \in \mathcal{U}([-k, k])$.
3. Teacher-student set-ups where we train a student quadratic model to recover the predictions of a teacher quadratic model. Here the datapoints are drawn uniformly from a d -dimensional cube and the labels are produced from the teacher model.

To explain the teacher-student set-up, we generate the labels using teacher (meta-)feature functions:

$$y_\alpha = \phi_{\text{teacher},\alpha}^\top \boldsymbol{\theta}_{\text{teacher}} + \frac{\zeta_{\text{teacher}}}{2} \boldsymbol{\theta}_{\text{teacher}}^\top \boldsymbol{\psi}_{\text{teacher},\alpha} \boldsymbol{\theta}_{\text{teacher}}, \tag{188}$$

and the predictions are generated by student (meta-)feature functions:

$$z_\alpha = \phi_{\text{student},\alpha}^\top \boldsymbol{\theta}_{\text{student}} + \frac{\zeta_{\text{student}}}{2} \boldsymbol{\theta}_{\text{student}}^\top \boldsymbol{\psi}_{\text{student},\alpha} \boldsymbol{\theta}_{\text{student}}. \tag{189}$$

The student (meta-)feature functions are defined by projecting the corresponding teacher (meta-)feature functions:

$$\phi_{\text{student}}(\mathbf{x}_\alpha) = \mathbf{Q} \phi_{\text{teacher}}(\mathbf{x}_\alpha), \tag{190}$$

$$\boldsymbol{\psi}_{\text{student}}(\mathbf{x}_\alpha) = \mathbf{Q} \boldsymbol{\psi}_{\text{teacher}}(\mathbf{x}_\alpha) \mathbf{Q}^\top, \tag{191}$$

where \mathbf{Q} is a projector which maps from the teacher to the student hidden dimension, $\mathbb{R}^{n_{\text{teacher}}} \rightarrow \mathbb{R}^{n_{\text{student}}}$. The two ζ 's are defined independently and chosen such that $z_{\alpha,0}, y_\alpha = O(1)$.

D.1.2. Pure quadratic model. In figure 4 we plot the results for a pure quadratic model trained with a random dataset. We take $\mathbf{x}_\alpha \sim \mathcal{U}([1/2, 1/2]^2)$ (i.e. the interior of a unit, two-dimensional cube) and the labels are $y_\alpha \sim \mathcal{U}([-1/2, 1/2])$. We take the activation function g in (185) to be the identity function. We also take the weights $\boldsymbol{\theta} \in \mathbb{R}^{500}$. We draw the positive eigenvalues of \mathbf{W}^i from $\mathcal{U}([.9, 1.1])$ and take the negative eigenvalues to be exactly -1 times the positive eigenvalues. Finally, we train the model for 100 epochs.

We observe this model has all the expected features of the catapult phase, the loss has the characteristic spike at early times and both the top eigenvalue of the NTK and the weight norm decrease significantly. The blue and red dashed lines in figures 4(d) and (e) correspond to the predictions from method 1 and 2 of appendix C.1, respectively. Note that here the prediction from method 1 is trivial. In order to get a non-trivial result we need the red line to sit to the right of $\eta \lambda_{\max}(H_{\alpha\beta,0}) = 2$, see the bound in (138). Therefore, we cannot use method 1 here to argue that the model has a catapult phase. On the other hand, the bound from method 2 is non-trivial because the blue dashed line sits to the right of $\eta \lambda_{\max}(H_{\alpha\beta,0}) = 2$, see (154). The results of figure 4 illustrate how our bounds are sufficient, but not necessary, for the existence of the catapult phase.

In figure 5 we plot the results for a pure quadratic model trained using the teacher student set-up. We take the teacher meta-feature function to have rank 500 and the student meta-feature function to have rank 400. In order to produce a non-linear function, in (185) we take the activation function g to be the tanh function. We also take

