

## 409 A Experimental Details

### 410 A.1 MLPs

411 In Figure 8 we used a 3-layer MLP learning a Gegenbauer polynomial  $Q_2(\beta \cdot \mathbf{x})$  in  $D = 5$  dimensions.  
412 Here  $\beta$  was a randomly chosen unit vector in  $\mathbb{R}^D$ . We implemented  $\mu\text{P}$  parameterization by hand.  
413 The output layer of the network was rescaled by  $\alpha_0/\sqrt{N}$ , consistent with  $\mu\text{P}$ . We chose  $\alpha_0 = 1000$   
414 to put us in the lazy regime. We set the the learning rate to be  $5N/(1 + \alpha_0^2)$ . General arguments  
415 based on kernel scale indicate that the learning rate should be scaled as  $\alpha_0^{-2}$  at large  $\alpha$ .

416 In Figure 13 we used a 3-layer MLP learning a Gegenbauer polynomial  $Q_2(\beta \cdot \mathbf{x})$  in  $D = 25$   
417 dimensions. We set the learning rate to be nearly as high as possible before a loss explosion.

### 418 A.2 Vision

#### 419 A.2.1 CIFAR-5m

420 All plots except Figure 5a and 5b. We trained with standard CIFAR data augmentation of random crop  
421 (RandomCrop(32, padding=4) in pytorch) and horizontal flip (RandomHorizontalFlip() in pytorch).  
422 As base network (for  $\mu\text{P}$ ) we used ResNet18 where BatchNorm was replaced with LayerNorm (to  
423 maintain the consistency of the neural network between train and test). We used the SGD optimizer  
424 with learning rate of .05 with cosine decay over 20000 steps, .9 momentum and batch size of 250.

425 For Figure 5a, we used the above setup, but with a learning rate of 0.01 and a much higher batch size  
426 of 2000, so as to replicate the edge of stability phenomenon [45] which only occurs at high batch  
427 sizes. For Figure 5b, we used a learning rate of 0.3 and batch size of 32, so as to show the behavior of  
428 high learning rate and small batch size on train loss.

#### 429 A.2.2 CIFAR-10 Multiple Passes

430 In Figure 10, we show the dynamics and representational consistency of ResNets trained on CIFAR-  
431 10 for several epochs. The architecture is a ResNet-18 with base-shape width set at  $N = 64$  channels.  
432 The model is trained with SGD with learning rate 0.1 and cosine annealing schedule. The batch-size  
433 used is 128.

#### 434 A.2.3 ImageNet

435 In all ImageNet experiments, we used a training subset of the ImageNet-1k dataset consisting of  
436  $2^{20} = 1048576$  labeled images and a test subset consisting of 1024 labeled images. Both subsets  
437 were randomly sampled from the full ImageNet-1k training and validation datasets, respectively. To  
438 extend the duration in training in which the network remains in the online regime beyond one epoch,  
439 we heavily augmented the images in the training dataset using PyTorch’s AutoAugment transform  
440 with the default policy, AutoAugmentPolicy.IMAGENET.

441 We again used the ResNet-18 architecture with  $\mu\text{P}$  parameterization relative to the ResNet-18  
442 network with base-shape width  $N = 64$  channel [14]. All architectures and training procedures were  
443 implemented in Jax and used the auxiliary Flax and Optax packages, respectively.

444 Figures 2(b) and 6(b) were trained using the Adam optimizer with the following learning rate schedule:  
445 linear warm-up for 0.5 epochs from learning rate  $8 \times 10^{-5}$  to  $8 \times 10^{-3}$ , followed by cosine decay  
446 over 49.5 epochs to  $8 \times 10^{-5}$ .

### 447 A.3 Language

#### 448 A.3.1 Wikitext-103 Language Modeling

449 For all Wikitext-103 tasks, we adopted the  $\mu\text{P}$  transformer as defined in the  $\mu\text{P}$  package [14]. In  
450 the plots shown in the main text, we used a depth-4 transformer, with  $d_{model}, d_k, d_v = N$  and  
451  $d_{ffn} = 4N$ . We performed a single pass through the train set in order to stay in the realistic online  
452 regime. We used a masked language modeling with sequence length  $S$  at varying input sequence  
453 lengths  $S$ . For Figure 1(d) we used the  $S \times S$  attention matrix of an  $S = 128$  transformer. In Figure  
454 4(e) we used the attention matrix of an  $S = 35$  transformer. We chose this different length simply to

455 illustrate the consistent message across sequence lengths. We used a batch size of  $B = 32$  for all  
456 experiments. The residual stream was thus a tensor of shape  $(S, B, d_{model})$ .

