

THE INFLUENCE OF LEARNING RULE ON REPRESENTATION DYNAMICS IN WIDE NEURAL NETWORKS

Blake Bordelon & Cengiz Pehlevan

School of Engineering and Applied Science

Harvard University

Cambridge, MA 02138, USA

{blake_bordelon, cpehlevan}@g.harvard.edu

ABSTRACT

It is unclear how changing the learning rule of a deep neural network alters its learning dynamics and representations. To gain insight into the relationship between learned features, function approximation, and the learning rule, we analyze infinite-width deep networks trained with gradient descent (GD) and biologically-plausible alternatives including feedback alignment (FA), direct feedback alignment (DFA), and error modulated Hebbian learning (Hebb), as well as gated linear networks (GLN). We show that, for each of these learning rules, the evolution of the output function at infinite width is governed by a time varying effective neural tangent kernel (eNTK). In the lazy training limit, this eNTK is static and does not evolve, while in the rich mean-field regime this kernel’s evolution can be determined self-consistently with dynamical mean field theory (DMFT). This DMFT enables comparisons of the feature and prediction dynamics induced by each of these learning rules. In the lazy limit, we find that DFA and Hebb can only learn using the last layer features, while full FA can utilize earlier layers with a scale determined by the initial correlation between feedforward and feedback weight matrices. In the rich regime, DFA and FA utilize a temporally evolving and depth-dependent NTK. Counterintuitively, we find that FA networks trained in the rich regime exhibit more feature learning if initialized with smaller correlation between the forward and backward pass weights. GLNs admit a very simple formula for their lazy limit kernel and preserve conditional Gaussianity of their preactivations under gating functions. Error modulated Hebb rules show very small task-relevant alignment of their kernels and perform most task relevant learning in the last layer.

1 INTRODUCTION

Deep neural networks have now attained state of the art performance across a variety of domains including computer vision and natural language processing (Goodfellow et al., 2016; LeCun et al., 2015). Central to the power and transferability of neural networks is their ability to flexibly adapt their layer-wise internal representations to the structure of the data distribution during learning.

In this paper, we explore how the learning rule that is used to train a deep network affects its learning dynamics and representations. Our primary motivation for studying different rules is that exact gradient descent (GD) training with the back-propagation algorithm is thought to be biologically implausible (Crick, 1989). While many alternatives to standard GD training were proposed (Whittington & Bogacz, 2019), it is unclear how modifying the learning rule changes the functional inductive bias and the learned representations of the network. Further, understanding the learned representations could potentially offer more insight into which learning rules account for representational changes observed in the brain (Poort et al., 2015; Kriegeskorte & Wei, 2021; Schumacher et al., 2022). Our current study is a step towards these directions.

The alternative learning rules we study are error modulated Hebbian learning (Hebb), Feedback alignment (FA) (Lillicrap et al., 2016) and direct feedback alignment (DFA) (Nøkland, 2016). These rules circumvent one of the biologically implausible features of GD: the weights used in the backward pass computation of error signals must be dynamically identical to the weights used on the

forward pass, known as the weight transport problem. Instead, FA and DFA algorithms compute an approximate backward pass with independent weights that are frozen through training. Hebb rule only uses a global error signal. While these learning rules do not perform exact GD, they are still able to evolve their internal representations and eventually fit the training data. Further, experiments have shown that FA and DFA can scale to certain problems such as view-synthesis, recommendation systems, and small scale image problems (Launay et al., 2020), but they do not perform as well in convolutional architectures with more complex image datasets (Bartunov et al., 2018). However, significant improvements to FA can be achieved if the feedback-weights have partial correlation with the feedforward weights (Xiao et al., 2018; Moskovitz et al., 2018; Boopathy & Fiete, 2022).

We also study gated linear networks (GLNs), which use frozen gating functions for nonlinearity (Fiat et al., 2019). Variants of these networks have bio-plausible interpretations in terms of dendritic gates (Sezener et al., 2021). Fixed gating can mitigate catastrophic forgetting (Veness et al., 2021; Budden et al., 2020) and enable efficient transfer and multi-task learning Saxe et al. (2022).

Here, we explore how the choice of learning rule modifies the representations, functional biases and dynamics of deep networks at the infinite width limit, which allows a precise analytical description of the network dynamics in terms of a collection of evolving kernels. At infinite width, the network can operate in the *lazy* regime, where the feature embeddings at each layer are constant through time, or the *rich/feature-learning* regime (Chizat et al., 2019; Yang & Hu, 2021; Bordelon & Pehlevan, 2022). The richness is controlled by a scalar parameter related to the initial scale of the output function.

In summary, our novel contributions are the following:

1. We identify a class of learning rules for which function evolution is described by a dynamical effective Neural Tangent Kernel (eNTK). We provide a dynamical mean field theory (DMFT) for these learning rules which can be used to compute this eNTK. We show both theoretically and empirically that convergence to this DMFT occurs at large width N with error $O(N^{-1/2})$.
2. We characterize precisely the inductive biases of infinite width networks in the lazy limit by computing their eNTKs at initialization. We generalize FA to allow partial correlation between the feedback weights and initial feedforward weights and show how this alters the eNTK.
3. We then study the *rich* regime so that the features are allowed to adapt during training. In this regime, the eNTK is dynamical and we give a DMFT to compute it. For deep linear networks, the DMFT equations close algebraically, while for nonlinear networks we provide a numerical procedure to solve them.
4. We compare the learned features and dynamics among these rules, analyzing the effect of richness, initial feedback correlation, and depth. We find that rich training enhances gradient-pseudogradient alignment for both FA and DFA. Counterintuitively, smaller initial feedback correlation generates more dramatic feature evolution for FA. The GLN networks have dynamics comparable to GD, while Hebb networks, as expected, do not exhibit task relevant adaptation of feature kernels, but rather evolve according to the input statistics.

1.1 RELATED WORKS

GLNs were introduced by Fiat et al. (2019) as a simplified model of ReLU networks, allowing the analysis of convergence and generalization in the lazy kernel limit. Veness et al. (2021) provided a simplified and biologically-plausible learning rule for deep GLNs which was extended by Budden et al. (2020) and provided an interpretation in terms of dendritic gating Sezener et al. (2021). These works demonstrated benefits to continual learning due to the fixed gating. Saxe et al. (2022) derived exact dynamical equations for a GLN with gates operating at each node and each edge of the network graph. Krishnamurthy et al. (2022) provided a theory of gating in recurrent networks.