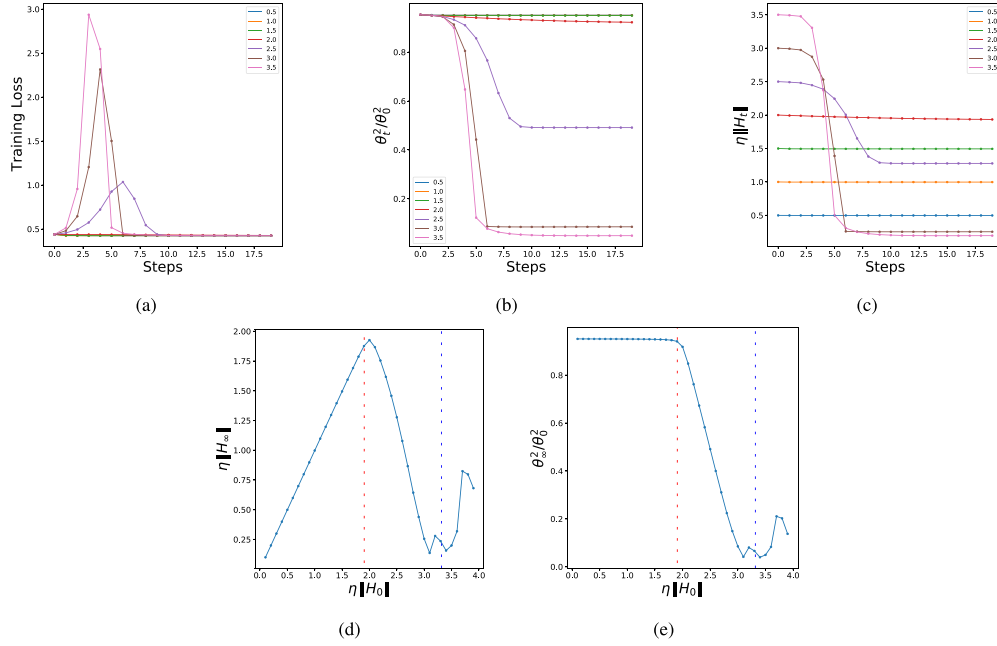


Figure 4. Results for the pure quadratic model trained on random two-dimensional data and one-dimensional labels. The plots are the same as in figure 1: (a)–(c) give the time evolution of the loss, the weight norm, and the NTK while (d)–(e) give the final value of the NTK and weight norm as a function of the (normalized) learning rate. The only difference is we have one dashed red line and one dashed blue line corresponding to predictions from method 1 and method 2 of appendix C.1, respectively.

the eigenvalues of \mathbf{W}^i to be ± 1 for all i . We draw the datapoints $x_\alpha \sim \mathcal{U}([-1/2, 1/2])$ and take the training and test set to have size 32 and 1000, respectively.

Here we observe that both method 1 (the red dashed line) and method 2 (the blue dashed line) of appendix C.1 give non-trivial predictions. Moreover, our experimental results match our theoretical predictions since the model converges for super-critical learning rates to the left of both dashed lines in figures 5(d)–(f). For this experiment the predictions from both methods agree to a high degree of precision and are virtually indistinguishable on the plots.

D.1.3. Quadratic model with bias. In figure 6 we summarize the results for a quadratic model with bias trained on the toy dataset $(x, y) = (1, 0)$. Specifically, we take the meta-feature hidden dimension to be $n_\psi = 100$ and the feature hidden dimension to be $n_\phi = 10$. For this model we take the activation function in (185) to be the identity function, $g(x) = x$. Finally, we take the eigenvalues of the meta-feature function to be $\lambda_i(\psi) = \pm 1$.

In this model we observe the expected behavior from our theoretical analysis. The blue, vertical dashed line in figures 6(d) and (e) correspond to the theoretical prediction given in (77). We observe that the model does indeed converge for super-critical learning




Figure 5. Results for the pure quadratic model trained using a teacher-student set-up.

rates to the left of this line. The model also converges for a small interval to the right of this line, which means our sufficiency condition can likely be weakened.

In figure 7 we perform a similar experiment, but for the quadratic model with bias trained using a teacher-student set-up. We take the inputs to be $1d$ and draw them from $x_\alpha \sim \mathcal{U}([-1/2, 1/2])$. We take the training set and test set to have size 32 and 1000, respectively. The teacher meta-feature function has rank $n_\psi = 200$ and the teacher feature functions have dimension $n_\phi = 20$. We take the student meta-feature function to have rank $n_\psi = 150$ and the student feature functions to have dimension $n_\phi = 10$. In (185) we take $g(x) = \tanh(x)$. Finally in (185) we take $\lambda_a(\mathbf{W}^i) = \pm 1$ for all a and i .

In figure 7 the vertical, dashed lines corresponds to the prediction from (163). We see agreement between the theoretical and experimental results since the model converges for super-critical learning rates to the left of the vertical lines. As with our other experiments, we observe that the bounds we derive are sufficient, but not necessary, to ensure convergence.

D.2. Generic two-layer homogeneous MLPs

Here we will study the two-layer homogeneous net trained on the toy dataset $(x, y) = (1, 0)$ and on random datasets. These examples will be simple enough to illustrate the learning dynamics of these models in the catapult phase. In the next section we will study the evolution of the weight norm in ReLU MLPs when trained on more realistic datasets.