457 We used the Adam optimizer with a learning rate of 0.0001. We also ran the same configuration with  
458 SGD and a learning rate of 0.5 and observed the same behavior. See section B for further plots and  
459 details.

460 For figure 3, we used the Wikitext-103 validation set in order to measure the evolution of the  
461 predictions on masked logits. In 3f), we averaged the mean squared error from the widest transformer  
462 by using 100 test points.

### 463 A.3.2 C4 Language Modelling

464 Figure 1 (b) we trained with base network being a 125m parameter transformer model on  
465 2.5 billion tokens using the Mosaic ML’s LLM codebase (<https://web.archive.org/web/20230519184343/https://github.com/mosaicml/examples/tree/main/examples/llm>).  
466 See [https://web.archive.org/web/20230519183813/https://github.com/mosaicml/](https://web.archive.org/web/20230519183813/https://github.com/mosaicml/examples/blob/main/examples/llm/yamls/mosaic_gpt/125m.yaml)  
467 [examples/blob/main/examples/llm/yamls/mosaic\\_gpt/125m.yaml](https://web.archive.org/web/20230519183813/https://github.com/mosaicml/examples/blob/main/examples/llm/yamls/mosaic_gpt/125m.yaml) for the full hyperpa-  
468 rameter details. We were limited by time and computational resources in our ability to explore further  
469 details of the C4 transformer model.  
470

## 471 B Further Plots of Convergence

472 In this section, we show additional figures illustrating convergence of network quantities across  
473 widths that we did not have space for in the main text.

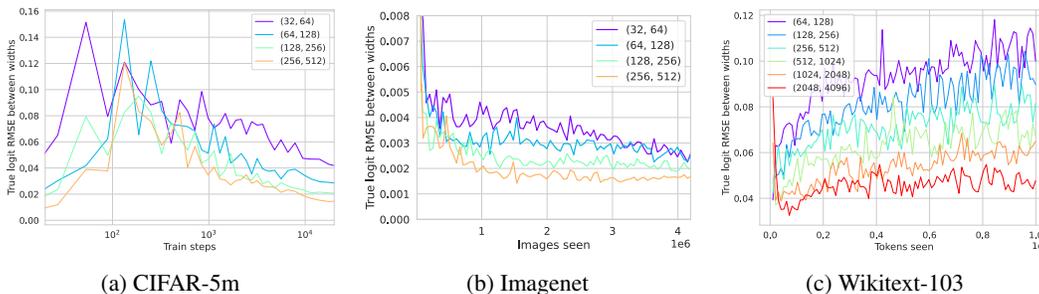


Figure 9: Analog of Figure 3 but comparing networks of successive widths rather than comparing all networks to the widest. Again, we see that as the network width grows, the difference between successive networks shrinks.

### 474 B.1 Vision

475 In Figure 9 (a), b) we plot the analogue of the first two columns of Figure 3, but now instead of  
476 computing logit RMSE to the widest network, we compute it between networks of successive widths  
477 for vision tasks.

478 A simple setting in which convergence properties are particularly clear and simple to study is for a  
479 ResNet learning CIFAR-10 and going over multiple passes of the dataset. In Figure 10 we plot a  
480 20-epoch pass over CIFAR 10, and study the generalization error, initial and final preactivations in  
481 the last layer, and final layer kernels across widths. The training error begins to exhibit pathologies  
482 after sufficiently many epochs, related to the discussion in section 3.2.

483 Next in figure 11, we show a higher-resolution plot of the kernel Gram matrices across widths and  
484 across layers for the CIFAR-5M ResNet after a pass through the data. The larger resolution allows  
485 one to see that even the fine-grained details in the structure of the Gram matrix are consistent across  
486 widths.