Lillicrap et al. (2016) showed that, in a two layer linear network the forward weights will evolve to align to the frozen feedback weights under the FA dynamics, allowing convergence of the network to a loss minimizer. This result was extended to deep networks by Frenkel et al. (2019), who also introduced a variant of FA where only the direction of the target is used. Refinetti et al. (2021) studied DFA in a two-layer student-teacher online learning setup, showing that the network first undergoes an alignment phase before converging to one the degenerate global minima of the loss. They argued that FA’s worse performance in CNNs is due to the inability of the forward pass gradients to align under the block-Toeplitz connectivity structure that arises from enforced weight sharing (d’Ascoli et al., 2019). Garg & Vempala (2022) analyzed matrix factorization with FA,

proving that, when overparameterized, it converges to a minimizer under standard conditions, albeit more slowly than GD. Cao et al. (2020) analyzed the kernel and loss dynamics of linear networks trained with learning rules from a space that includes GD, contrastive Hebbian, and predictive coding rules, showing strong dependence of hierarchical representations on learning rule.

Recent works have utilized DMFT techniques to analyze typical performance of algorithms trained on high-dimensional random data (Agoritsas et al., 2018; Mignacco et al., 2020; Celentano et al., 2021; Gerbelot et al., 2022). In the present work, we do not average over random datasets, but rather over initial random weights and treat data as an input to the theory. Wide NNs have been analyzed at infinite width in both lazy regimes with the NTK (Jacot et al., 2018; Lee et al., 2019) and rich feature learning regimes (Mei et al., 2018). In the feature learning limit, the evolution of kernel order parameters have been obtained with both Tensor Programs framework (Yang & Hu, 2021) and with DMFT (Bordelon & Pehlevan, 2022). Song et al. (2021) recently analyzed the lazy infinite width limit of two layer networks trained with FA and weight decay, finding that only one layer effectively contributes to the two-layer NTK. Boopathy & Fiete (2022) proposed alignment based learning rules for networks at large width in the lazy regime, which performs comparably to GD and outperform standard FA. Their *Align-Ada* rule corresponds to our ρ -FA with $\rho = 1$ in *lazy* large width networks.

2 EFFECTIVE NEURAL TANGENT KERNEL FOR A LEARNING RULE

We denote the output of a neural network for input $\mathbf{x}_\mu \in \mathbb{R}^D$ as f_μ . For concreteness, in the main text we will focus on scalar targets $f_\mu \in \mathbb{R}$ and MLP architectures. Other architectures such as multi-class outputs and CNN architectures with infinite channel count can also be analyzed as we show in the Appendix C. For the moment, we let the function be computed recursively from a collection of weight matrices $\boldsymbol{\theta} = \text{Vec}\{\mathbf{W}^0, \mathbf{W}^1, \dots, \mathbf{w}^L\}$ in terms of preactivation vectors $\mathbf{h}_\mu^\ell \in \mathbb{R}^N$ where,

$$f_\mu = \frac{1}{\gamma_0 N} \mathbf{w}^L \cdot \phi(\mathbf{h}_\mu^L), \quad \mathbf{h}_\mu^{\ell+1} = \frac{1}{\sqrt{N}} \mathbf{W}^\ell \phi(\mathbf{h}_\mu^\ell), \quad \mathbf{h}_\mu^1 = \frac{1}{\sqrt{D}} \mathbf{W}^0 \mathbf{x}_\mu \quad (1)$$

where nonlinearity ϕ is applied element-wise. The scalar parameter γ_0 controls how rich the network training is: small γ_0 corresponds to lazy learning while large γ_0 generates large changes to the features (Chizat et al., 2019). For gated linear networks, we follow Fiat et al. (2019) and modify the forward pass equations by replacing $\phi(\mathbf{h}_\mu^\ell)$ with a multiplicative gating function $\dot{\phi}(\mathbf{m}_\mu^\ell) \mathbf{h}_\mu^\ell$ where gating variables $\mathbf{m}_\mu^\ell = \frac{1}{\sqrt{D}} \mathbf{M}^\ell \mathbf{x}_\mu$ are fixed through training with $M_{ij} \sim \mathcal{N}(0, 1)$. To minimize loss $\mathcal{L} = \sum_\mu \ell(f_\mu, y_\mu)$, we consider learning rules to the parameters $\boldsymbol{\theta}$ of the form

$$\frac{d}{dt} \mathbf{w}^L = \gamma_0 \sum_\mu \phi(\mathbf{h}_\mu^L(t)) \Delta_\mu, \quad \frac{d}{dt} \mathbf{W}^\ell = \frac{\gamma_0}{\sqrt{N}} \sum_\mu \Delta_\mu \tilde{\mathbf{g}}_\mu^{\ell+1} \phi(\mathbf{h}_\mu^\ell)^\top, \quad \frac{d}{dt} \mathbf{W}^0 = \frac{\gamma_0}{\sqrt{D}} \sum_\mu \Delta_\mu \tilde{\mathbf{g}}_\mu^1 \mathbf{x}_\mu^\top \quad (2)$$

where the error signal is $\Delta_\mu(t) = -\frac{\partial \mathcal{L}}{\partial f_\mu} |_{f_\mu(t)}$. The last layer weights \mathbf{w}^L are always updated with their true gradient. This corresponds to the biologically-plausible and local delta-rule, which merely correlates the error signals Δ_μ and the last layer features $\phi(\mathbf{h}_\mu^L)$ (Widrow & Hoff, 1960). In intermediate layers, the *pseudo-gradient* vectors $\tilde{\mathbf{g}}_\mu^\ell$ are determined by the choice of the learning rule. For concreteness, we provide below the recursive definitions of $\tilde{\mathbf{g}}^\ell$ for our five learning rules of interest.

