

514 **A Experimental details**

515 **Datasets:** We considered the MNIST [10], Fashion-MNIST [64], and CIFAR-10 [42] datasets. We
 516 standardized the images and used one-hot encoding for the labels.

517 **Models:** We considered fully connected networks (FCNs), Myrtle family CNNs [58] and ResNets
 518 (version 1) [28] trained using the JAX [7], and Flax libraries [29]. We use d and w to denote the
 519 depth and width of the network. Below, we provide additional details of the models and clarify what
 520 width corresponds to for CNNs and ResNets.

521 1. **FCNs:** We considered ReLU FCNs with constant width w in Neural Tangent Parameteriza-
 522 tion (NTP) / Standard Parameterization (SP), initialized at criticality [52]. The models do
 523 not include bias or normalization. The forward pass of the pre-activations from layer l to
 524 $l + 1$ is given by

$$h_i^{l+1} = \gamma^l \sum_j^w W_{ij}^l \phi(h_j^l), \quad (2)$$

525 where $\phi(\cdot)$ is the ReLU activation and γ^l is a constant. For NTP, $\gamma^l = 2/\sqrt{w}$ and the weights
 526 W^l are initialized using normal distribution, i.e., $W_{ij}^l \sim \mathcal{N}(0, 1)$. For SP, $\gamma^l = 1$ and the
 527 weights W^l are initialized as $W_{ij}^l \sim \mathcal{N}(0, 2/w)$. For the last layer, we have $\gamma^L = 1/\sqrt{w}$ for
 528 NTP and $W_{ij}^L \sim \mathcal{N}(0, 1/w)$ for SP.

529 We considered $d \in \{4, 8, 16\}$ and $w \in \{256, 512, 1024, 2048\}$. For $d/w \gtrsim 1/16$, the
 530 dynamics is noisier, and it becomes challenging to separate the underlying deterministic
 531 dynamics from random fluctuations (see Appendix E).

532 2. **CNNs:** We considered Myrtle family ReLU CNNs [58] without any bias or normalization
 533 in Standard Parameterization (SP), initialized using He initialization [28]. The above model
 534 uses a fixed number of channels in each layer, which we refer to as the width of the network.
 535 In this case, the forward pass equations for the pre-activations from layer l to layer $l + 1$ are
 536 given by

$$h_i^{l+1}(\alpha) = \sum_j^w \sum_{\beta \in \ker} W_{ij}^{l+1}(\beta) \phi(h_i^l(\alpha + \beta)), \quad (3)$$

537 where α, β label the spacial location. The weights are initialized as $W_{ij}^l(\beta) \sim \mathcal{N}(0, 2/k^2 w)$,
 538 where k is the filter size. We considered $d \in \{5, 7, 10\}$ and $w \in \{64, 128, 256\}$.

539 3. **ResNets:** We considered version 1 ResNet [28] implementations from `Flax examples`
 540 without Batch Norm or regularization. For ResNets, width corresponds to the number of
 541 channels in the first block. For example, the standard ResNet-18 has four blocks with widths
 542 $[w, 2w, 4w, 8w]$, with $w = 64$. We refer to w as the width or the widening factor. We
 543 considered ResNet-18 and ResNet-34 with $w \in \{32, 64, 128\}$.

544 All the models are trained with the average loss over the batch $\mathcal{D}_B = \{(x_\mu, y_\mu)\}_{\mu=1}^B$, i.e.,
 545 $L(x, y_{\mathcal{D}_B}) = 1/B \sum_{\mu=1}^B \ell(x_\mu, y_\mu)$, where $\ell(\cdot)$ is the loss function. This normalization, along with
 546 initialization, ensures that the loss is $\mathcal{O}(1)$ at initialization.

547 **Bias:** Throughout this work, we have primarily focused on models without any bias for simplicity. In
 548 Appendix K, we demonstrate that bias does not have an appreciable impact on the results.

549 **Batch size:** We use a batch size of 512 and scale the learning rate as $\eta = c/\lambda_0^H$ in all our experiments,
 550 unless specified. Appendix J shows results for a smaller batch size $B = 32$.

551 **Learning rate:** We scale the learning rate constant as $c = 2^x$, with $x \in \{-1.0, \dots, x_{max}\}$ in steps of
 552 0.1. Here, x_{max} is related to the maximum learning rate constant as $c_{max} = 2^{x_{max}}$.

553 **Sharpness measurement:** We measure sharpness using the power iteration method with 20 iterations.
 554 We found that 20 iterations suffice both for MSE and cross-entropy loss. For MSE loss, we use $m =$
 555 2048 randomly selected training examples for evaluating sharpness at each step. In comparison, we

556 found that cross-entropy requires a large number of training examples to obtain a good approximation
557 of sharpness. Given the computational constraints, we use 4096 training examples to approximate
558 sharpness for cross-entropy loss.

559 **Averages over initialization and SGD runs:** All the critical constants depend on both the random
560 initializations and the SGD runs. In our experiments, we found that the fluctuations from initial-
561 ization at large d/w outweigh the randomness coming from different SGD runs. Thus, we focus on
562 initialization averages in all our experiments.

563 A.1 Compute usage

564 We utilized different computational resources depending on the task complexity. For less demanding
565 tasks, we performed computation for a total of 2800 hours, utilizing a seventh of an NVIDIA A100
566 GPU. For more computationally intensive tasks, we utilized a full NVIDIA A100 GPU for a total
567 300 hours.

568 A.2 Reproducibility

569 We provide a notebook in the supplementary material that reproduces the main results of the paper.

570 A.3 Details of Figures in the main text:

571 **Figure 1:** A shallow CNN ($d = 5, w = 128$) in SP trained on the CIFAR-10 dataset with MSE
572 loss for 1000 epochs using SGD with learning rates $\eta = c/\lambda_0^H$ and batch size $B = 512$. We measure
573 sharpness at every step for the first epoch, every epoch between 10 and 100 epochs, and every 10
574 epochs beyond 100.

575 **Figure 2:** (top panel) A wide ($d = 5, w = 512$) and (bottom panel) a deep CNN ($d = 10, w = 128$)
576 in SP trained on the CIFAR-10 dataset with MSE loss for $t = 10$ steps using vanilla SGD with
577 learning rates $\eta = c/\lambda_0^H$ and batch size $B = 512$.

578 **Figure 3:** (a) FCNs in NTP with $d = 8$ and $w \in \{256, 512, 1024, 2048\}$ trained on the MNIST
579 dataset, (b) CNNs in SP with $d = 7$ and $w \in \{64, 128, 256, 512\}$ trained on the Fashion-MNIST
580 dataset, (c) ResNet in SP with $d = 18$ and $w \in \{32, 64, 128\}$ trained on the CIFAR-10 dataset
581 (without batch normalization). Each data point in the figure represents an average of ten distinct
582 initialization, and the solid lines represent a two-degree polynomial $y = a + bx + cx^2$ fitted to the
583 raw data points. Here, where $x = 1/w$, and y can take on one of three values: c_{loss}, c_{sharp} and
584 c_{max} . We show the error bars in Appendix C only as they may give an impression that the inequality
585 $c_{loss} \leq c_{sharp} \leq c_{max}$ can be violated.

586 **Figure 4:** (a) FCNs in NTP with $d \in \{4, 8, 16\}$ and $w \in \{256, 512, 1024, 2048\}$ trained on the
587 MNIST dataset, (b) CNNs in SP with $d \in \{5, 7, 10\}$ and $w \in \{64, 128, 256, 512\}$ trained on the
588 Fashion-MNIST dataset, (c) ResNet in SP with $d \in \{18, 34\}$ and $w \in \{32, 64, 128\}$ trained on the
589 CIFAR-10 dataset (without batch normalization).