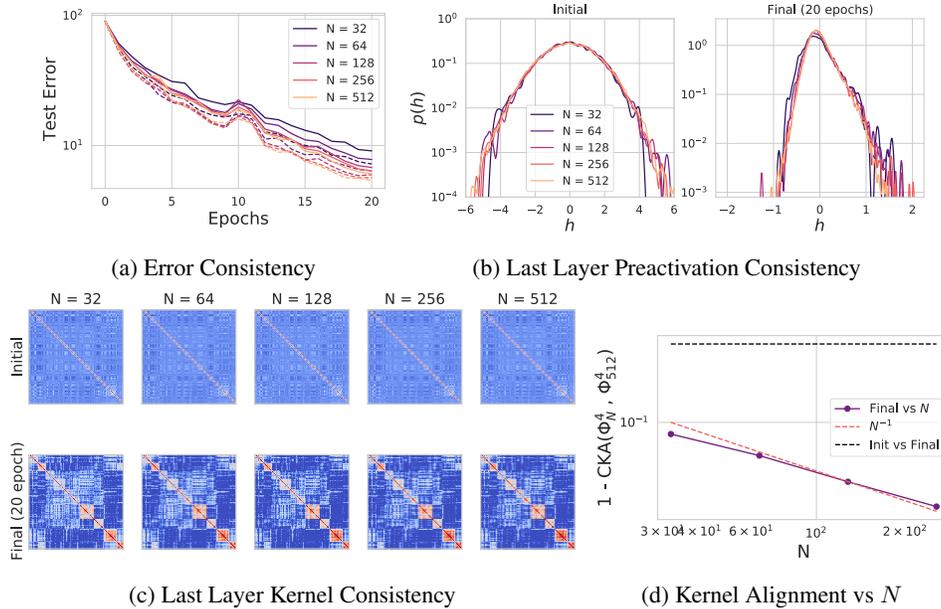


Figure 10: Training on CIFAR-10 in a ResNet-18 for multiple epochs generates dynamic preactivation densities and feature kernels which converge at realistic widths. (a) The test classification error curves for single models (solid) and ensembled (dashed) converge for realistic widths. (b) At initialization preactivation distributions in the last hidden layer of the CNN for a randomly chosen data point are Gaussian (as expected) and are very consistent across model widths  $N$ . To obtain histograms we train an ensemble of  $E = 8$  independently initialized networks concatenate activation patterns across members of the ensemble. After 20 epochs of training (models are around  $\sim 95\%$  accuracy), the preactivation distributions for the same data point have become non-Gaussian (consistent with infinite width theory) but are still remarkably consistent for large widths. (c) The final layer’s feature kernel at initialization shows very little structure, but (d) after training networks of all widths converge to similar kernels. The plot in (d) compares ensemble averaged kernels with the  $N = 512$  ensembled kernel.

487 **B.2 Language**

488 In Figure 9(c) We plot the RMSE difference between the values placed on the correct logit by networks  
 489 of successive width (rather than comparing to the widest as in Figure 3). We again see that as the  
 490 widths grow the differences shrink.

491 Next, in Figure 12, we create an analog of the language column of Figure 3, this time for  $\mu P$   
 492 transformers of the same architecture and dataset but now optimized with vanilla SGD. The fact that  
 493 wider transformers perform better still holds, and one can clearly see narrower networks approaching  
 494 wider ones in their output logit values.

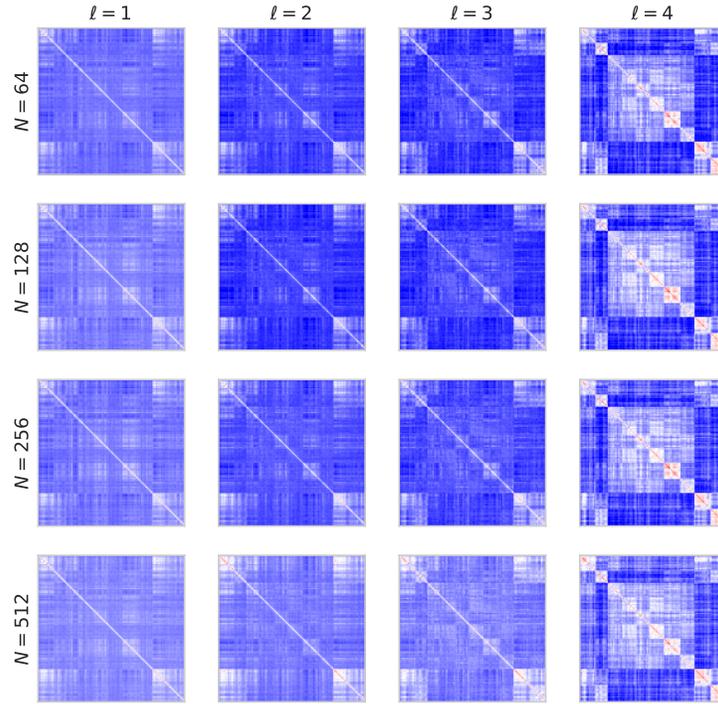


Figure 11: Convergence of layerwise representations in each layer (block)  $\ell$  of the ResNet-18 at large width  $N$  after training on CIFAR-5M.