$$\tilde{\mathbf{g}}_\mu^\ell = \begin{cases} \dot{\phi}(\mathbf{h}_\mu^\ell) \odot \left[\frac{1}{\sqrt{N}} \mathbf{W}^\ell(t)^\top \tilde{\mathbf{g}}_\mu^{\ell+1} \right], & \tilde{\mathbf{g}}_\mu^L = \dot{\phi}(\mathbf{h}_\mu^L) \odot \mathbf{w}^L & \text{GD} \\ \dot{\phi}(\mathbf{h}_\mu^\ell) \odot \left[\frac{1}{\sqrt{N}} \left(\rho \mathbf{W}^\ell(0) + \sqrt{1 - \rho^2} \tilde{\mathbf{W}}^\ell \right)^\top \tilde{\mathbf{g}}_\mu^{\ell+1} \right], & \tilde{W}_{ij}^\ell \sim \mathcal{N}(0, 1) & \rho\text{-FA} \\ \dot{\phi}(\mathbf{h}_\mu^\ell) \odot \tilde{\mathbf{z}}^\ell, & \tilde{z}_i^\ell \sim \mathcal{N}(0, 1) & \text{DFA} \\ \dot{\phi}(\mathbf{m}_\mu^\ell) \odot \left[\frac{1}{\sqrt{N}} \mathbf{W}^\ell(t)^\top \tilde{\mathbf{g}}_\mu^{\ell+1} \right], & \tilde{\mathbf{g}}^L = \dot{\phi}(\mathbf{m}_\mu^L) \odot \mathbf{w}^L(t) & \text{GLN} \\ \Delta_\mu(t) \phi(\mathbf{h}_\mu^\ell(t)) & & \text{Hebb} \end{cases} \quad (3)$$

While GD uses the instantaneous feedforward weights on the backward pass, ρ -FA uses the weight matrices which do not evolve throughout training. These weights have correlation ρ with the initial forward pass weights $\mathbf{W}^\ell(0)$. This choice is motivated by the observation that partial correlation between forward and backward pass weights at initialization can improve training (Liao et al., 2016;

Xiao et al., 2018; Moskowitz et al., 2018), though the cost is partial weight transport at initialization. However, we consider partial correlation at initialization more biologically plausible than the demanding weight transport at each step of training, like in GD. For DFA, the weight vectors \tilde{z}^ℓ are sampled randomly at initialization and do not evolve in time. For GLN, the gating variables \mathbf{m}_μ^ℓ are frozen through time but the exact feedforward weights are used in the backward pass. Lastly, we modify the classic Hebb rule (Hebb, 1949) to get $\Delta \mathbf{W}^\ell \propto \sum_\mu \Delta_\mu(t)^2 \phi(\mathbf{h}_\mu^{\ell+1}) \phi(\mathbf{h}_\mu^\ell)^\top$, which weighs each example by its current error. Unlike standard Hebbian updates, this learning rule gives stable dynamics without regularization (App. G). For all rules, the evolution of the function is determined by a time-dependent eNTK $K_{\mu\nu}$ which is defined as

$$\begin{aligned} \frac{\partial f_\mu}{\partial t} &= \frac{\partial f_\mu}{\partial \boldsymbol{\theta}} \cdot \frac{d\boldsymbol{\theta}}{dt} = \sum_\nu \Delta_\nu K_{\mu\nu}(t, t), \quad K_{\mu\nu}(t, s) = \sum_{\ell=0}^L \tilde{G}_{\mu\nu}^{\ell+1}(t, s) \Phi_{\mu\nu}^\ell(t, s) \\ \tilde{G}_{\mu\nu}^\ell(t, s) &= \frac{1}{N} \mathbf{g}_\mu^\ell(t) \cdot \tilde{\mathbf{g}}_\nu^\ell(s), \quad \Phi_{\mu\nu}^\ell(t, s) = \frac{1}{N} \phi(\mathbf{h}_\mu^\ell(t)) \cdot \phi(\mathbf{h}_\nu^\ell(s)), \end{aligned} \quad (4)$$

where the base cases $\tilde{G}_{\mu\nu}^{L+1}(t, s) = 1$ and $\Phi_{\mu\nu}^0(t, s) = \frac{1}{D} \mathbf{x}_\mu \cdot \mathbf{x}_\nu$ are time-invariant. The kernel $\tilde{\mathbf{G}}^\ell$ computes an inner product between the true gradient signals $\mathbf{g}_\mu^\ell = \gamma_0 N \frac{\partial f_\mu}{\partial \mathbf{h}_\mu^\ell}$ and the pseudo-gradient $\tilde{\mathbf{g}}_\nu^\ell$ which is set by the chosen learning rule. We see that because $\tilde{\mathbf{G}}^\ell$ is not necessarily symmetric, \mathbf{K} is also not necessarily symmetric. The matrix $\tilde{\mathbf{G}}^\ell$ quantifies pseudo-gradient / gradient alignment.

3 DYNAMICAL MEAN FIELD THEORY FOR VARIOUS LEARNING RULES

For each of these learning rules considered, the infinite width $N \rightarrow \infty$ limit of network learning can be described by a dynamical mean field theory (DMFT) (Bordelon & Pehlevan, 2022). At infinite width, the dynamics of the kernels Φ^ℓ and $\tilde{\mathbf{G}}^\ell$ become deterministic over random Gaussian initialization of parameters $\boldsymbol{\theta}$. The activity of neurons in each layer become i.i.d. random variables drawn from a distribution defined by these kernels, which themselves are averages over these single-site distributions. Below, we provide DMFT formulas which are valid for all of our learning rules