In figure 8 we show the results for the two-layer MLP trained on the toy dataset $(x, y) = (1, 0)$. For the activation function we take $a_+ = 1$ and $a_- = 3/4$. We also take

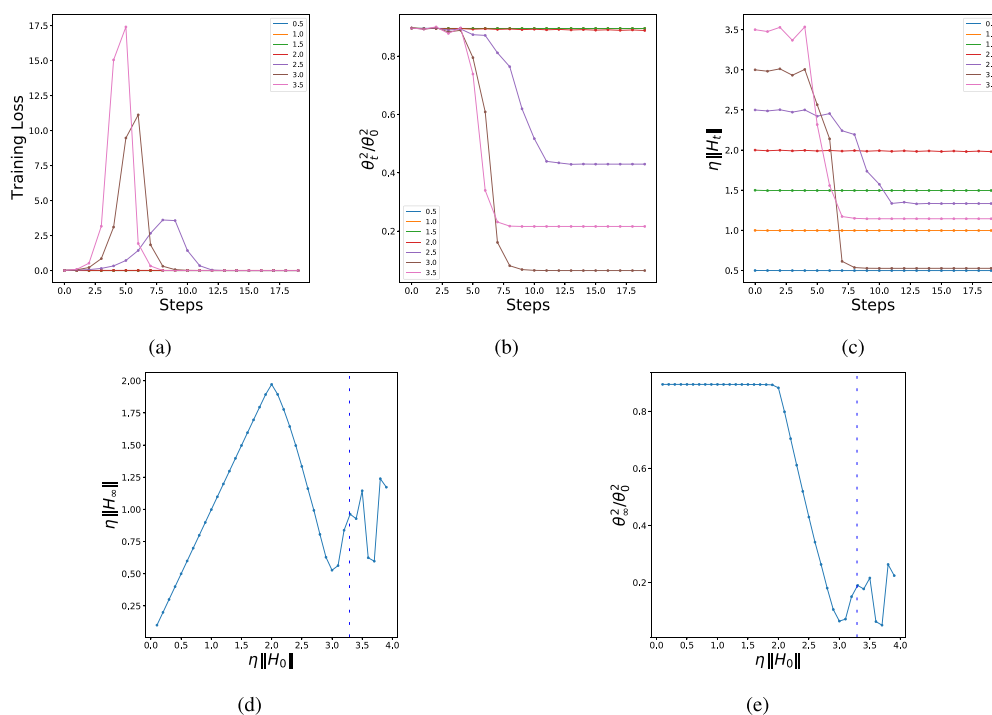


Figure 6. Results for the quadratic model with bias trained on the toy dataset $(x, y) = (1, 0)$.

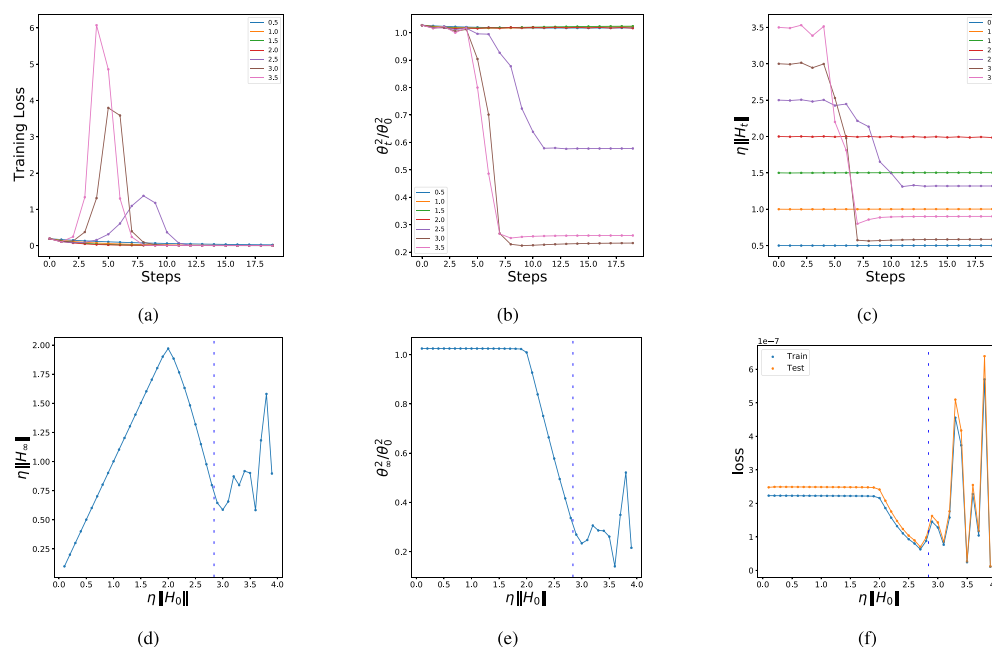


Figure 7. Results for the quadratic model with bias trained using a teacher-student set-up.