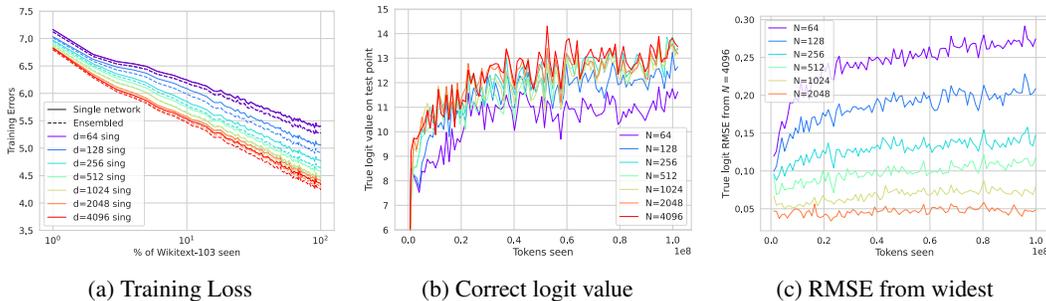


Figure 12: An analog of Figure 3 for  $\mu$ P transformers trained with SGD. a) Training loss. It is interesting that in  $\mu$  parameterization the SGD optimized network is competitive with the Adam-optimized network. b) Value placed on the correct logit for a specific masked token. c) RMSE of correct logit value from the widest network.

## 495 C Defining $\mu$ P and SP

496 There are several detailed discussions about  $\mu$ P vs SP scaling [10, 12, 7, 8, 11]. The aim of this  
 497 section is to simply give an accessible and conceptual overview of their distinction, as well as a  
 498 motivation for  $\mu$ P from the perspective of keeping features moving in time even at infinite width. .

499 There are several equivalent ways of parameterizing neural networks that give rise to the same dyn-  
 500 amical effects, whether in  $\mu$ -parameterization or standard parameterization. We give the definitions

501 in the case of a single-output feed-forward network and demonstrate that SP and  $\mu\text{P}$  give rise to  
 502  $O_N(N^{-1/2}), O_N(1)$  feature movement at initialization, respectively.

503 Generalizations to other architectures (ResNets, Transformers) are straightforward. For a detailed  
 504 discussion see [10] and also [11].

### 505 C.1 SP

506 We assume all hidden layers have equal width  $N$ . Let the input space have dimension  $D$ . Let  $\mu$  be  
 507 the index of the training point in the dataset. At each layer  $\ell$ , the pre-activation  $\mathbf{h}_\mu^{\ell+1}$  in layer  $\ell + 1$  is  
 508 given by

$$\mathbf{h}_\mu^{\ell+1} = \frac{1}{\sqrt{N}} \mathbf{W}^\ell \cdot \phi(\mathbf{h}_\mu^\ell), \quad (2)$$

509 where  $\phi$  is an element-wise non-linearity, often taken to be the ReLU function. Here the  $N^{-1/2}$  out  
 510 front allows  $\mathbf{h}_\mu^{\ell+1}$  to be  $O_N(1)$  at initialization as  $N \rightarrow \infty$  by the law of large numbers. The output  
 511 of the network  $f_\mu$  is then given by:

$$f_\mu = \frac{\alpha}{\sqrt{N}} \mathbf{w}^L \cdot \phi(\mathbf{h}_\mu^\ell). \quad (3)$$

512 Here again the  $N^{-1/2}$  scaling again yields that  $f_\mu$  will be  $O_N(1)$  as  $N \rightarrow \infty$ . In SP,  $\alpha$  is taken to be  
 513 1, but we will keep it explicit as it plays an important role in distinguishing the parameterizations. It  
 514 is the laziness parameter identified in [28]. The change in the function is given by

$$\frac{df_\mu}{dt} = -\eta \sum_\nu K_{\mu\nu} \ell'(f_\nu, y_\nu). \quad (4)$$

515 Here  $K_{\mu\nu} = \nabla_\theta f_\mu \cdot \nabla_\theta f_\nu$  is the NTK gram matrix.  $y^\nu$  is the true label.  $\ell$  is the loss function (e.g.  
 516 MSE or crossentropy) and  $\ell'$  is its derivative with respect to the first argument. The NTK is easily  
 517 seen to be order  $\alpha^2$  and  $\ell'$  is order 1 at small  $\alpha$ . In order to have the change in the function be  $O(1)$   
 518 we set  $\eta = \eta_0/\alpha^2$ .

519 Using the chain rule, one can directly see that the pre-activations evolve as [8, 3, 29]

$$\frac{d\mathbf{h}^\ell}{dt} \sim \eta \frac{\alpha}{\sqrt{N}} = \frac{\eta_0}{\alpha\sqrt{N}}. \quad (5)$$

520 Thus, at large  $N$  and  $\alpha = 1$  the pre-activations of this network evolve as  $O(N^{-1/2})$ . Consequently,  
 521 at infinite width the feature do not evolve and infinitely wide networks in standard parameterization  
 522 become kernel machines with the static and initialization-independent infinite-width NTK.