$$\begin{aligned} h_\mu^\ell(t) &= u_\mu^\ell(t) + \gamma_0 \int_0^t ds \sum_{\nu=1}^P [A_{\mu\nu}^{\ell-1}(t, s) g_\nu^\ell(s) + C_{\mu\nu}^{\ell-1}(t, s) \tilde{g}_\nu^\ell(s) + \Phi_{\mu\nu}^{\ell-1}(t, s) \Delta_\nu(s) \tilde{g}_\nu^\ell(s)] \\ z_\mu^\ell(t) &= r_\mu^\ell(t) + \gamma_0 \int_0^t ds \sum_{\nu=1}^P [B_{\mu\nu}^\ell(t, s) + \tilde{G}_{\mu\nu}^{\ell+1}(t, s) \Delta_\nu(s)] \phi(h_\nu^\ell(s)), \quad g_\mu^\ell(t) = \dot{\phi}(h_\mu^\ell(t)) z_\mu^\ell(t) \\ \{u_\mu^\ell(t)\} &\sim \mathcal{GP}(0, \Phi^{\ell-1}), \quad \Phi_{\mu\nu}^\ell(t, s) = \langle \phi(h_\mu^\ell(t)) \phi(h_\nu^\ell(s)) \rangle, \quad A_{\mu\nu}^\ell(t, s) = \gamma_0^{-1} \left\langle \frac{\delta}{\delta r_\nu^\ell(s)} \phi(h_\mu^\ell(t)) \right\rangle \\ \{r_\mu^\ell(t)\} &\sim \mathcal{GP}(0, \mathbf{G}^{\ell+1}), \quad \tilde{G}_{\mu\nu}^\ell(t, s) = \langle g_\mu^\ell(t) \tilde{g}_\nu^\ell(s) \rangle, \quad B_{\mu\nu}^\ell(t, s) = \gamma_0^{-1} \left\langle \frac{\delta}{\delta u_\nu^{\ell+1}(s)} g_\mu^{\ell+1}(t) \right\rangle \end{aligned} \quad (5)$$

The definitions of $\tilde{g}_\mu^\ell(t)$ depend on the learning rule and are described in Table 1. The $z_\mu^\ell(t)$ is the *pre-gradient field* defined so that $g_\mu^\ell(t) = \dot{\phi}(h_\mu^\ell(t)) z_\mu^\ell(t)$. The dependence of these DMFT equations on data comes from the base case $\Phi_{\mu\nu}^0(t, s) = \frac{1}{D} \mathbf{x}_\mu \cdot \mathbf{x}_\nu$ and error signal $\Delta_\mu = -\frac{\partial \mathcal{L}}{\partial f_\mu}$.

Rule	GD	ρ -FA	DFA	GLN	Hebb
$\tilde{g}_\mu^\ell(t)$	$\dot{\phi}(h_\mu^\ell(t)) z_\mu^\ell(t)$	$\dot{\phi}(h_\mu^\ell(t)) \tilde{z}_\mu^\ell(t)$	$\dot{\phi}(h_\mu^\ell(t)) \tilde{z}^\ell$	$\dot{\phi}(m_\mu^\ell) z_\mu^\ell(t)$	$\Delta_\mu(t) \phi(h_\mu^\ell(t))$

Table 1: The field definitions for each learning rule. For ρ -FA, the field has definition $\tilde{z}_\mu^\ell(t) = \rho v_\mu^\ell(t) + \sqrt{1 - \rho^2} \tilde{\zeta}_\mu^\ell(t) + \gamma_0 \int_0^t ds \sum_\nu D_{\mu\nu}^\ell(t, s) \phi(h_\nu^\ell(s))$ where $\{v_\mu^\ell(t), \tilde{\zeta}_\mu^\ell(t)\}$ are Gaussian with $\langle r_\mu^\ell(t) v_\nu^\ell(s) \rangle = \tilde{G}_{\mu\nu}^{\ell+1}(t, s)$. The $\tilde{\zeta}^\ell$ field is an independent Gaussian with correlation $\langle \tilde{\zeta}_\mu^\ell(t) \tilde{\zeta}_\nu^\ell(s) \rangle = \langle \tilde{g}_\mu^{\ell+1}(t) \tilde{g}_\nu^{\ell+1}(s) \rangle = \tilde{G}_{\mu\nu}^{\ell+1}(t, s)$. For DFA, the \tilde{z}^ℓ field is static $\tilde{z}^\ell \sim \mathcal{N}(0, 1)$. For GLN, we use $\{m_\mu^\ell\} \sim \mathcal{N}(0, \mathbf{K}^x)$ as a gating variable. $C^\ell = 0$ except for ρ -FA with $\rho > 0$.

We see that, for $\{\text{GD}, \rho\text{-FA}, \text{DFA}, \text{Hebb}\}$ the distribution of $h_\mu^\ell(t), z_\mu^\ell(t)$ are Gaussian throughout training only in the lazy $\gamma_0 \rightarrow 0$ limit for general nonlinear activation functions $\phi(h)$. However, conditional on $\{m_\mu^\ell\}$, the $\{h^\ell, z^\ell\}$ fields are all Gaussian for GLNs. For all algorithms except $\rho\text{-FA}$, $C^\ell = 0$. For $\rho\text{-FA}$ we have $C_{\mu\alpha}^\ell(t, s) = \gamma_0^{-1} \left\langle \frac{\delta}{\delta v_\alpha^\ell(s)} \phi(h_\mu^\ell(t)) \right\rangle$.