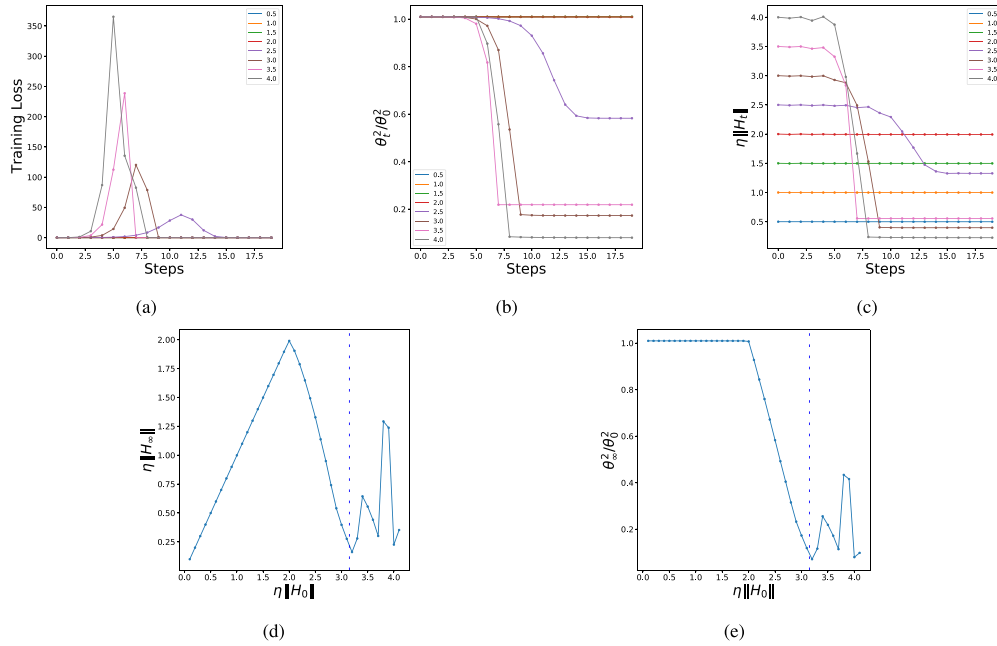


Figure 8. Results for the homogeneous net with $a_+ = 1$ and $a_- = 3/4$ trained on the toy dataset.

the width of the hidden layer to be 1024 and train the model until it converges. In comparison to figure 2, we can note that our theoretical prediction (95), which corresponds to the dashed blue lines in figures 8(d) and (e), has moved to the right. This result is not surprising since we derived the bound on the learning rate (95) using the upper bound on the NTK given in (97), which we reproduce below:

$$H_t \leq \frac{a_+^2}{n} \theta_t^2. \tag{192}$$

In general, the closer this upper bound is to being saturated, the weaker the upper bound (95) on η will be. The upper bound (192) is saturated when $a_+ = a_-$, in which case the bound on η (95) becomes optimal and agrees with the bound derived in Lewkowycz *et al* (2020) for the two-layer linear net. On the other hand, as we decrease a_- the NTK decreases, so the bound (192) becomes less tight, which in turn causes the upper bound on η to decrease.

Next, we test the prediction of appendix C.3, see (180), by studying homogeneous MLPs on random data. Here we will set $a_+ = 1$ and $a_- = 1/2$ and set the width to be $n = 1024$. We will draw both the datapoints x_α and the labels y_α from $\mathcal{U}([-1/2, 1/2])$. We take the training set to have size 32 and train the model for 1000 epochs. The results are summarized in figure 9. Here we observe that our upper bound is close to 2.25. Although our result is still non-trivial (it indicates that the model can converge for super-critical learning rates) the fact the model empirically converges up to $\eta \lambda_{\max}(H_{\alpha\beta,0}) = 4$ indicates there is room for improvement.

We can also note that most of the results in figure 9 are qualitatively similar to the results in figure 2, indicating that adding multiple datapoints does not significantly

alter the training dynamics, at least for random data. One difference however is that in figure 2(e) the model converged for learning rates up to $\eta\lambda_{\max}(H_{\alpha\beta,0}) = 4.5$, while here the model stops converging around $\eta\lambda_{\max}(H_{\alpha\beta,0}) = 4$. This indicates that where the catapult phase ends and the divergent phase begins must be data-dependent.

Finally, we will study the convergence of a two-layer ReLU net trained on a single datapoint in order to test the predictions of appendix B.4. We take the width of the model to be $n = 1024$ and the dataset to be $(x, y) = (4, 2)$. The results are summarized in figure 10. We observe that our theoretical predictions, the dashed vertical lines, again agree with the data since the model converges to the left of the line $\eta\lambda_{\max}(H_{\alpha\beta,0}) = 4$. In particular, θ_t^2 receives large negative updates when $\eta\lambda_{\max}(H_{\alpha\beta,0}) \leq 4$, but can receive positive corrections when $\eta\lambda_{\max}(H_{\alpha\beta,0}) \approx 4.5$, see figure 10(b). We can also note that around this value, $\eta\lambda_{\max}(H_{\alpha\beta,0}) \approx 4.5$, the model becomes trivial (the NTK vanishes) and the neurons die. A similar phenomenon was also observed in Lewkowycz *et al* (2020) for two-layer ReLU nets trained on one data-point, although they found the model became trivial around $\eta\lambda_{\max}(H_{\alpha\beta,0}) \approx 12$.