523 In many machine learning libraries, the factors of  $1/\sqrt{N}$  are not explicitly placed in front of each  
 524 multiplication with the weight matrices. Rather, the weight matrices themselves are drawn from a  
 525 distribution  $\mathbf{W}^\ell \sim \mathcal{N}(0, \frac{1}{N} \mathbf{1})$ . Although this gives identical forward pass, this changes the  $N$  scaling  
 526 of the gradients in the backward pass. As long as the learning rate is appropriately rescaled to account  
 527 for this, the dynamics are equivalent to the SP parameterization discussed above.

### 528 C.2 $\mu\text{P}$

529 One of the simplest ways to define the  $\mu$ -parameterization is to take  $\alpha = 1/\sqrt{N}$ . This implies that  
 530 we simply replace the final layer of the network by:

$$f_\mu = \frac{1}{N} \mathbf{w}^L \phi(\mathbf{h}_\mu^\ell). \quad (6)$$

531 As the prior analysis shows, in order to have  $df_\mu/dt$  be  $O(1)$  at initialization, we take  $\eta = N\eta_0$ , so  
 532 the learning rate in this definition scales extensively with  $N$ . In this setting, we now have that

$$\frac{d\mathbf{h}^\ell}{dt} \sim O_N(1). \quad (7)$$

533 In [10, 12], gives an equivalent definition of  $\mu\text{P}$  that gives rise to the same dynamics but keeps the  
 534 learning rate to be  $O_N(1)$ . We use this version of  $\mu\text{P}$  in the experiments that we run, simply because  
 535 that is what is used in the package [14]. Consequently, our learning rate does not need to be changed  
 536 as width grows.

537 **D  $\mu$ P versus Standard Parameterization**

538 **D.1 MLPs**

539 In figure 13, we show a 3-layer MLP learning a quadratic polynomial. In subfigure a) use a batch  
 540 size of 10 and a learning rate going as  $\eta = 5N/(1 + \alpha_0^2)$  with  $\alpha_0 = 1$ . The output layer is scaled  
 541 as  $\alpha_0/\sqrt{N}$ , putting us in the rich regime. The learning rate has been picked to be nearly as large  
 542 as possible at this batch size in order to maximize the large loss curve fluctuations yielded by large  
 543 learning rate effects. In subfigure c) we do not rescale the output layer, and have a width-indepdent  
 544 learning rate going as  $\eta = 50/(1 + \alpha_0^2)$  with  $\alpha_0 = 1$ . This puts is in the large-learning rate regime  
 545 for a standard parameterized network. See Appendix A for more details.

546 We plot the learning curves across widths and find striking agreement, even at the fine-grained level  
 547 of fluctuations due to small batch size and large learning rate effects. Although this is not exactly the  
 548 full-batch edge-of-stability effect reported in [45], the large oscillations may be similar to a small  
 549 batch size analog. We plot the absolute difference from the widest network in subfigure b) to highlight  
 550 the strong agreement across widths.

551 In subfigure c), we have the same network but in standard parameterization. The narrower networks  
 552 now learn features more quickly, leading to inconsistent dynamics across widths.

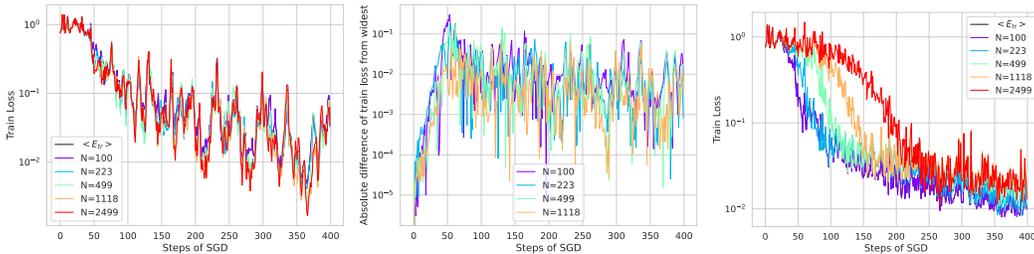


Figure 13: 3-Layer MLP learning a quadratic polynomial  $y = Q_2(\beta \cdot \mathbf{x})$  on the sphere. Data is provided in an online setting in the same order across widths, as in the realistic experiments in the main text. Large learning rate small batch size effects in MLPs are consistent across large widths. a) For  $\mu$ P networks, the loss curves match across widths, even accounting for fluctuations due to batch size or large learning rate edge-of-stability-like effects. b) Plot of the difference in training error from the widest network. c) The same network but in standard parameterization. Dynamics are no longer consistent across widths, and wider networks approach a lazy limit.