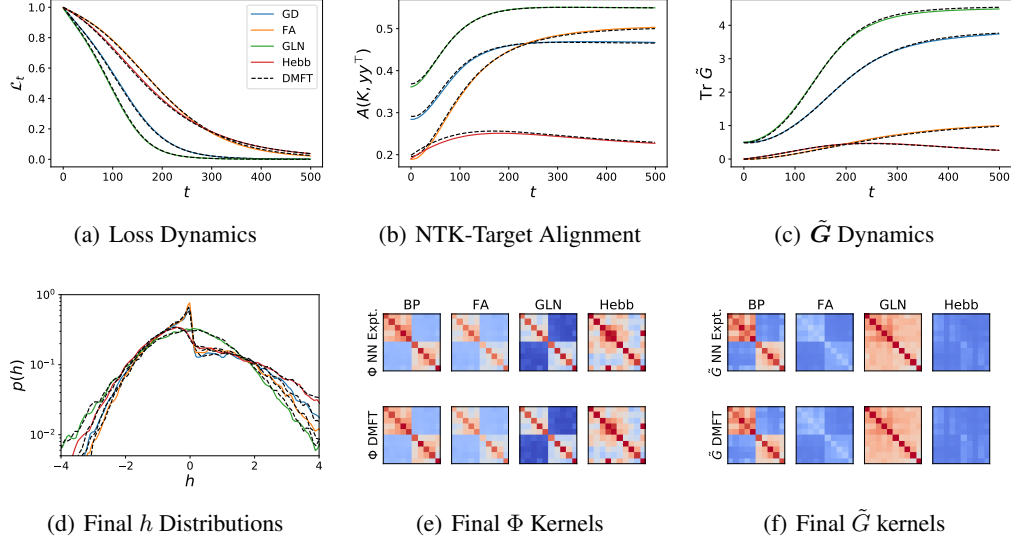


Figure 1: The DMFT predicts feature dynamics of wide networks trained with gradient descent (GD), feedback alignment (FA) with $\rho = 0$, gated linear network (GLN), and a error-modulated $\beta = 1$ Hebb rule (Hebb) in the feature learning regime. (a) The loss dynamics in a two layer ($L = 1, N = 2000$) network trained with these learning rules at richness $\gamma_0 = 2$. The network is trained on a collection of $P = 10$ random vectors in $D = 50$ dimensions. (b) The cosine similarity of the eNTK with the targets $A(\mathbf{K}, \mathbf{y}\mathbf{y}^\top) = \frac{\mathbf{y}^\top \mathbf{K} \mathbf{y}}{|\mathbf{K}|_F |\mathbf{y}|^2}$ reveals increasing alignment for all algorithms. Though FA starts with the lowest alignment, its final NTK task alignment exceeds that of GD. (c) The dynamics of the gradient-pseudogradient kernel $\tilde{\mathbf{G}}$ also reveals increasing correlation of g with \tilde{g} . FA starts with $\tilde{\mathbf{G}} = 0$ but $\tilde{\mathbf{G}}$ increases to non-zero value. (d) The distribution of hidden layer preactivations after training reveals non-Gaussian statistics for both GD and FA, but approximately Gaussian statistics for GLN. (e-f) The final Φ and $\tilde{\mathbf{G}}$ kernels from theory and experiment.

As described in prior results on the GD case (Bordelon & Pehlevan, 2022), the above equations can be solved self-consistently in polynomial (in train-set size P and training steps T) time. With an estimate of the dynamical kernels $\{\Phi_{\mu\nu}^\ell(t, s), \tilde{G}_{\mu\nu}^\ell(t, s), G_{\mu\nu}^\ell(t, s)\}$, one computes the eNTK $K_{\mu\nu}^\ell(t)$ and error dynamics $\Delta_\mu(t)$. From these objects, we can sample the stochastic processes $\{h^\ell, z^\ell, \tilde{z}^\ell\}$ which can then be used to derive new refined estimates of the kernels. This procedure is repeated until convergence. This algorithm can be found in App. A. An example of such a solution is provided in Figure 1 for two layer ReLU networks trained with GD, FA, GLN, and Hebb. We show that our self-consistent DMFT accurately predicts training and kernel dynamics, as well as the density of preactivations $\{h_\mu(t)\}$ and final kernels $\{\Phi_{\mu\nu}, \tilde{G}_{\mu\nu}\}$ for each learning rule. We observe substantial differences in the learned representations (Figure 1e), all predicted by our DMFT.

3.1 LAZY OR EARLY TIME STATIC-KERNEL LIMITS

When $\gamma_0 \rightarrow 0$, we see that the fields $h_\mu^\ell(t)$ and $z_\mu^\ell(t)$ are equal to the Gaussian variables $u_\mu^\ell(0)$ and $r_\mu^\ell(0)$. In this limit, the eNTK $K_{\mu\nu}$ remains static and has the form summarized in Table 2 in terms of the initial feature kernels Φ^ℓ and gradient kernels G^ℓ . We derive these kernels in Appendix D.

The feature $P \times P$ matrices $\Phi^\ell, \mathbf{G}^\ell$ in Table 2 are computed recursively as

$$\Phi^\ell = \langle \phi(\mathbf{u})\phi(\mathbf{u})^\top \rangle_{\mathbf{u} \sim \mathcal{N}(0, \Phi^{\ell-1})}, \quad \mathbf{G}^\ell = \mathbf{G}^{\ell+1} \odot \langle \dot{\phi}(\mathbf{u})\dot{\phi}(\mathbf{u})^\top \rangle_{\mathbf{u} \sim \mathcal{N}(0, \Phi^{\ell-1})} \quad (6)$$

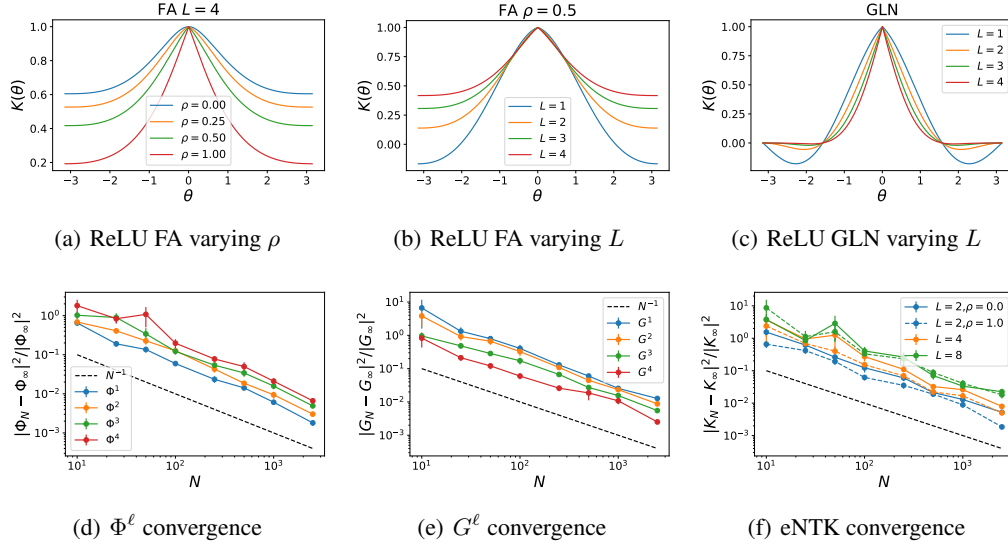


Figure 2: The lazy infinite width limits of the various learning rules can be fully summarized with their initial eNTK. (a) The kernels of ρ -aligned ReLU FA and ReLU GLN for inputs separated by angle θ . (a) The kernels for varying ρ in ρ -aligned FA. Larger ρ has a sharper peak in the kernel around $\theta = 0$. The $\rho \rightarrow 0$ limit recovers the NNGP kernel Φ^L while the $\rho \rightarrow 1$ limit gives the backprop NTK. (b) Deeper networks with partial alignment $\rho = 0.5$. (c) ReLU-GLN kernel sharpens with depth. (d)-(e) The relative error of the infinite width Φ^ℓ , G^ℓ kernels in a width N ReLU neural network. The late layer Φ^ℓ and early layer G^ℓ kernels have highest errors since finite size effects accumulate on forward and backward passes respectively. (f) Finite width corrections to eNTK are larger for small ρ and large depth L . All square errors go as $|K_N - K_\infty|^2 \sim O_N(1/N)$.