590 **Figure 5:** Normalized sharpness measured at $c\tau = 200$ against the learning rate constant for 7-layer
591 CNNs in SP trained on the CIFAR-10 dataset, with $w \in \{128, 256, 512\}$. Each data point is an
592 average over five random initialization. Smoothing details are provided in Appendix I.2

593 **Figure 6:** The phase diagram of the uv model trained with MSE loss using gradient descent with (a)
594 the top eigenvalue of Hessian λ_t^H , (b) the trace of Hessian $\text{tr}(H_t)$ and (c) the square of the Frobenius
595 norm $\text{tr}(H_t^T H_t)$ used as a measure of sharpness. In (a), the learning rate is scaled as $\eta = c/\lambda_0^H$,
596 while in (b) and (c), the learning rate is scaled as $\eta = k/\text{tr}(H_0)$. The vertical dashed line shows $c = 2$
597 ($k = 2$) for reference. Each data point is an average over 500 random initializations.

598 **Figure 7:** Training trajectories of the uv model with (a, b) large ($w = 512$) and (c, d) small ($w = 2$)
599 width, trained for $t = 10$ training steps on a single example $(x, y) = (1, 0)$ with MSE loss using
600 vanilla gradient descent with learning rates (a, c) $c = 0.5$ and (b, d) $c = 2.50$.




Figure 8: Training trajectories of the uv model with (a) large ($w = 512$) and (b) a small ($w = 2$) widths trained for $t = 10$ training steps on a single example $(x, y) = (1, 0)$. For the wide network, $c_{loss} = 2.1$, $c_{sharp} = 2.6$, $c_{max} = 4.0$, and for the narrow network, $c_{loss} = 3.74$, $c_{sharp} = 4.63$, $c_{max} = 4.93$.

601 B Additional results for the uv model

602 B.1 Details of the model

603 Consider a two-layer linear network in (NTP) with unit input-output dimensions

$$f(x) = \frac{1}{\sqrt{w}} v^T u x, \quad (4)$$

604 where $x, f(x) \in \mathbb{R}$. Here, $u, v \in \mathbb{R}^w$ are trainable parameters, with each element initialized using
 605 the normal distribution, $u_i, v_i \sim \mathcal{N}(0, 1)$ for $i \in \{1, \dots, w\}$. The model is trained using MSE loss
 606 on a single training example $(x, y) = (1, 0)$, which simplifies the loss to

$$\mathcal{L}(u, v) = \frac{1}{2} f^2. \quad (5)$$

607 The trace of the Hessian $\text{tr}(H)$ has a simple expression in terms of the norms of the weight vectors

$$\text{tr}(H) = \frac{x^2}{w} (\|u\|^2 + \|v\|^2), \quad (6)$$

608 which is equivalent to the NTK for this model. The Frobenius norm of the Hessian $\|H\|_F$ can be
 609 written in terms of the loss \mathcal{L} and $\text{tr}(H)$

$$\|H\|_F^2 = \text{tr}(H)^2 + 2f^2 \left(1 + \frac{2}{w}\right) = \text{tr}(H)^2 + 4\mathcal{L} \left(1 + \frac{2}{w}\right) \quad (7)$$

610 The gradient descent updates of the model trained using MSE loss on a single training example
 611 $(x, y) = (1, 0)$ are given by




Figure 9: (a) Normalized sharpness measured at $\tau = 100$ steps against the learning rate constant for the uv model trained on $(x, y) = (1, 0)$, with varying widths. Each data point is an average of over 500 initializations, where the shaded region depicts the standard deviation around the mean trend. (b, c) Smooth estimations of the first two derivatives, χ_τ and χ'_τ , of the, averaged normalized sharpness wrt the learning rate constant. The vertical dashed lines denote c_{crit} estimated for each width, using the maximum of χ'_τ . Here, we have removed the points beyond $c = 3.5$ for the calculation of derivatives to avoid large fluctuations near the divergent phase. Smoothing details are described in Appendix I.2

$$v_{t+1} = v_t - \eta f_t \frac{1}{\sqrt{w}} u_t x \quad (8)$$

$$u_{t+1} = u_t - \eta f_t \frac{1}{\sqrt{w}} v_t x \quad (9)$$

612 The update equations in function space can be written in terms of the trace of the Hessian $\text{tr}(H)$.

$$f_{t+1} = f_t \left(1 - \eta \text{tr}(H_t) + \frac{\eta^2 f_t^2}{w} \right) \quad (10)$$

$$\text{tr}(H_{t+1}) = \text{tr}(H_t) + \frac{\eta f_t^2}{w} (\eta \text{tr}(H_t) - 4).$$

613 Figure 8 shows the training trajectories of the uv model trained on $(x, y) = (1, 0)$ using MSE loss
614 for 10 training steps. The model shows similar dynamics to those presented in Section 2. It is worth
615 mentioning that the above equations have been analyzed in [47] at large width. In the following
616 subsections, we extend their analysis by incorporating the higher-order terms to analyze the effect of
617 finite width.

618 B.2 The intermediate saturation regime

619 The uv model trained on $(x, y) = (1, 0)$ does not show the progressive sharpening and late-time
620 regimes (iii) and (iv) described in Section 1. Hence, we can measure sharpness at the end of training
621 to analyze how it is reduced upon increasing the learning rate and to compare it with the intermediate
622 saturation regime results in Section 3.

623 Figure 9(a) shows the normalized sharpness measured at $\tau = 100$ steps for various widths. This
624 behavior reproduces the results observed in the intermediate saturation regime in Section 3. In
625 particular, we can see stages (1) and (2), where $\lambda_\tau^H / \lambda_0^H$ starts off fairly independent of learning rate
626 constant c , and then dramatically reduces when $c > 2$; stage (3), where $\lambda_\tau^H / \lambda_0^H$ plateaus at a small
627 value as a function of c is too close to the divergent phase in this model to be clearly observed. The
628 corresponding derivatives of the averaged normalized sharpness, χ_τ , and χ'_τ , are shown in Figure
629 9(b, c). The vertical dashed lines denote c_{crit} estimated for each width, using the maximum of χ'_τ .
630 We observe that $c_{crit} = 2$ for all widths.




Figure 10: (a, b) The averaged loss at the first step $\langle \mathcal{L}_1 / \mathcal{L}_0 \rangle$ against the learning rate constant k for varying widths obtained from (a) inequality [16](#) and (b) numerical experiments. The intersection of $\langle \mathcal{L}_1 / \mathcal{L}_0 \rangle$ with the horizontal line $y = 1$ depicts k_{loss} . The two vertical lines $k = 2$ and $k = 4$ mark the endpoints of k_{loss} at small and large widths. The shaded region in (b) shows the standard deviation around the mean trend. (c) The scaling of λ_0^H and $\text{tr}(H_0)$ with width.

631 B.3 Opening of the sharpness reduction phase in the wv model

632 This section shows that $\mathcal{O}(1/w)$ terms in Equation [10](#) effectively lead to the opening of the sharpness
 633 reduction phase with $1/w$ in the wv model. In Appendix [B.2](#) we demonstrated that for the wv model,
 634 $c_{crit} = 2$ for all values of widths. Hence, it suffices to show that c_{loss} increases from the value 2
 635 as $1/w$ increases. We do so by finding the smallest k such that the averaged loss over initializations
 636 increases during early training.