553 **D.2 Vision**

554 Next, we focus on a vision task and compare the large learning rate small batch size effects in SP to  
 555 the  $\mu$ P parameterized network in Figure 5a. By contrast to that figure, we see significantly different  
 556 dynamics and batch variation across widths. In Figure 14 a) we plot the early time behavior of a  
 557 CIFAR-5m task at large learning rate. The large learning rate effects cause the loss to substantially  
 558 oscillate, but the oscillations across widths are inconsistent by contrast to 5a b). Further, at late  
 559 times in Figure 14, the sharp spikes in the loss function due to large learning rate effects become  
 560 substantially different across widths. Indeed, in SP some widths converge for a given learning rate  
 561 while others do not. This trend has already been well-studied in [12]. We again stress that our  
 562 observation is that not only are the final losses similar across widths in  $\mu$ P (as observed in [12]), but  
 563 that the individual batch and large learning rate fluctuations agree across widths at early times in  $\mu$ P  
 564 as well.

565 **D.3 Language**

566 Finally, we present a complementary set of figures to those in the right columns of Figures 3 and 12  
 567 for transformers of the same architecture on Wikitext-103 but in standard parameterization.

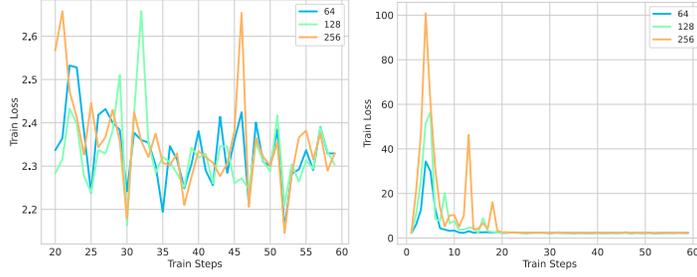


Figure 14: Different widths have different loss curves. a) Early time dynamics of the loss across widths is not consistent. b) Dynamics of the loss across widths at later times also does not appear consistent. There are explosions that happen at different times and scales across widths.

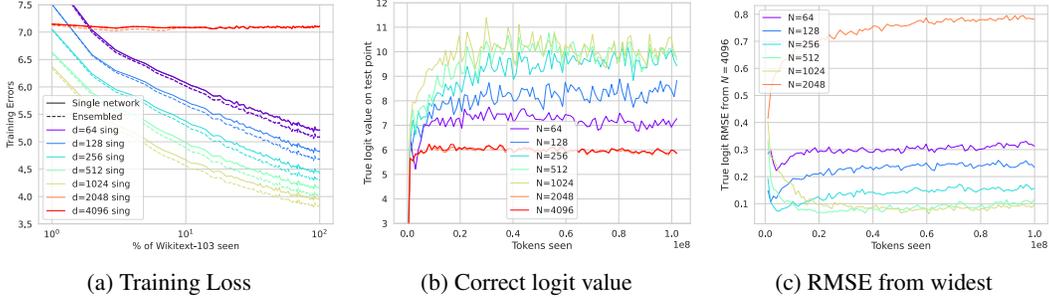


Figure 15: An analog of Figure 3. SP transformers trained with Adam. a) Training loss. At large widths, the learning rate chosen is too big for the network to properly learn, and the loss is flat. This is consistent with what is observed in [12] — the optimal learning rate in SP changes with width. b) Value placed on the correct logit for a specific masked token. c) RMSE of correct logit value from the widest network. In both of these plots, the monotonic behavior across width evident under the  $\mu P$  parameterization is violated. Even after discarding the networks that do not converge under SP, the behavior remains non-monotonic across width.

## 568 E Overview of Finite Width Corrections to Feature Learning Networks

569 In this section we review some basic ideas from the mean field theory of feature learning neural  
 570 networks. We first describe the predictions that mean field theory makes about infinite width networks  
 571 before describing finite size corrections to the dynamics of learning. To eliminate unnecessary  
 572 complexity, we will focus on MLP layers, but these arguments can be easily extended to CNN and  
 573 self-attention layers as well. We start by defining a MLP in a parameterization equivalent to  $\mu P$

$$f_\mu = \frac{1}{N} \mathbf{w}^L \cdot \phi(\mathbf{h}_\mu^\ell), \quad \mathbf{h}_\mu^{\ell+1} = \frac{1}{\sqrt{N}} \mathbf{W}^\ell \phi(\mathbf{h}_\mu^\ell), \quad \mathbf{h}_\mu^1 = \frac{1}{\sqrt{D}} \mathbf{W}^0 \mathbf{x}_\mu. \quad (8)$$