Rule	GD	ρ -FA	DFA	GLN	Hebb
$K_{\mu\nu}$	$\sum_{\ell=0}^L G_{\mu\nu}^{\ell+1} \Phi_{\mu\nu}^\ell$	$\sum_{\ell=0}^L \rho^{L-\ell} G_{\mu\nu}^{\ell+1} \Phi_{\mu\nu}^\ell$	$\Phi_{\mu\nu}^L$	$\left[\langle \dot{\phi}(m_\mu) \dot{\phi}(m_\nu) \rangle \right]^L K_{\mu\nu}^x$	$\Phi_{\mu\nu}^L$

Table 2: The initial eNTK $K_{\mu\nu}$ for each learning rule. The GD kernel is the usual initial NTK of Jacot et al. (2018). For ρ -aligned FA, each layer ℓ 's contribution to the eNTK is suppressed by a factor $\rho^{L-\ell}$. For DFA and Hebb, only the last layer feature kernel Φ^L contributes to the NTK. For GLN, each layer has an identical contribution.

with base cases $\Phi^0 = \mathbf{K}^x$ and $\mathbf{G}^{L+1} = \mathbf{11}^\top$. We provide interpretations of this result below.

- Backpropagation (GD) and $\rho = 1$ FA recover the usual depth L NTK, with contributions from every layer $K_{\mu\nu} = \sum_{\ell} G_{\mu\nu}^{\ell+1} \Phi_{\mu\nu}^\ell$ at initialization. This kernel governs both training dynamics and test predictions in the lazy limit $\gamma_0 \rightarrow 0$ (Jacot et al., 2018; Lou et al., 2022).
- $\rho = 0$ FA, DFA and Hebb are equivalent to using the NNGP kernel $K_{\mu\nu} \sim \Phi_{\mu\nu}^L$, giving the Bayes posterior mean (Matthews et al., 2018; Lee et al., 2018; Hron et al., 2020). In the $\gamma_0, \rho \rightarrow 0$ limit, only the dynamics of the readout weights \mathbf{w}^L contribute to the evolution of f_μ since error signals cannot successfully propagate backward and gradients cannot align with pseudo-gradients (App D). The standard $\rho = 0$ FA will be indistinguishable from merely training \mathbf{w}^L with the delta-rule unless the network is trained in the rich feature learning regime $\gamma_0 > 0$, where \tilde{G}^ℓ can evolve. This effect was also noted in two layer networks by Song et al. (2021).
- ρ -FA weighs each layer ℓ with scale $\rho^{L-\ell}$, since each layer's pseudo-gradient is only partially correlated with the true gradient, giving recursion $\tilde{G}^\ell = \rho \tilde{G}^{\ell+1}$ with base case $\tilde{G}^{L+1} = G^{L+1}$.
- GLN's kernel in lazy limit is determined by the Gaussian gating variables $\{m_\mu^\ell\} \sim \mathcal{N}(0, \mathbf{K}^x)$.

We visualize these kernels for deep ReLU networks and ReLU GLNs for normalized inputs $|\mathbf{x}|^2 = |\mathbf{x}'|^2 = D$, by plotting the kernel as a function of the angle θ separating two inputs $\cos(\theta) =$

$\frac{1}{D} \mathbf{x}^\top \mathbf{x}'$. We find that the kernels develop a sharp discontinuity at the origin $\theta = 0$, which becomes more exaggerated as ρ and L increase. We further show that the square difference of width N kernels and infinite width kernels go as $O(N^{-1})$. We derive this scaling with a perturbative argument in App. H, which enables analytical prediction of leading order finite size effects (Figure 7). In the lazy $\gamma_0 \rightarrow 0$ limit, these kernels define the eNTK and the network prediction dynamics.

3.2 FEATURE LEARNING ENABLES GRADIENT/PSEUDO-GRADIENT ALIGNMENT AND KERNEL/TASK ALIGNMENT

In the last section, we saw that, in the $\gamma_0 \rightarrow 0$ limit, all algorithms have frozen preactivations and pregradient features $\{h_\mu^\ell(t), z_\mu^\ell(t)\}$. A consequence of this fact is that FA and DFA cannot increase their gradient-pseudogradient alignment throughout training in the lazy limit $\gamma_0 = 0$. However, if we increase γ_0 , then the gradient features $g_\mu^\ell(t)$ and pseudo-gradients $\tilde{g}_\mu^\ell(t)$ evolve in time and can increase their alignment. In Figure 3, we show the effect of increasing γ_0 on alignment dynamics in a depth 4 tanh network trained with DFA. In (b), we see that larger γ_0 is associated with high task-alignment of the last layer feature kernel Φ^L , which becomes essentially rank one and aligned to $\mathbf{y}\mathbf{y}^\top$. The asymptotic cosine similarity between gradients and pseudogradients also increase with γ_0 . The eNTK also becomes aligned with the task relevant directions (shown in Figure 3 c), like has been observed in GD training (Baratin et al., 2021; Shan & Bordelon, 2021; Geiger et al., 2021; Atanasov et al., 2022). We see that width N networks have a dynamical eNTK $\mathbf{K}_N(t)$ which deviates from the DMFT eNTK $\mathbf{K}_\infty(t)$ by $O(1/N)$ in square loss. DMFT is more predictive for larger γ_0 networks, suggesting a reduction in finite size variability due to task-relevant feature evolution.