D.3. ReLU MLPs for images

D.3.1. Architectures and datasets. In this appendix we present results for ReLU MLPs without bias. The networks have the form:

$$z_\alpha = \frac{1}{n^{\frac{k+1}{2}}} \mathbf{v}^T \sigma_{\text{ReLU}}(\mathbf{W}_1 \sigma_{\text{ReLU}}(\dots \mathbf{W}_k \sigma_{\text{ReLU}}(\mathbf{U} \mathbf{x}_\alpha))), \tag{193}$$

where $\mathbf{U} \in \mathbb{R}^{n \times d}$, $\mathbf{W}_i \in \mathbb{R}^{n \times n}$, $\mathbf{v} \in \mathbb{R}^n$, and $\mathbf{x}_\alpha \in \mathbb{R}^d$. In practice we will set $k = 0$ or $k = 1$, i.e. we consider two- and three-layer networks. The components of all the weights are drawn from $\mathcal{N}(0, 1)$. We set the width to be $n = 1024$ for both $k = 0$ and $k = 1$.

We will train these MLPs on three datasets, which are two-class versions of MNIST, FMNIST, and CIFAR-10. For MNIST the model is trained to distinguish images of ‘0’ and ‘1’, for FMNIST it is trained to distinguish T-shirts and trousers, and on CIFAR-10 it is trained to distinguish airplanes and automobiles. For all cases we give the first image a label of -1 , the second image a label of 1 , and train the model to minimize the MSE. We take the training set to consist of 128 images from the original training set and take the test set to consist of all of the corresponding images in the original test set. Finally, we train the models until the training loss changes by less than 10^{-8} .

D.3.2. Two layers. Here we will study two-layer MLPs trained on the two-class versions of MNIST, FMNIST and CIFAR-10. The primary results for the two-class version of MNIST are given in figure 3. However, in figure 3 it is difficult to see how the weight norm behaves for $\eta\lambda_{\max}(H_{\alpha\beta,0}) \lesssim 4$, so in figure 11 we plot the evolution of θ_t^2 zoomed in to this region. We observe the expected behavior, for $\eta\lambda_{\max}(H_{\alpha\beta,0}) \lesssim 4$ the weight norm decreases during training, but for larger learning rates it receives large, positive corrections.

The results for the two-layer MLP trained on the two-class version of FMNIST are given in figure 12. For the most part, the results are qualitatively the same as for MNIST. One qualitative difference is that in figure 12(d) we observe the final value of $\eta\lambda_{\max}(H_{\alpha\beta,t})$ converges to 2 for almost all learning rates in the catapult phase. This is a

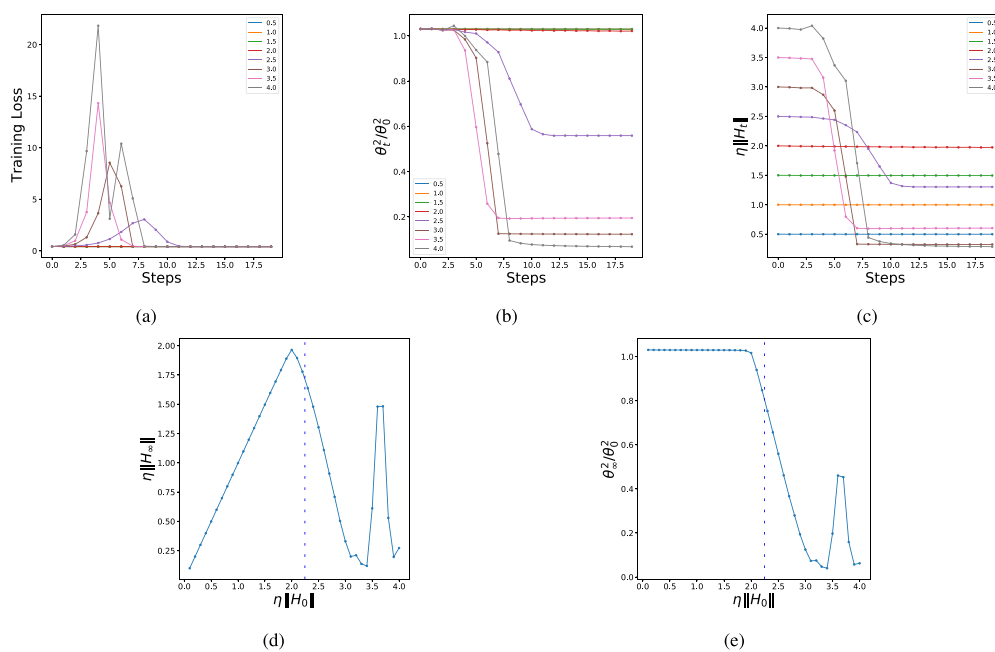


Figure 9. Results for the homogeneous net with $a_+ = 1$ and $a_- = 1/2$ trained on random data.