574 We will consider these networks trained from a random Gaussian initialization of the weights so  
 575 that  $\boldsymbol{\theta} = \text{Vec}\{\mathbf{w}^L, \dots, \mathbf{W}^0\}$  follows  $\boldsymbol{\theta} \sim \mathcal{N}(0, \mathbf{I})$  at initialization. This network is then trained with  
 576 some gradient based optimizer, leading to dynamical predictions  $f_\mu(t)$  and dynamical preactivations  
 577  $\mathbf{h}_\mu^\ell(t)$ . Because of the random initialization of weights, the outputs of the network and the precise  
 578 preactivations are random variables. However, at infinite width  $N \rightarrow \infty$ , a dramatic simplification of  
 579 the dynamics occurs.

### 580 E.1 The Infinite Width/Mean Field Limit

581 The predictions  $f_\mu(t)$  and internal representations of infinite width limit of neural networks admit a  
 582 description in terms of non-random initialization-independent dynamical feature kernels  $\Phi_{\mu\nu}^\ell(t, s)$   
 583 and gradient kernels  $G_{\mu\nu}^\ell(t, s)$  defined as

$$\Phi_{\mu\nu}^\ell(t, s) = \frac{1}{N} \phi(\mathbf{h}_\mu^\ell(t)) \cdot \phi(\mathbf{h}_\nu^\ell(s)), \quad G_{\mu\nu}^\ell(t, s) = \frac{1}{N} \mathbf{g}_\mu^\ell(t) \cdot \mathbf{g}_\nu^\ell(s), \quad (9)$$

584 where  $\mu, \nu$  index data points and  $t, s$  index training time and  $\mathbf{g}_\mu^\ell(t) = N \frac{\partial f_\mu}{\partial \mathbf{h}^\ell}$  are back-propagated  
 585 gradient signals [10, 11]. Further, all preactivation vectors  $\mathbf{h}_\mu^\ell(t) \in \mathbb{R}^N$  have entries that become iid  
 586 draws from a (potentially non-Gaussian) single site density  $p(h)$ , which converges as

$$\frac{1}{N} \sum_{i=1}^N \delta(h - h_i) \rightarrow p(h), \quad (10)$$

587 which should be understood in terms of integration of these densities against test functions. At infinite  
 588 width, the sums over neurons in a layer can be replaced by deterministic integrals over this single site  
 589 density  $\Phi_{\mu\nu}^\ell(t, s) = \int p(h_\mu^\ell(t), h_\nu^\ell(s)) \phi(h_\mu^\ell(t)) \phi(h_\nu^\ell(s)) dh_\mu^\ell(t) dh_\nu^\ell(s)$ .

## 590 E.2 Finite Width Effects

591 At finite width, the internal kernels  $\{\Phi_{\mu\nu}^\ell(t, s), G_{\mu\nu}^\ell(t, s)\}$  and predictions  $f_\mu(t)$  of the model become  
 592 initialization and width-dependent and deviate from their mean field dynamics. For Gaussian random  
 593 initialization of the weights of the network, the predictions and kernels fluctuate (from init to init)  
 594 with variance that scales asymptotically like  $\mathcal{O}(1/N)$  for width  $N$  (or  $1/d_{model}$  for transformer)  
 595 [36]. Further, the *ensemble averaged* values for the predictions  $\langle f_\mu(t) \rangle$  and kernels  $\langle \Phi_{\mu\nu}^\ell(t, s) \rangle$  differ  
 596 asymptotically from their infinite width values by  $\mathcal{O}(N^{-1})$ . Both of these two leading order effects  
 597 can influence the expected (train or test) loss of the model. At fixed width and late training time,  
 598 finite size effects beyond leading order can accumulate and become relevant, however theory predicts  
 599 that any observable average at width  $N$  admits an asymptotic series in powers of  $N^{-1}$  [36].