637 It follows from Equation [10](#) that the averaged loss increases in the first training step if the following
 638 holds

$$\left\langle \frac{\mathcal{L}_1}{\mathcal{L}_0} \right\rangle = \left\langle \left(1 - \eta \text{tr}(H_0) + \frac{\eta^2 f_0^2}{w} \right)^2 \right\rangle > 1, \quad (11)$$

639 where $\langle \cdot \rangle$ denotes the average over initializations. On scaling the learning rate with trace as $\eta =$
 640 $k / \text{tr}(H_0)$, we have

$$\left\langle \frac{\mathcal{L}_1}{\mathcal{L}_0} \right\rangle = \left\langle \left(1 - k + \frac{k^2}{w} \frac{f_0^2}{\text{tr}(H_0)^2} \right) \right\rangle > 1 \quad (12)$$

$$\left\langle \frac{\mathcal{L}_1}{\mathcal{L}_0} \right\rangle = \left((1 - k)^2 + 2(1 - k) \frac{k^2}{w} \left\langle \frac{f_0^2}{\text{tr}(H_0)^2} \right\rangle + \frac{k^4}{w^2} \left\langle \frac{f_0^4}{\text{tr}(H_0)^4} \right\rangle \right) > 1. \quad (13)$$

641 The required two averages have the following expressions as shown in Appendix [B.8](#):

$$\left\langle \frac{f_0^2}{\text{Tr}(H_0)^2} \right\rangle = \frac{w}{4(w + 1)} \quad (14)$$

$$\left\langle \frac{f_0^4}{\text{Tr}(H_0)^4} \right\rangle = \frac{3(w + 2)w^3}{16} \frac{\Gamma(w)}{\Gamma(w + 4)}. \quad (15)$$

642 Inserting the above expressions in Equation [13](#), on average the loss increases in the very first step if
 643 the following inequality holds

$$\left\langle \frac{\mathcal{L}_1}{\mathcal{L}_0} \right\rangle = \left((1 - k)^2 + \frac{k^2(1 - k)}{2(w + 1)} + \frac{3k^4}{16(w + 3)(w + 1)} \right) > 1 \quad (16)$$

644 The graphical representation of the above inequality shown in Figure 10(a) is in excellent agreement
645 with the experimental results presented in Figure 10(b).

646 Let us denote k'_{loss} as the minimum learning rate constant such that the average loss increases in the
647 first step. Similarly, let k_{loss} denote the learning rate constant if the loss increases in the first 10 steps.
648 Then, k'_{loss} increases from the value 2 as $1/w$ increases as shown in Figure 10(a). By comparison, the
649 trace reduces at any step if $\eta \text{tr}(H_t) < 4$. At initialization, this condition becomes $k < 4$. Hence, for
650 $k < k'_{loss}$, both the loss and trace monotonically decrease in the first training step. These arguments
651 can be extended to later training steps, revealing that the loss and trace will continue to decrease for
652 $k < k'_{loss}$.

653 Next, let η_{loss} denote the learning rate corresponding to c_{loss} . Then, we have $\eta_{loss} = \frac{c_{loss}}{\lambda_0^H} = \frac{k_{loss}}{\text{tr}(H_0)}$,
654 implying

$$c_{loss} = k_{loss} \frac{\lambda_0^H}{\text{tr}(H_0)}. \quad (17)$$

655 Figure 10(c) shows that $\lambda_0^H \geq \text{tr}(H_0)$ for all widths, implying $c_{loss} \geq k_{loss}$. Hence, c_{loss} increases
656 with $1/w$ as observed in Figure 6(a). In Appendix B.2 we demonstrated that for the uv model,
657 $c_{crit} = 2$ for all values of widths. Incorporating this with c_{loss} increases with $1/w$, we have sharpness
658 reduction phase opening up as $1/w$ increases.

659 B.4 Opening of the loss catapult phase at finite width

660 In this section, we use the Frobenius norm of the Hessian $\|H\|_F$ as a proxy for the sharpness and
661 demonstrate the emergence of the loss-sharpness catapult phase at finite width. In particular, We
662 analyze the expectation value $\langle \text{tr}(H^T H) \rangle$ after the first training step near $k = k_{loss}$ and show that
663 $k_{loss} \leq k_{frob}$, with the difference increasing with $1/w$. First, we write $\text{tr}(H_t^T H_t)$ in terms of \mathcal{L}_t and
664 $\text{tr}(H_t)$

$$\text{tr}(H_t^T H_t) = \text{tr}(H_t)^2 + 4 \left(1 + \frac{2}{w}\right) \mathcal{L}_t. \quad (18)$$

665 Next, using Equations 1 we write down the change in $\text{tr}(H_t^T H_t)$ after the first training step in terms
666 of $\text{tr}(H_0)$ and \mathcal{L}_0

$$\begin{aligned} \Delta \text{tr}(H_1^T H_1) &= \text{tr}(H_1^T H_1) - \text{tr}(H_0^T H_0) = \text{tr}(H_1)^2 - \text{tr}(H_0)^2 + 4 \left(1 + \frac{2}{w}\right) (\mathcal{L}_1 - \mathcal{L}_0) \\ \Delta \text{tr}(H_1^T H_1) &= \frac{\eta f_0^2}{w} (\eta \text{tr}(H_0) - 4) \left[\frac{\eta f_0^2}{w} (\eta \text{tr}(H_0) - 4) + 2 \text{tr}(H_0) \right] + 4 \left(1 + \frac{2}{w}\right) (\mathcal{L}_1 - \mathcal{L}_0) \end{aligned} \quad (19)$$

667 Next, we substitute $\eta = k / \text{tr}(H_0)$ to obtain the above equation as a function of k

$$\Delta \text{tr}(H_1^T H_1) = \frac{k(k-4)}{w} \left[\frac{k(k-4)}{w} \frac{f_0^4}{\text{tr}(H_0)^2} + 2f_0^2 \right] + 4 \left(1 + \frac{2}{w}\right) (\mathcal{L}_1 - \mathcal{L}_0) \quad (20)$$

668 Finally, we calculate the expectation value of $\langle \Delta \text{tr}(H_1^T H_1) \rangle$

$$\langle \Delta \text{tr}(H_1^T H_1) \rangle = \frac{k(k-4)}{w} \left[\frac{k(k-4)}{w} \left\langle \frac{f_0^4}{\text{tr}(H_0)^2} \right\rangle + 2 \langle f_0^2 \rangle \right] + 4 \left(1 + \frac{2}{w}\right) \langle \mathcal{L}_1 - \mathcal{L}_0 \rangle, \quad (21)$$




Figure 11: The relationship between the critical constants for the uv model trained on a single training examples $(x, y) = (1, 0)$ with MSE loss using gradient descent. Each data point corresponds to a random initialization

669 by estimating $\left\langle \frac{f_0^4}{\text{tr}(H_0)^2} \right\rangle$ using the approach demonstrated in the previous section

$$\left\langle \frac{f_0^4}{\text{tr}(H_0)^2} \right\rangle = \frac{3w}{4(w+3)}. \quad (22)$$

670 Inserting $\left\langle \frac{f_0^4}{\text{tr}(H_0)^2} \right\rangle$ in Equation 21 along with $\langle f_0^2 \rangle = 1$, we have

$$\langle \Delta \text{tr}(H_1^T H_1) \rangle = \underbrace{\frac{k(k-4)}{w} \left[\frac{3k(k-4)}{4(w+3)} + 2 \right]}_{I(k,w)} + 4 \left(1 + \frac{2}{w} \right) \langle \mathcal{L}_1 - \mathcal{L}_0 \rangle \quad (23)$$

671 At infinite width, the above equation reduces to $\langle \Delta \text{tr}(H_1^T H_1) \rangle = 4 \langle \mathcal{L}_1 - \mathcal{L}_0 \rangle$, and hence, $k_{frob} =$
672 k_{loss} . For any finite width, $I(k, w) < 0$ for $0 < k < 4$. At $k \leq k_{loss}$, $\mathcal{L}_1 - \mathcal{L}_0 \leq 0$, and therefore
673 $\langle \Delta \text{tr}(H_1^T H_1) \rangle < 0$. In order for the sharpness to catapult, we require $\langle \Delta \text{tr}(H_1^T H_1) \rangle > 0$ and
674 therefore $k_{frob} > k_{loss}$. As $1/w$ increases $|I(k, w)|$ also increases, which means a higher value of
675 $\mathcal{L}_1 - \mathcal{L}_0$ is required to reach a point where $\langle \Delta \text{tr}(H_1^T H_1) \rangle \geq 0$. Thus $k_{frob} - k_{loss}$ increases with
676 $1/w$.