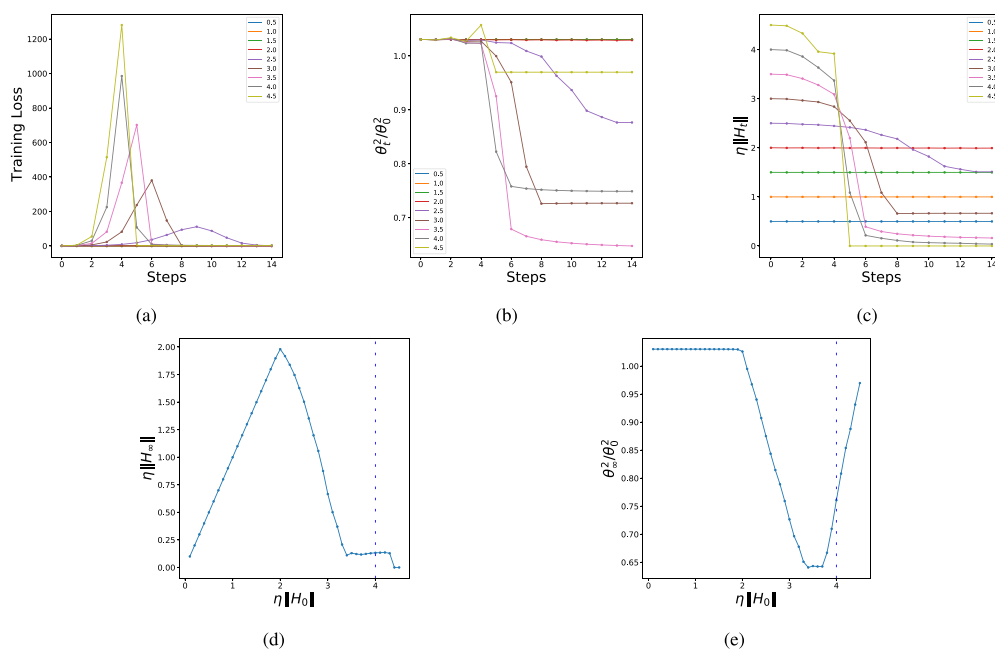


Figure 10. Results for a two-layer ReLU net trained on the dataset $(x, y) = (4, 2)$.




Figure 11. Evolution of θ_t^2 in a two-layer ReLU MLP trained on the two-class version of MNIST. Here we have zoomed into the region $\eta\lambda_{\max}(H_{\alpha\beta,0}) \leq 5$ to see how the evolution in θ_t^2 changes qualitatively as we increase the learning rate. For the evolution of these quantities for larger learning rates see figure 3.