### 600 E.2.1 Trainability at Finite Size

601 The  $\mathcal{O}(N^{-1})$  correction to feature and gradient kernels can lead to non-trivial corrections to the  
 602 loss dynamics. Working in continuous time, we can define the neural tangent kernel (NTK) as  
 603  $K_{\mu\nu}(t) = \sum_\ell G_{\mu\nu}^{\ell+1}(t, t) \Phi_{\mu\nu}^\ell(t, t)$ , where base cases are  $\Phi_{\mu\nu}^0(t, s) = \frac{1}{D} \mathbf{x}_\mu \cdot \mathbf{x}_\nu$  and  $G_{\mu\nu}^{L+1}(t, s) = 1$ .  
 604 Following the approximation to online dynamics with MSE loss in Section 4, we consider a gradient  
 605 flow on the average dynamical NTK

$$\frac{d}{dt} \Delta(t) = - \langle \mathbf{K}(t) \rangle \Delta(t) \implies \Delta(t) = \mathcal{T} \exp \left( - \int_0^t ds \langle \mathbf{K}(s) \rangle \right) \mathbf{y}, \quad (11)$$

606 where  $\mathcal{T}$  is the time-ordering operator. We now consider the leading correction to the average NTK  
 607 around infinite width  $\langle \mathbf{K}(t) \rangle = \mathbf{K}_\infty(t) + \frac{1}{N} \mathbf{K}^1(t) + \mathcal{O}(N^{-2})$ . With this correction, we see that  
 608 the dynamics of errors  $\Delta$

$$\Delta(t) = \mathcal{T} \exp \left( - \int_0^t ds \mathbf{K}_\infty(s) - \frac{1}{N} \int_0^t ds \mathbf{K}^1(s) + \mathcal{O}(N^{-2}) \right) \mathbf{y}. \quad (12)$$

609 The fact that the  $\frac{1}{N} \mathbf{K}^1$  correction is integrated over time and placed in the matrix exponential  
 610 indicates that small corrections to NTK dynamics can lead to large dynamical amplification of logit  
 611 corrections. This fact was pointed out in another work [36] which tried to motivate a study of  
 612 perturbation theory in logarithms of the transition matrix  $\mathbf{T}(t)$  defined as

$$\frac{d}{dt} \mathbf{T}(t) = - \langle \mathbf{K}(t) \rangle \mathbf{T}(t), \quad \mathbf{T}(0) = \mathbf{I}, \quad \mathbf{R}(t) = \log \mathbf{T}(t). \quad (13)$$

613 The solution to this can be used to construct the errors at a later time  $\Delta(t) = \exp(\mathbf{R}(t)) \mathbf{y}$ .

## 614 F Offline Training

615 Figure 16 depicts the loss curve for a ConvNeXt-T (tiny) model trained on ImageNet in the typical,  
 616 offline setting — where data is encountered repeatedly across many epochs. As the networks overfit  
 617 the training data — in Figure 16, beyond 40,000 training steps or five epochs — the loss curves  
 618 diverge dramatically for different-width networks. Width consistency subsequently erodes.

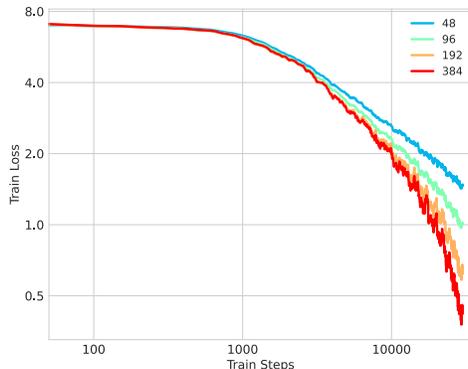


Figure 16: Consistency of loss curves across widths in the beginning and separation as loss becomes sufficiently small in offline learning.

## 619 G Use of Compute

620 For most experiments, we used Nvidia A100 SXM4 40GB and 80 GB GPUs on an academic cluster.

621 For the Wikitext-103 tasks, each width included 4 ensembles loaded onto an A100 GPU that ran for  
 622 a range between 1 to 3 days. For each sweep over widths this corresponds to about 8 A100-days.  
 623 Accounting for sweeps over different sequence lengths, optimizers, and parameterizations, this  
 624 corresponds to about 50 A100-days.

625 All MLP tasks, including the calculation of empirical NTKs and their spectral properties were done  
 626 in 15-30 minute Colab sessions using the basic GPUs provided.

627 The CIFAR-10 ResNet experiments in Figure 10 were done using a total of less than 1 A100-day of  
 628 compute across all widths and ensembles.

629 The ImageNet ResNet experiments vectorize training over between one to four same-width neural  
 630 networks on one A100 GPU. Each experiment training a collection of networks for 30 epochs takes  
 631 between one to three A100-days. Overall, these experiments expended roughly 30 A100-days.

632 For the CIFAR-5m experiments in Figure 2 and 3, across all widths, it required a few hours of A100  
 633 GPU. For Figure 6 and 7, as these were ensembled across multiple runs, these required close to 1-2  
 634 days of A100-GPUs. Figure 5a was just run for a few 100 steps of the training, so didn't use much  
 635 compute power.