677 B.5 The early training trajectories

678 Figure 8 shows the early training trajectories of the uv model with large ($w = 512$) and small ($w = 2$)
679 widths. The dynamics depicted show several similarities with early training dynamics of real-world
680 models shown in Figure 2. At small widths, the loss catapults at relatively higher learning rates
681 (specifically, at $c_{loss} = 3.74$, which is significantly higher than the critical value of $c_{crit} = 2$).

682 B.6 Relationship between critical constants

683 Figure 11 shows the relationship between various critical constants for the uv model. The data show
684 that the inequality $c_{loss} \leq c_{sharp} \leq c_{max}$ holds for every random initialization of the uv model.

685 B.7 Phase diagrams with error bars

686 This section shows the variation in the phase diagram boundaries of the uv model shown in Figure 6(a,
687 b). Figure 12 shows these phase diagrams. Each data point is an average of over 500 initializations.
688 The horizontal bars around each data point indicate the region between 25% and 75% quantile.

689 B.8 Derivation of the expectation values

690 Here, we provide the detailed derivation of the averages $\left\langle \frac{f_0^2}{\text{Tr}(H_0)^2} \right\rangle$ and $\left\langle \frac{f_0^4}{\text{Tr}(H_0)^4} \right\rangle$. We begin by
691 finding the average $\left\langle \frac{f_0^2}{\text{Tr}(H_0)^2} \right\rangle$




Figure 12: The phase diagram of the uv model trained with MSE loss using gradient descent with (a) sharpness λ_t^H (b) trace of Hessian tr_0^H and (c) the square of the Frobenius norm $\text{tr}(H_t^T H_t)$ used as a measure of sharpness. In (a), the learning rate is scaled as $\eta = c/\lambda_0^H$, while in (b) and (c), the learning rate is scaled as $\eta = k/\text{tr}(H_0)$. Each data point denotes an average of over 500 initialization, and the smooth curve represents a 2-degree polynomial fitted to the raw data. The horizontal bars around the average data point indicate the region between 25% and 75% quantile.

$$\left\langle \frac{f_0^2}{\text{Tr}(H_0)^2} \right\rangle = w \int_{-\infty}^{\infty} \prod_{i=1}^w \left(\frac{dv_i du_i}{2\pi} \right) \exp\left(-\frac{\|u\|^2 + \|v\|^2}{2}\right) \frac{\sum_{j,k=1}^w u_j v_j u_k v_k}{(\|u\|^2 + \|v\|^2)^2}, \quad (24)$$

692 where $\|\cdot\|$ denotes the norm of the vectors.

693 The above integral is non-zero only if $j = k$. Hence, it is a sum of w identical integrals. Without any
694 loss of generality, we solve this integral for $j = 1$ and multiply by w to obtain the final result, i.e.,

$$\left\langle \frac{f_0^2}{\text{Tr}(H_0)^2} \right\rangle = w^2 \int_{-\infty}^{\infty} \prod_{i=1}^w \left(\frac{dv_i du_i}{2\pi} \right) \exp\left(-\frac{\|u\|^2 + \|v\|^2}{2}\right) \frac{u_1^2 v_1^2}{(\|u\|^2 + \|v\|^2)^2} \quad (25)$$

695 Consider a transformation of $u, v \in \mathbb{R}^w$ into w dimensional spherical coordinates such that

$$u_1 = r_u \cos \varphi_{u_1}, \quad v_1 = r_v \cos \varphi_{v_1}, \quad (26)$$

696 which yields,

$$\left\langle \frac{f_0^2}{\text{Tr}(H_0)^2} \right\rangle = \frac{w^2}{(2\pi)^w} \int dr_u dr_v d\Omega_{u,w} d\Omega_{v,w} r_u^{w-1} r_v^{w-1} \exp\left(-\frac{r_u^2 + r_v^2}{2}\right) \frac{r_u^2 \cos^2 \varphi_{u_1} r_v^2 \cos^2 \varphi_{v_1}}{(r_u^2 + r_v^2)^2} \quad (27)$$

$$\left\langle \frac{f_0^2}{\text{Tr}(H_0)^2} \right\rangle = \frac{w^2}{(2\pi)^w} \int dr_u dr_v \exp\left(-\frac{r_u^2 + r_v^2}{2}\right) \frac{r_u^2 r_v^2}{(r_u^{w+1} + r_v^{w+1})^2} \int d\Omega_{u,w} d\Omega_{v,w} \cos^2 \varphi_{u_1} \cos^2 \varphi_{v_1} \quad (28)$$

$$\left\langle \frac{f_0^2}{\text{Tr}(H_0)^2} \right\rangle = \frac{w^2}{(2\pi)^w} \underbrace{\int dr_u dr_v \exp\left(-\frac{r_u^2 + r_v^2}{2}\right) \frac{r_u^2 r_v^2}{(r_u^{w+1} + r_v^{w+1})^2}}_{I_r} \left(\underbrace{\int d\Omega_w \cos^2 \varphi_1}_{I_\varphi} \right)^2, \quad (29)$$

697 where $d\Omega$ denotes the w dimensional solid angle element. Here, we denote the radial and angular
698 integrals by I_r and I_φ respectively. The radial integral I_r is

$$I_r = \int_0^\infty dr_u dr_v \frac{r_u^{w+1} r_v^{w+1}}{(r_u^2 + r_v^2)^2} \exp\left(-\frac{r_u^2 + r_v^2}{2}\right). \quad (30)$$

699 Let $r_u = R \cos \theta$ and $r_v = R \sin \theta$ with $R \in [0, \infty)$ and $\theta \in [-\frac{\pi}{2}, \frac{\pi}{2}]$, then we have

$$I_r = \int_0^\infty dR R^{2w-1} e^{-R^2/2} \int_0^{\pi/2} d\theta \cos^{w+1} \theta \sin^{w+1} \theta \quad (31)$$

$$I_r = \frac{\sqrt{\pi} \Gamma(w) \Gamma(\frac{w+2}{2})}{2^3 \Gamma(\frac{w+3}{2})}, \quad (32)$$

700 where $\Gamma(\cdot)$ denotes the Gamma function. The angular integral I_φ is

$$I_\varphi = \int d\varphi_1 d\varphi_2 \dots d\varphi_{w-1} \sin^{w-2} \varphi_1 \cos^2 \varphi_1 \sin^{w-3} \varphi_2 \dots \sin \varphi_{w-2} \quad (33)$$

$$I_\varphi = \int d\varphi_1 d\varphi_2 \dots d\varphi_{w-1} \sin^{w-2} \varphi_1 \sin^{w-3} \varphi_2 \dots \sin \varphi_{w-2} \frac{\int_0^\pi d\varphi_1 \sin^{w-2} \varphi_1 \cos^2 \varphi_1}{\int_0^\pi d\varphi_1 \sin^{w-2} \varphi_1} \quad (34)$$

$$I_\varphi = \frac{\pi^{w/2}}{\Gamma(\frac{w+2}{2})}. \quad (35)$$

701 Plugging in Equations [32](#) and [35](#) into Equation [29](#), we obtain a very simple expression

$$\left\langle \frac{f_0^2}{Tr(H_0)^2} \right\rangle = \frac{w^2}{2^{w+3}} \frac{\sqrt{\pi} \Gamma(w)}{\Gamma(\frac{w+2}{2}) \Gamma(\frac{w+3}{2})} = \frac{w}{4(w+1)}. \quad (36)$$