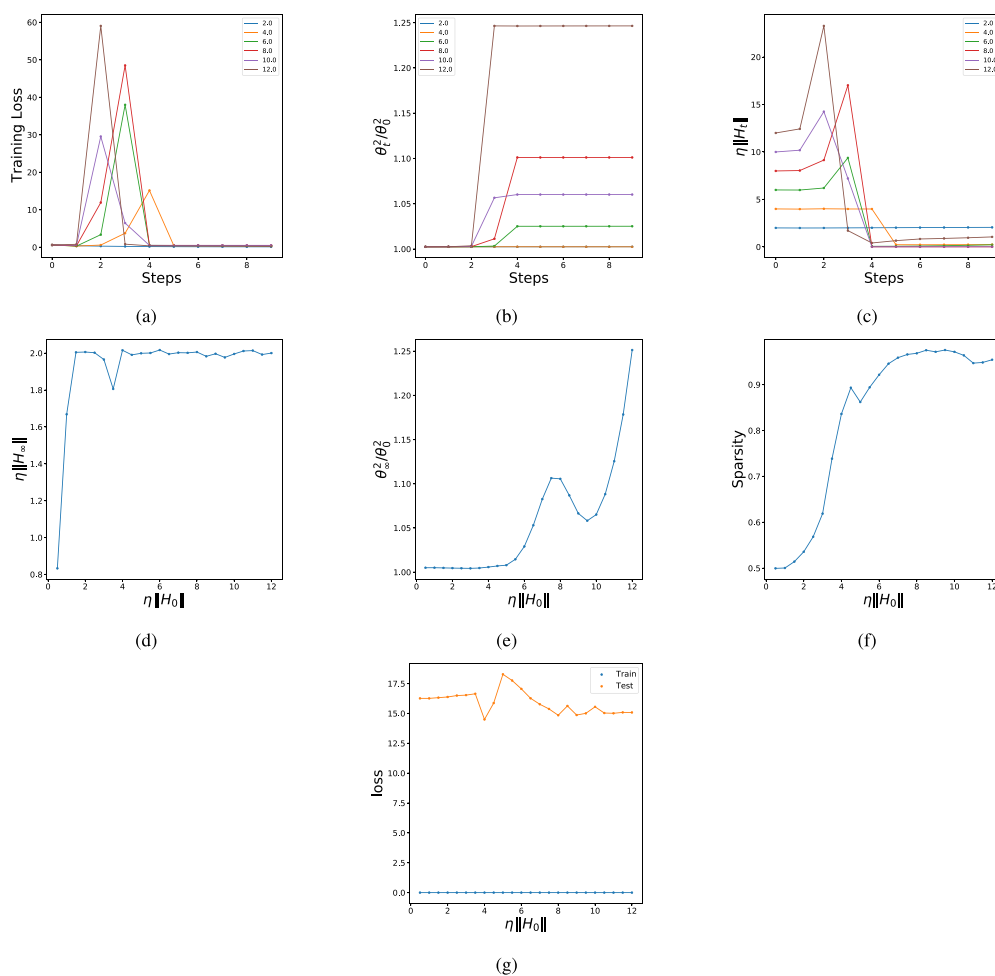


Figure 12. Results for a two-layer ReLU net trained on the two-class version of FMNIST. The plots (a)–(g) are the same as in figure 3.




Figure 13. Results for a two-layer ReLU net trained on the two-class version of CIFAR-10. The plots (a)–(g) are the same as in figure 3.

pronounced example of how training a model with full-batch gradient descent causes the top eigenvalue of the NTK to hover around the EOS (Cohen *et al* 2021, Agarwala *et al* 2022). We also observe, as before, that the weight norm θ_t^2 increases for $\eta\lambda_{\max}(H_{\alpha\beta,0}) \gtrsim 4$ and the increase becomes particularly pronounced when $\eta\lambda_{\max}(H_{\alpha\beta,0}) \gtrsim 6$. Finally, as with MNIST, we observe the activation map becomes increasingly sparse as we increase the learning rate. The analogous results for CIFAR-10 are shown in figure 13. The results are almost identical in form to what we observed for FMNIST. In particular, we observe that the trained model has a sparse activation map and that the top eigenvalue of the trained NTK hovers around the EOS.

D.3.3. Three layers. In this appendix we repeat the above experiments for ReLU nets with two hidden layers. The results for the three-layer ReLU MLP trained on




Figure 14. Results for the three-layer ReLU net trained on the two-class version of MNIST.

the two-class versions of MNIST, FMNIST, and CIFAR-10 are given in figures 14–16, respectively.

The results are largely the same as before: increasing the learning rate produces a sparser activation map and the top eigenvalue of the trained NTK hovers around the EOS. We also observe that, above a critical learning rate, the weight norm θ_t^2 can increase over the course of training. One difference, in comparison to the two-layer case, is that the models diverge for a smaller learning rate. We also observe that the activation map for the first ReLU layer is sparser than the activation map for the second layer when the model is trained on the two-class versions of FMNIST and CIFAR-10. The first and second layers are comparably sparse when the model is trained on MNIST.

Catapult dynamics and phase transitions in quadratic nets




Figure 15. Results for the three-layer ReLU net trained on the two-class version of FMNIST.




Figure 16. Results for the three-layer ReLU net trained on the two-class version of CIFAR-10.

References

- Agarwala A, Pedregosa F and Pennington J 2022 Second-order regression models exhibit progressive sharpening to the edge of stability (arXiv:2210.04860)
- Andriushchenko M, Varre A, Pillaud-Vivien L and Flammarion N 2022 Sgd with large step sizes learns sparse features (arXiv:2210.05337)
- Arnaboldi L, Krzakala F, Loureiro B and Stephan L 2023 Escaping mediocrity: how two-layer networks learn hard generalized linear models with sgd (arXiv:2305.18502)
- Arora S, Du S S, Hu W, Li Z, Salakhutdinov R and Wang R 2019 On exact computation with an infinitely wide neural net (arXiv:1904.11955)
- Bai Y, Krause B, Wang H, Xiong C and Socher R 2020 Taylorized training: towards better approximation of neural network training at finite width *CoRR* (arXiv:2002.04010)
- Candes E J, Li X and Soltanolkotabi M 2015 Phase retrieval via wirtinger flow: Theory and algorithms *IEEE Trans. Inf. Theory* **61** 1985–2007
- Cohen J M, Kaur S, Li Y, Kolter J Z and Talwalkar A 2021 Gradient descent on neural networks typically occurs at the edge of stability (arXiv:2103.00065)
- Davis D, Drusvyatskiy D and Paquette C 2020 The nonsmooth landscape of phase retrieval *IMA J. Numer. Anal.* **40** 2652–95
- Dinh L, Pascanu R, Bengio S and Bengio Y 2017 Sharp minima can generalize for deep nets (arXiv:1703.04933)
- Dyer E and Gur-Ari G 2019 Asymptotics of wide networks from Feynman diagrams *CoRR* (arXiv:1909.11304)