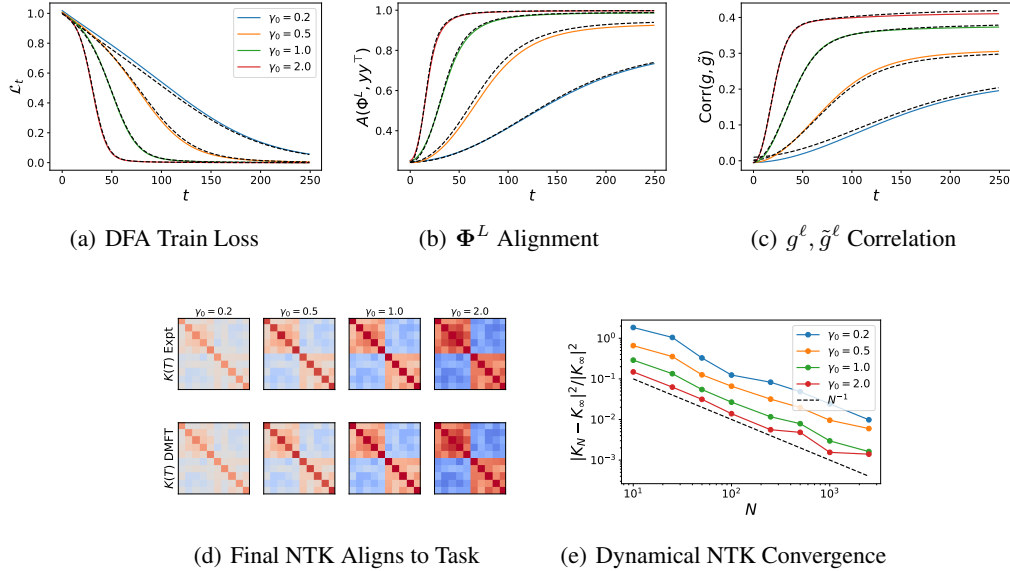


Figure 3: Feature Learning enables alignment for a depth 4 ($L = 3$ hidden layers) tanh network trained with direct feedback alignment (DFA) with varying γ_0 . (a) Training loss for DFA networks with width $N = 4000$ with varying richness γ_0 shows that feature learning accelerates training, as predicted by DMFT (black). (b) The alignment (cosine similarity) of the last layer kernel Φ^L with the target function reveals successful task dependent feature learning at large γ_0 . (c) The dynamics of pseudo-grad./grad. correlation $\text{corr}(\mathbf{g}, \tilde{\mathbf{g}}) = \frac{1}{LP} \sum_{\ell, \mu} \frac{\mathbf{g}_\mu^\ell(t) \cdot \tilde{\mathbf{g}}_\mu^\ell(t)}{|\mathbf{g}_\mu^\ell(t)| |\tilde{\mathbf{g}}_\mu^\ell(t)|}$ averaged over layers ℓ and datapoints μ . Larger γ_0 generates more significant alignment between pseudogradients and gradients. (d) The final NTKs as a function of γ_0 reveals increasing clustering of the data points by class. (e) The error of the DMFT approximation for K 's dynamics as a function of N : $\frac{\langle |\mathbf{K}_N(t) - \mathbf{K}_\infty(t)|^2 \rangle_t}{\langle |\mathbf{K}_\infty(t)|^2 \rangle_t} \sim O(N^{-1})$, where the averages are computed over the time interval of training. This error is smaller for larger feature learning strength γ_0 .

3.3 DEEP LINEAR NETWORK KERNEL DYNAMICS

When $\gamma_0 > 0$ the kernels and features in the network evolve according to the DMFT equations. For deep linear networks we can analyze the equations for the kernels in closed form without sampling since the correlation functions close algebraically (App. E). In Figure 4, we utilize our algebraic DMFT equations to explore ρ -FA dynamics in a depth 4 linear network. Networks with larger ρ train faster, which can be intuited by noting that the initial function time derivative $\frac{df}{dt}|_{t=0} \sim \sum_{\ell=0}^L \rho^{L-\ell} \sim \frac{1-\rho^{L+1}}{1-\rho}$ is an increasing function of ρ . We observe higher final gradient pseudogradient alignment in each layer with larger ρ , which is also intuitive from the initial condition $\tilde{G}^\ell(0) = \rho^{L-\ell}$. However, surprisingly, for large initial correlation ρ , the NTK achieves lower task alignment, despite having larger $\tilde{G}^\ell(t)$. We show that this is caused by smaller overlap of each layer’s feature kernel $\mathbf{H}^\ell(t)$ with $\mathbf{y}\mathbf{y}^\top$. Though this phenomenon is counterintuitive, we gain more insight in the next section by studying an even simpler two layer model.