702 The other integral $\left\langle \frac{f_0^4}{Tr(H_0)^4} \right\rangle$ can be obtained by generalizing the above approach as described below

$$\left\langle \frac{f_0^4}{Tr(H_0)^4} \right\rangle = w^2 \int_{-\infty}^\infty \prod_{i=1}^w \left(\frac{dv_i du_i}{2\pi} \right) \exp\left(-\frac{\|u\|^2 + \|v\|^2}{2}\right) \frac{\sum_{j,k,l,m=1}^w u_j v_j u_k v_k u_l v_l u_m v_m}{(\|u\|^2 + \|v\|^2)^4}. \quad (37)$$

703 The integral is zero if either $j = k$ and $l = m$ or $j = k = l = m$, which we consider separately.

704 Without loss of generality, we find the following integrals

$$\left\langle \frac{f_0^4}{Tr(H_0)^4} \right\rangle_{22} = w^2 \int_{-\infty}^\infty \prod_{i=1}^w \left(\frac{dv_i du_i}{2\pi} \right) \exp\left(-\frac{\|u\|^2 + \|v\|^2}{2}\right) \frac{u_1^2 u_2^2 v_1^2 v_2^2}{(\|u\|^2 + \|v\|^2)^4} \quad (38)$$

$$\left\langle \frac{f_0^4}{Tr(H_0)^4} \right\rangle_4 = w^2 \int_{-\infty}^\infty \prod_{i=1}^w \left(\frac{dv_i du_i}{2\pi} \right) \exp\left(-\frac{\|u\|^2 + \|v\|^2}{2}\right) \frac{u_1^4 v_1^4}{(\|u\|^2 + \|v\|^2)^4}, \quad (39)$$

705 which have the following expressions

$$\left\langle \frac{f_0^4}{Tr(H_0)^4} \right\rangle_{22} = \frac{w^2}{16} \frac{\Gamma(w)}{\Gamma(w+4)} \quad (40)$$

$$\left\langle \frac{f_0^4}{Tr(H_0)^4} \right\rangle_4 = \frac{9w^2}{16} \frac{\Gamma(w)}{\Gamma(w+4)}, \quad (41)$$




Figure 13: Phase diagrams of FCNs in NTP with varying depths trained on the MNIST dataset using MSE loss.




Figure 14: Phase diagrams of FCNs in NTP with varying depths trained on the Fashion-MNIST dataset.




Figure 15: Phase diagrams of FCNs in NTP with varying depths trained on the CIFAR-10 dataset.

706 where $\Gamma(\cdot)$ denotes the gamma function. On combining the expressions with their multiplicities, we
 707 obtain the final result

$$\left\langle \frac{f_0^4}{\text{Tr}(H_0)^4} \right\rangle = 3w(w-1) \left\langle \frac{f_0^4}{\text{Tr}(H_0)^4} \right\rangle_{22} + w \left\langle \frac{f_0^4}{\text{Tr}(H_0)^4} \right\rangle_4 \quad (42)$$

$$\left\langle \frac{f_0^4}{\text{Tr}(H_0)^4} \right\rangle = \frac{3(w+2)w^3}{16} \frac{\Gamma(w)}{\Gamma(w+4)} \quad (43)$$

708 C Phase diagrams of early training

709 This section describes experimental details and shows additional phase diagrams of early training.
 710 The results include (1) FCNs in NTP trained on MNIST, Fashion-MNIST, and CIFAR-10 datasets,
 711 (2) CNNs in SP trained on Fashion-MNIST and CIFAR-10, and (3) ResNets in SP trained on CIFAR-




Figure 16: Phase diagrams of Convolutional Neural Networks (CNNs) in SP with varying depths trained on the Fashion-MNIST dataset.




Figure 17: Phase diagrams of Convolutional Neural Networks (CNNs) in SP with varying depths trained on the CIFAR-10 dataset.




Figure 18: Phase diagrams of Resnets in SP with different depths trained on the CIFAR-10 dataset.

712 10 datasets using MSE loss. Figures 13 to 18 show these results. The depths and widths are the same
 713 as specified in Appendix A. Each data point is an average over 10 initializations. The horizontal bars
 714 around the average data point indicate the region between 25% and 75% quantile. Phase diagrams
 715 for cross-entropy results are shown in Appendix F.

716 **Additional experimental details** : We train each model for $t = 10$ steps using SGD with learning
 717 rates $\eta = c/\lambda_0^H$ and batch size of 512, where $c = 2^x$ with $x \in \{0.0, \dots, x_{max}\}$ in steps of 0.1.
 718 Here, x_{max} is related to the maximum trainable learning rate constant as $c_{max} = 2^{x_{max}}$. We have
 719 considered 10 random initializations for each model. As mentioned in Appendix A, we do not
 720 consider averages over SGD runs as the randomness from initialization outweighs it. Hence, we
 721 obtain 10 values for each of the critical values in the following results. For each initialization, we
 722 compute the critical constants using Definitions 1 and 3. To avoid a random increase in loss and
 723 sharpness due to fluctuations, we round off the values of λ_t^H/λ_0^H and $\mathcal{L}_t/\mathcal{L}_0$ to their second decimal
 724 places before comparing with 1. We denote the average values using data points and variation using




Figure 19: Phase diagrams of FCNs in NTP with varying widths trained on the CIFAR-10 dataset.




Figure 20: The relationship between various critical constants for various models and datasets. Each data point corresponds to a model with random initialization. The dashed line denotes the values where $x = y$.

725 horizontal bars around the average data points, which indicate the region between 25% and 75%
 726 quantile. The smooth curves are obtained by fitting a two-degree polynomial $y = a + bx + cx^2$ with
 727 $x = 1/w$ and y can take on one of three values: c_{loss} , c_{sharp} and c_{max} .

728 **Phase diagrams with depth** Figure 19: shows the phase diagrams with depth for FCNs in NTP
 729 trained on the CIFAR-10 dataset. The phase diagrams look qualitatively similar compared to the $1/w$
 730 phase diagrams.

731 D Relationship between various critical constants

732 Figure 20 illustrates the relationship between the early training critical constants for models and
 733 datasets. The experimental setup is the same as in Appendix C. Typically, we find that $c_{loss} \leq$
 734 $c_{sharp} \leq c_{max}$ holds true. However, there are some exceptions, which are observed at high values
 735 of d/w (see 20 (d, e)), where the trends of the critical constants converge, and large fluctuations can
 736 cause deviations from the inequality.




Figure 21: Training loss trajectories of ReLU FCNs with $d = 16$ trained on the CIFAR-10 dataset with MSE loss using SGD with learning rate $\eta = c/\lambda_0^H$ and batch size $B = 512$.




Figure 22: Sharpness trajectories of ReLU FCNs with $d = 16$ trained on the CIFAR-10 dataset with MSE loss using SGD with learning rate $\eta = c/\lambda_0^H$ and batch size $B = 512$.

737 E The effect of d/w on the noise in dynamics

738 In this section, we demonstrate that for FCNs with $d/w \gtrsim 1/16$, the dynamics becomes noise-
 739 dominated. This aspect makes it challenging to distinguish the underlying deterministic dynamics
 740 from random fluctuations. To demonstrate this, we consider FCNs trained on CIFAR-10 using MSE
 741 and cross-entropy loss and use 4096 training examples for estimating sharpness.

742 Figures 21 and 22 show the training loss and sharpness of FCNs with $d = 16$ and varying widths,
 743 trained on CIFAR-10 using MSE loss. We observe that the sharpness dynamics becomes noisier for
 744 $w \lesssim 64$.




Figure 23: Training loss trajectories of ReLU FCNs with $d = 16$ trained on the CIFAR-10 dataset with cross-entropy loss using SGD with learning rate $\eta = c/\lambda_0^H$ and batch size $B = 512$.