- Goodfellow I, Bengio Y, Courville A and Bengio Y 2016 *Deep Learning* vol 1 (MIT Press)
- Hoefler T, Alistarh D, Ben-Nun T, Dryden N and Peste A 2021 Sparsity in deep learning: pruning and growth for efficient inference and training in neural networks (arXiv:2102.00554)
- Huang J and Yau H 2019 Dynamics of deep neural networks and neural tangent hierarchy *CoRR* (arXiv:1909.08156)
- Huang W, Du W, Xu R Y D and Liu C 2020 Implicit bias of deep linear networks in the large learning rate phase *CoRR* (arXiv:2011.12547)
- Iyer G, Hanin B and Rolnick D 2022 Maximal initial learning rates in deep relu networks (arXiv:2212.07295)
- Jacot A, Gabriel F and Hongler C 2018 Neural tangent kernel: convergence and generalization in neural networks (arXiv:1806.07572)
- Kaur S, Cohen J and Lipton Z C 2022 On the maximum hessian eigenvalue and generalization (arXiv:2206.10654)
- Keskar N S, Mudigere D, Nocedal J, Smelyanskiy M and Tang P T P 2016 On large-batch training for deep learning: generalization gap and sharp minima *CoRR* (arXiv:1609.04836)
- Lee J, Bahri Y, Novak R, Schoenholz S S, Pennington J and Sohl-Dickstein J 2017 Deep neural networks as Gaussian processes (arXiv:1711.00165)
- Lee J, Xiao L, Schoenholz S, Bahri Y, Novak R, Sohl-Dickstein J and Pennington J 2019 Wide neural networks of any depth evolve as linear models under gradient descent *Advances in Neural Information Processing Systems* vol 32 pp 8572–83
- Lewkowycz A, Bahri Y, Dyer E, Sohl-Dickstein J and Gur-Ari G 2020 The large learning rate phase of deep learning: the catapult mechanism (arXiv:2003.02218)
- Li Y, Wei C and Ma T 2019 Towards explaining the regularization effect of initial large learning rate in training neural networks *CoRR* (arXiv:1907.04595)
- Li Z, You C, Bhojanapalli S, Li D, Rawat A S, Reddi S J, Ye K, Chern F, Yu F, Guo R and Kumar S 2022 Large models are parsimonious learners: activation sparsity in trained transformers (arXiv:2210.06313)
- Martin S, Bach F and Biroli G 2024 On the impact of overparameterization on the training of a shallow neural network in high dimensions *Int. Conf. on Artificial Intelligence and Statistics* (PMLR) pp 3655–63
- Mignacco F, Urbani P and Zdeborová L 2021 Stochasticity helps to navigate rough landscapes: comparing gradient-descent-based algorithms in the phase retrieval problem *Mach. Learn.: Sci. Technol.* **2** 035029
- Paszke A *et al* 2019 Pytorch: An imperative style, high-performance deep learning library *Advances in Neural Information Processing Systems* p 32
- Roberts D A, Yaida S and Hanin B 2022 *The Principles of Deep Learning Theory* vol 46 (Cambridge University Press)
- Sarao Mannelli S, Biroli G, Cammarota C, Krzakala F, Urbani P and Zdeborová L 2020a Complex dynamics in simple neural networks: Understanding gradient flow in phase retrieval *Advances in Neural Information Processing Systems* vol 33 pp 3265–74
- Sarao Mannelli S, Vanden-Eijnden E and Zdeborová L 2020b Optimization and generalization of shallow neural networks with quadratic activation functions *Advances in Neural Information Processing Systems* vol 33 pp 13445–55
- Sohl-Dickstein J, Novak R, Schoenholz S S and Lee J 2020 On the infinite width limit of neural networks with a standard parameterization (arXiv:2001.07301)
- Sun J, Qu Q and Wright J 2018 A geometric analysis of phase retrieval *Found. Comput. Math.* **18** 1131–98
- Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez A N, Kaiser L and Polosukhin I 2017 Attention is all you need *Proc. 31st Int. Conf. on Neural Information Processing Systems* pp 6000–10
- Wang Y, Chen M, Zhao T and Tao M 2021 Large learning rate tames homogeneity: convergence and balancing effect (arXiv:2110.03677)
- Yaida S 2019 Non-Gaussian processes and neural networks at finite widths (arXiv:1910.00019)
- Yaida S 2020 Non-Gaussian processes and neural networks at finite widths *Mathematical and Scientific Machine Learning* (PMLR) pp 165–92
- Zhu L, Liu C, Radhakrishnan A and Belkin M 2022 Quadratic models for understanding neural network dynamics (arXiv:2205.11787)