Figure 24: Sharpness trajectories of ReLU FCNs with $d = 16$ trained on the CIFAR-10 dataset with cross-entropy loss using SGD with learning rate $\eta = c/\lambda_0^H$ and batch size $B = 512$.

745 Figures 23 and 24 shows the training dynamics with loss switched to cross-entropy, while keeping
 746 the initialization and SGD batch sequence the same as in the MSE loss case. In comparison to MSE
 747 loss, the training loss and sharpness dynamics show a higher level of noise, especially for $w \lesssim 256$.
 748 As a result, it becomes difficult to characterize the training dynamics for $d/w \gtrsim 1/16$.

749 F Crossentropy

750 In this section, we provide additional results for models trained with cross-entropy (xent) loss and
 751 compare them with MSE results. Broadly speaking, models trained with cross-entropy loss show
 752 similar characteristics to those trained with MSE loss, such as, (i) sharpness reduction during early
 753 training, (ii) an increase in critical constants c_{loss} , c_{sharp} with d and $1/w$, (iii) $c_{loss} \leq c_{sharp} \leq c_{max}$.




Figure 25: The phase diagrams of early training of FCNs trained on the CIFAR-10 dataset using (a, b, c) MSE and (d, e, f) cross-entropy loss. Each data point is an average over 10 initializations, and solid lines represent a smooth curve fitted to raw data points. The horizontal bars around the averaged data point indicates the region between 25% and 75% quantile. For cross-entropy phase diagrams, the $c = 2$ line is shown for reference only and does not relate to c_{crit} .




Figure 26: Comparison of the relationship between critical constants for FCNs in SP trained on CIFAR-10 using MSE and cross-entropy loss. Each data point corresponds to a randomly initialized model with depths and widths mentioned in Appendix A.

754 However, the dynamics of models trained with cross-entropy loss is noisier compared to MSE as
 755 shown in the previous section, and characterizing these dynamics can be more complex. In the
 756 following experiments, we consider models trained on the CIFAR-10 dataset and used 4096 training
 757 examples to estimate sharpness.

758 F.1 Phase diagrams

759 Figure 25 compares the phase diagrams of FCNs in SP trained on the CIFAR-10 dataset, using both
 760 MSE and cross-entropy loss. The estimated critical constants for cross-entropy loss are generally
 761 more noisy, as quantified by the confidence intervals. In comparison to phase diagrams of models
 762 trained with MSE loss, we observe a few notable differences. First, the loss starts to catapult at a
 763 value appreciably larger than $c = 2$ at large widths. Primarily, $4 \lesssim c_{loss} \lesssim 8$. Additionally, c_{max}
 764 generally decreases with $1/w$. This decreasing trend becomes less sharp at large depths.

765 Despite these differences, the phase diagrams for both loss functions share various similarities. First,
 766 we observe sharpness reduces during early training for $c < c_{sharp}$ (see the first row of Figure 24).
 767 Next, we observe that the inequality $c_{loss} \leq c_{sharp} \leq c_{max}$ generally holds for both loss functions
 768 as demonstrated in Figure 26, barring some exceptions.

769 Figure 27 shows the phase diagrams for CNNs and ResNets trained on the CIFAR-10 dataset using
 770 cross-entropy loss. The observed critical constants are much noisier as quantified by the confidence
 771 intervals. Nevertheless, the phase diagram shows similar trends as mentioned above. For large $1/w$
 772 models, we found that progressive sharpening begins after 5 – 10 training steps. For these cases,
 773 we only use the first 5 steps to measure sharpness to avoid progressive sharpening. For CNNs, we
 774 observed that the dynamics becomes difficult to characterize for $w \lesssim 32$ and $d \gtrsim 10$, due to large
 775 fluctuations. Consequently, we’ve opted not to include these particular results.




Figure 27: Phase diagrams of (a, b) CNNs and (c) ResNets trained on the CIFAR-10 dataset with cross-entropy loss using SGD with $\eta = c/\lambda_0^H$ and $B = 512$.

776 **F.2 Intermediate saturation regime**




Figure 28: Sharpness measured at $c\tau = 100$ against the learning rate constant for FCNs trained on the CIFAR-10 dataset using cross-entropy loss, with varying depths and widths. Each curve is an average over ten initializations, where the shaded region depicts the standard deviation around the mean trend. The vertical dashed line shows $c = 2$ for reference.

777 Figure 28 shows the normalized sharpness measured at $c\tau = 100$ for FCNs trained on CIFAR-10
 778 using cross-entropy loss.³ Similar to MSE loss, we observe an abrupt drop in sharpness at large
 779 learning rates. However, this abrupt drop occurs at $2 \lesssim c_{crit} \lesssim 4$. The estimated sharpness is noisier
 780 (compare with Figure 37), which hinders a reliable estimation of c_{crit} . We speculate that we require a
 781 large number of averages for a reliable estimation of c_{crit} for cross-entropy loss. We leave the precise
 782 characterization of c_{crit} for cross-entropy loss for future work.

³The time step $\tau = 100/c$ is in the middle of the intermediate saturation regime for most of the models. For further details on estimating sharpness, see Appendix L.I




Figure 29: Comparison of the early training dynamics of (a, b) vanilla, (c, d) centered, and (e, f) zero-initialized FCNs (with depth = 8 and width = 512), trained on the CIFAR-10 dataset with MSE loss using gradient descent for 20 steps.

783 G The effect of setting model output to zero at initialization

784 In this section, we demonstrate the effect of network output $f(x; \theta_t)$ at initialization on the early
 785 training dynamics. In particular, we set the network output to zero at initialization, $f(x; \theta_0) = 0$,
 786 by (1) ‘centering’ the network by its initial value $f_c(x; \theta_t) = f(x; \theta) - f(x; \theta_0)$ or (2) setting the
 787 last layer weights to zero at initialization. We show that both (1) and (2) remove the opening of the
 788 sharpness reduction phase with $1/w$. Resultantly, the average onset of loss catapult occurs at $c_{loss} \approx 2$,
 789 independent of depth and width.

790 Throughout this section, we use ‘vanilla’ networks to refer to networks initialized in the standard
 791 way. For simplicity, we train FCNs using full batch gradient descent with MSE loss using a subset
 792 consisting of 4096 examples of the CIFAR-10 dataset.

793 G.1 The effect of centering networks

794 Given a network function $f(x; \theta_t)$, we define the centered network $f_c(x; \theta_t)$ as

$$f_c(x; \theta_t) = f(x; \theta_t) - f(x; \theta_0), \quad (44)$$




Figure 30: The phase diagrams of early training dynamics of (a, b, c) centered and (d, e, f) zero-init networks trained on CIFAR-10 using MSE using gradient descent. Each data point is an average over 10 initializations. The horizontal bars around the average data point indicate the region between 25% and 75% quantile.

795 where $f(x; \theta_0)$ is the network output at initialization. By construction, the network output is zero at
 796 initialization. It is noteworthy that centering a network is an unusual way of training deep networks
 797 as it doubles the cost of training because of two forward passes.

798 Figure 29 compares the training loss and sharpness dynamics of vanilla networks and centered
 799 networks. Unlike vanilla networks, we do not observe a decrease in sharpness for $c < c_{loss}$ during
 800 early training. Rather, we observe a slight increase in sharpness. To distinguish this slight increase
 801 from sharpness catapult, we introduce a threshold ϵ , comparing normalized sharpness $\lambda_t^H / \lambda_0^H$ with
 802 $1 + \epsilon$, to define a sharpness catapult.⁴ As demonstrated in Appendix G.3, the uv model trained on a
 803 single training example (x, y) with $y \neq 0$ sheds lights on this initial increase in sharpness.

804 Interestingly, irrespective of depth and width, we observe that loss catapults at $c_{loss} \approx 2$, as
 805 demonstrated in the phase diagrams in Figure 30(a, b, c). These findings suggest a strong correlation
 806 between a large network output at initialization $\|f(x; \theta_0)\|$ and the opening of the sharpness reduction
 807 phase discussed in Section 2.

808 G.2 The effect of setting the last layer to zero

809 An alternative way to train networks with $f(x; \theta_0) = 0$ is by setting the last layer to zero at
 810 initialization. The principle of criticality at initialization [52, 55, 65] does not put any constraints on
 811 the last layer weights. Hence, setting the last layer to zero does not affect signal/gradient propagation
 812 at initialization. Yet, setting the last layer to zero results in initialization in a flat curvature region
 813 at initialization, resulting in access to larger learning rates. We refer to these networks as ‘zero-init’
 814 networks.

815 Figure 29 compares the training dynamics of zero-init networks with vanilla and centered networks.
 816 We observe that the dynamics is quite similar to the centered networks: (i) sharpness does not reduce
 817 for small learning rates and (ii) loss catapults $c_{loss} \approx 2$, irrespective of depth and width. Figure 30(d,
 818 e, f) show the phase diagrams of networks with zero-initialized networks. Like centered networks,
 819 the critical constants do not scale with depth and width. Again, suggesting that a large network output

⁴In experiments, we set $\epsilon = 0.05$. We use the same threshold for zero-init networks.

820 at initialization $\|f(x; \theta_0)\|$ is related to the opening of the sharpness reduction phase in the early
 821 training results shown in Section 2

822 G.3 Insights from uv model trained on (x, y)

823 In this section, we gain insights into the effect of setting network output to zero at initialization using
 824 uv model trained on an example (x, y) . In particular, we show that loss catapults at $k_{loss} = 2$ and
 825 sharpness increases during early training.

826 Consider the uv model trained on a single training example (x, y) with $y \neq 0$ ⁵

$$f(x) = \frac{1}{\sqrt{w}} \sum_i^w u_i v_i x.$$

827 This simplifies the loss function to

$$\mathcal{L} = \frac{1}{2} (f(x) - y)^2 = \frac{1}{2} \Delta f^2, \quad (45)$$

828 where Δf is the residual. The trace of the Hessian $\text{tr}(H)$ is

$$\text{tr}(H) = \frac{x^2}{w} (\|v\|^2 + \|u\|^2). \quad (46)$$

829 The Frobenius norm can be written in terms of the trace and the network output

$$\|H\|_F^2 = \lambda^2 + 2x^2 \Delta f^2 \left(1 + \frac{2f}{w \Delta f}\right). \quad (47)$$

830 The function and residual updates are given by

$$f_{t+1} = f_t - \eta \text{tr}(H_t) + \frac{\eta^2 x^2}{w} f_t \Delta f_t^2 \quad (48)$$

$$\Delta f_{t+1} = \Delta f_t \left(1 - \eta \text{tr}(H_t) + \frac{\eta^2 x^2}{w} f_t \Delta f_t\right). \quad (49)$$

831 Similarly, we can obtain the trace update equations

$$\text{tr}(H_{t+1}) = \text{tr}(H_t) + \frac{\eta \Delta f_t^2 x^2}{w} \left(\eta \text{tr}(H_t) - 4 \frac{f_t}{\Delta f_t}\right). \quad (50)$$

832 Let us analyze them for the networks with zero output at initialization. The loss at the first step
 833 increases if

$$\left\langle \frac{L_1}{L_0} \right\rangle = \left\langle \left(1 - \eta \text{tr}(H_0) + \frac{\eta^2 x^2}{n} f_0 \Delta f_0\right)^2 \right\rangle > 1 \quad (51)$$

$$(52)$$

834 Setting $f_0 = 0$ and scaling the learning rate as $\eta = k / \text{tr}(H_0)$, we see that the loss increases at the
 835 first step if $k > 2$.

⁵Note that for $y = 0$, the network is already at a minimum.

$$\left\langle \frac{L_1}{L_0} \right\rangle = \langle (1 - k)^2 \rangle > 1 \quad (53)$$

836 Next, we analyze the change in trace during the first training step. Setting $f_0 = 0$, we observe that
 837 the trace increases for all learning rates

$$\text{tr}(H_1) = \text{tr}(H_0) + \frac{\eta^2 x^2}{w} \Delta f_0^2 \text{tr}(H_0), \quad (54)$$

838 modulated by the learning rate and width. Finally, we analyze the change in Frobenius norm in the
 839 first training step at $k = k_{loss}$, which implies $\Delta f_1^2 = \Delta f_0^2$,

$$\langle \Delta \|H_1\|^2 \rangle = \langle \text{tr}(H_1)^2 - \text{tr}(H_0)^2 + 2x^2 (\Delta f_1^2 - \Delta f_0^2) \rangle. \quad (55)$$

840 As $\text{tr}(H)$ increases in the first training step, $\|H\|_F$ also increases in the first training step.




Figure 31: The early training dynamics of FCNs with a fixed output scale trained on the CIFAR-10 dataset with MSE loss using gradient descent.

841 H The effect of output scale on the training dynamics

842 Given a neural network function $f(x)$ with depth d and width w , we define the scaled network as
843 $f_s(x) = \alpha f(x)$, where α is referred to as the output scale. In this section, we empirically study the
844 impact of the output scale on the early training dynamics. In particular, we show that a large (resp.
845 small) value of $\|f(x; \theta_0)\|$ relative to the one-hot encodings of the labels causes the sharpness to
846 decrease (resp. increase) during early training. Interestingly, we still observe an increase in $\langle c_{loss} \rangle$
847 with d and $1/w$, unlike the case of initializing network output to zero, highlighting the unique impact
848 of output scale on the dynamics. For simplicity, we train FCNs using gradient descent with MSE loss
849 using a subset consisting of 4096 examples of the CIFAR-10 dataset, as in the previous section.

850 H.1 The effect of fixed output scale at initialization

851 In this section, we study the training dynamics of models trained with a fixed output scale at
852 initialization. Given a network output function $f(\theta)$, we define the ‘scaled network’ as

$$f_s(\theta) = \frac{s f(\theta)}{\|f(\theta_0)\|}, \quad (56)$$

853 where s is a scalar, fixed throughout training. By construction, the network output norm $\|f_s(\theta_0)\|$
854 equals s . For standard initialization, $s = \|f(\theta_0)\| = \mathcal{O}(\sqrt{k})$, where k are the number of classes.

855 Figure 31 shows the training dynamics of FCNs for three different values of the output scale s . The
856 training dynamics of networks with $s = 1.0$ and $s = 10.0$ share qualitative similarities. In contrast,
857 networks initialized with a smaller output scale ($s = 0.1$) exhibit distinctly different dynamics. In
858 particular, we observe that for large output scales ($s \gtrsim 0.5$) sharpness decreases during early training,
859 while sharpness increases for small output scales⁶. Furthermore, the training dynamics tends to
860 be noisier at small output scales, making it difficult to characterize catapult dynamics amidst these
861 fluctuations. In summary, the training dynamics of networks with small output scale deviate from the
862 training dynamics discussed in the main text, particularly as the sharpness quickly increases during
863 early training.

864 Figure 32 shows the trends of various critical constants with width for FCNs for three different values
865 of s . Similar to vanilla networks, we observe that c_{loss} increases with d and $1/w$. In comparison,
866 sharpness decreases (increases) for large (small) values of s . These experiments suggest that the
867 output scale primarily influences the increase/decrease in sharpness during early training and does not
868 affect the scaling of c_{loss} with depth and width.

869 Note that we do not generate phase diagrams for these experiments as the training dynamics of
870 networks with small output scales at initialization deviate from the training dynamics discussed in the
871 main text.

872 H.2 Scaling the output scale with width

873 In this section, we study the training dynamics of models with an output scale scaled with width
874 as $\alpha = w^{-\sigma}$, which is commonly used in the literature [18, 6, 4]. We consider three distinct σ
875 values $\{-0.5, 0.0, 0.5\}$, where $\sigma = -0.5$ represents the lazy regime, $\sigma = 0.5$ corresponds to feature
876 learning (rich) regime and $\sigma = 0.0$ corresponds to standard (vanilla) initialization.

877 Figure 33 shows the training loss and sharpness trajectories of FCNs trained on for different σ values.
878 We observe that the training trajectories in the lazy regime look identical to standard initialization. In
879 comparison, the training trajectories in the feature learning regime is distinctly different. We observe
880 that in the standard and lazy regimes, sharpness decreases during early training, whereas sharpness
881 tends to increase in the feature learning regime and eventually oscillates around the edge of stability
882 regime. Moreover, we observe that sharpness can catapult before the training loss in the feature
883 learning regime (compare catapult peaks in 33(e, f)). These results are in parallel to the fixed output
884 scale networks studied in the previous section.

⁶We empirically observed that sharpness reduces for output scales as small as $s \sim 0.5$, which is relatively small compared to \sqrt{k} .




Figure 32: The phase diagrams of early training dynamics for ReLU FCNs with fixed output scale trained on a subset of the CIFAR-10 dataset using MSE loss using gradient descent. Each data point is an average over 10 initializations. The horizontal bars around the average data point indicate the region between 25% and 75% quantile.

885 Figure 34 summarizes the early training dynamics of FCNs with different σ values. We observe
 886 similar results as in the previous section. The output scale affects the initial increase/decrease of
 887 sharpness but does not affect the scaling trend of c_{loss} with depth and width. Moreover, we observe a
 888 systematic pattern of c_{max} scaling with width. In the lazy regime, we observe that c_{max} increases
 889 with $1/w$, while c_{max} decreases with $1/w$ in the feature learning regime.

890 I Sharpness curves in the intermediate saturation regime

891 This section shows additional results for Section 3 for MSE loss. Cross-entropy results are shown in
 892 Appendix F. Figures 35 to 39 show the normalized sharpness curves for different depths and widths.

893 I.1 Estimating the sharpness

894 This paragraph describes the procedure for measuring the sharpness to study the effect of the learning
 895 rate, depth, and width in the intermediate saturation regime. We measure the sharpness λ_τ^H at a time
 896 τ in the middle of the intermediate saturation regime. We choose τ so that $c\tau \approx 200$, for learning
 897 rates $c = 2^x$, where $x \in [-1.0, 4.0]$ in steps of 0.1. The value 200 is chosen such that τ is in the
 898 middle of the intermediate saturation regime. Next, we measure sharpness over a range of steps
 899 $t \in [\tau - 5, \tau + 5]$ and average over t to reduce fluctuations. We repeat this process for various
 900 initializations and obtain the average sharpness.

901 I.2 Estimating the critical constant c_{crit}

902 This subsection explains how to estimate c_{crit} from sharpness measured at time τ . First, we normalize
 903 the sharpness with its initial value, and then average over random initializations. Next, we estimate
 904 the critical point c_{crit} using the second derivative of the order parameter curve. Even if the obtained
 905 averaged normalized sharpness curve is somewhat smooth, the second derivative may become
 906 extremely noisy as minor fluctuations amplify on taking derivatives. This can cause difficulties in




Figure 33: The early training dynamics of FCNs with output scale $\alpha = w^{-\sigma}$ trained on the CIFAR-10 dataset with MSE loss using gradient descent.

907 obtaining c_{crit} . We resolve this issue by estimating the smooth derivatives of the averaged order
 908 parameter with the Savitzky–Golay filter [56] using its scipy implementation [60]. The estimated
 909 c_{crit} is shown by vertical lines in the sharpness curves in Figures 35 to 39

910 J The effect of batch size on the reported results

911 J.1 The early transient regime

912 Figure 40 shows the phase diagrams of early training dynamics of FCNs with $d = 4$ trained on the
 913 CIFAR-10 dataset using two different batch sizes. The phase diagram obtained is consistent with the
 914 findings presented in Section 2, except for one key difference. Specifically, we observe that when
 915 d/w is small and small batch sizes are used for training, sharpness may increase from initialization
 916 at relatively smaller values of c . This is reflected in Fig. 40 by $\langle c_{sharp} \rangle$ moving to the left as B
 917 is reduced from 512 to 128. However, this initial increase in sharpness is small compared to the
 918 sharpness catapult observed at larger batch sizes. We found that this increase at small batch sizes is
 919 due to fluctuations in gradient estimation that can cause sharpness to increase above its initial value
 920 by chance.




Figure 34: The phase diagrams of early training dynamics for ReLU FCNs with varying depths and output scale.




Figure 35: Sharpness measured at $c\tau = 200$ against the learning rate constant for FCNs trained on the MNIST dataset, with varying depths and widths. Each curve is an average over ten initializations, where the shaded region depicts the standard deviation around the mean trend. The vertical lines denote c_{crit} estimated using the maximum of χ'_τ .

921 J.2 The intermediate saturation regime

922 Figure 41 shows the normalized sharpness, measured at $c\tau = 200$, and its derivatives for various
 923 widths and batch sizes. The results are consistent with those in Section 3, with a lowering in the peak
 924 heights of the derivatives χ and χ' at small batch sizes. The lowering of the peak heights means the
 925 full width at half maximum increases, which implies a broadening of the transition around c_{crit} at
 926 smaller batch sizes.

927 K The effect of bias on the reported results

928 In this section, we show that FCNs with bias show similar results as presented in the main text. We
 929 considered FCNs in SP initialized with He initialization [28].




Figure 36: Sharpness measured at $c\tau = 200$ against the learning rate constant for FCNs trained on the Fashion-MNIST dataset, with varying depths and widths. Each curve is an average over ten initializations, where the shaded region depicts the standard deviation around the mean trend. The vertical lines denote c_{crit} estimated using the maximum of χ'_τ .




Figure 37: Sharpness measured at $c\tau = 200$ against the learning rate constant for FCNs trained on the CIFAR-10 dataset, with varying depths and widths. Each curve is an average over ten initializations, where the shaded region depicts the standard deviation around the mean trend. The vertical lines denote c_{crit} estimated using the maximum of χ'_τ .




Figure 38: Sharpness measured at $c\tau = 200$ against the learning rate constant for Myrtle-CNNs trained on the CIFAR-10 dataset, with varying depths and widths. Each curve is an average of over ten initializations, where the shaded region depicts the standard deviation around the mean trend. The vertical lines denote c_{crit} estimated using the maximum of χ'_τ .

930 Figure 42 shows the phase diagrams of early training for FCNs with bias trained on the CIFAR-10
 931 dataset. We observe a similar phase diagram compared to the no-bias case (compare with Figure 25).




Figure 39: Sharpness measured at $c\tau = 200$ against the learning rate constant for ResNets trained on the CIFAR-10 dataset, with varying depths and widths. Each curve is an average of over ten initializations, where the shaded region depicts the standard deviation around the mean trend. The vertical lines denote c_{crit} estimated using the maximum of χ'_τ .




Figure 40: The phase diagram of early training for FCNs with $d = 4$ trained on the CIFAR-10 dataset with MSE loss using SGD with different batch sizes.




Figure 41: (a) Normalized sharpness measured at $c\tau = 200$ against the learning rate constant for FCNs with $d = 4$ trained on the CIFAR-10 dataset, with varying widths. Each data point is an average over 10 initializations, where the shaded region depicts the standard deviation around the mean trend. (b, c) Smooth estimations of the first two derivatives, χ_τ and χ''_τ , of the averaged normalized sharpness wrt the learning rate constant.




Figure 42: The phase diagram of early training for FCNs with bias trained on the CIFAR-10 dataset with MSE loss using SGD with different depths.