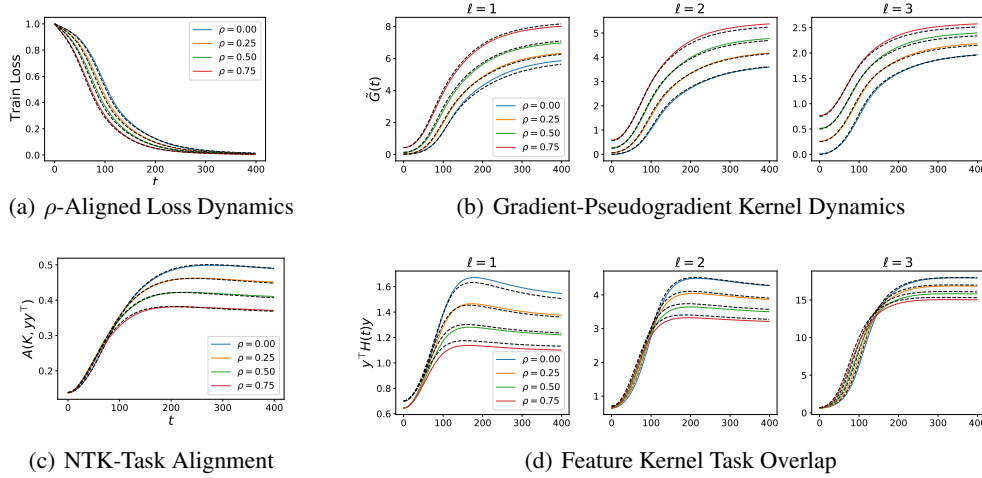


Figure 4: The initial feedback correlation ρ alters alignment dynamics in on the FA dynamics in a depth 4 ($L = 3$ hidden layer) linear network. (a) Larger ρ leads to faster initial training since the scale of the eNTK is larger. (b) Further, larger ρ leads to larger scales of $\tilde{G}(t) = \frac{1}{N} \mathbf{g}^\ell(t) \cdot \tilde{\mathbf{g}}^\ell(t)$. (c) However, smaller ρ leads to more alignment of the NTK $\mathbf{K}(t)$ with the task-relevant subspace, measured with cosine similarity $A(\mathbf{K}, \mathbf{y}\mathbf{y}^\top)$. (d) The feature kernel $\mathbf{H}(t)$ overlaps with \mathbf{y} reveal that $\mathbf{H}^\ell(t)$ aligns more significantly in the small ρ networks.

3.3.1 EXACTLY SOLVEABLE DYNAMICS IN TWO LAYER LINEAR NETWORK

We can provide exact solutions to the infinite width GD and ρ -FA dynamics in the setting of Saxe et al. (2013), specifically a two layer linear network trained with whitened data $K_{\mu\nu}^x = \delta_{\mu\nu}$. Unlike Saxe et al. (2013)’s result, however, we do not demand small initialization scale (or equivalently large γ_0), but rather provide the exact solution for all positive γ_0 . We will establish that large initial correlation ρ results in higher gradient/pseudogradient alignment but lower alignment of the hidden feature kernel $\mathbf{H}(t)$ with the task relevant subspace $\mathbf{y}\mathbf{y}^\top$.

We first note that when $\mathbf{K}^x = \mathbf{I}$, the GD or FA hidden feature kernel $\mathbf{H}(t)$ only evolves in the rank-one $\mathbf{y}\mathbf{y}^\top$ subspace. It thus suffices to track the projection of $\mathbf{H}(t)$ on this rank one subspace, which we call $H_y(t)$. In the App. F we derive dynamics for H_y for GD and ρ -FA

$$H_y(t) = \begin{cases} \tilde{G}(t) = \sqrt{1 + \gamma_0^2 (y - \Delta(t))^2}, & \frac{d\Delta}{dt} = -\sqrt{1 + \gamma_0^2 (y - \Delta(t))^2} \Delta(t) & \text{GD} \\ 2\tilde{G}(t) + 1 - 2\rho = 1 + a^2, & \frac{da}{dt} = \gamma_0 y - \frac{1}{2}a^3 - (1 + \rho)a & \rho\text{-FA} \end{cases} \quad (7)$$

We illustrate these dynamics in Figure 5. The fixed points are $H_y = \sqrt{1 + \gamma_0^2 y^2}$ for GD and for ρ -FA, $H_y = 1 + a^2$ where a is the smallest positive root of $\frac{1}{2}a^3 + (1 + \rho)a = \gamma_0 y$. For both GD and FA, we see that increasing γ_0 results in larger asymptotic values for H_y and \tilde{G} . For ρ -FA the fixed point of a ’s dynamics is a strictly decreasing function of ρ since $\frac{da}{d\rho} < 0$, showing that the final

value of H_y is smaller for larger ρ . On the contrary, we have that the final $\tilde{G} = \rho + \frac{1}{2}a^2$ is a strictly increasing function of ρ since $\frac{d}{d\rho}\tilde{G} = 1 - \frac{a^2}{\frac{3}{2}a^2 + (1+\rho)} > \frac{1}{3} > 0$. Thus, this simple model replicates the phenomenon of increasing \tilde{G} and decreasing H_y as ρ increases. For the Hebb rule with $\mathbf{K}^x = \mathbf{I}$, the story is different. Instead of aligning \mathbf{H} along the rank-one task relevant subspace, the dynamics instead decouple over samples, giving the following P separate equations

$$\frac{d}{dt}\Delta_\mu = -[H_{\mu\mu}(t) + \gamma_0\Delta_\mu(y_\mu - \Delta_\mu)]\Delta_\mu(t), \quad \frac{d}{dt}H_{\mu\mu} = 2\gamma_0\Delta_\mu(t)^2H_{\mu\mu}. \quad (8)$$

From this perspective, we see that the hidden feature kernel does not align to the task, but rather increases its entries in overall scale as is illustrated in Figure 5 (b).




Figure 5: The feature kernel dynamics and scaling with γ_0^2 for GD, ρ -FA, and Hebbian rules in an exactly solvable two layer linear network. (a) The loss dynamics for all algorithms reveals that $\rho = 0$ FA and Hebb rules have same early time dynamics and that $\rho = 1$ FA and GD have same early-time dynamics. However all loss curves become distinct at late times due to different eNTK dynamics. (b) The alignment of the kernel to the target function $H_y(t) = \frac{1}{|\mathbf{y}|^2} \mathbf{y}^\top \mathbf{H} \mathbf{y} / \text{Tr} \mathbf{H}(t)$ increases significantly for GD, and FA, but not for Hebb, reflecting the task-independence of the learned representation. (c) The movement of the feature kernel $\Delta H_y = \lim_{t \rightarrow \infty} H_y(t) - H_y(0)$ as a function of γ_0 for GD, and $\rho = 0, 1$ FA. At small feature learning strength, all algorithms have updates on the order of $\Delta H_y \sim \gamma_0^2$. At large γ_0 , GD has $\Delta H_y \sim \gamma_0$ while FA has $\Delta H_y \sim \gamma_0^{2/3}$. The $\rho = 1$ FA (green) has lower ΔH_y than the $\rho = 0$ FA across all γ_0 .

4 DISCUSSION

We provided an analysis of the training dynamics of a wide range of learning rules at infinite width. This set of rules includes (but is not limited to) GD, ρ -FA, DFA, GLN and Hebb as well as many others. We showed that each of these learning rules has an dynamical effective NTK which concentrates over initializations at infinite width. In the lazy $\gamma_0 \rightarrow 0$ regime, it suffices to compute the the initial NTK, while in the rich regime, we provide a dynamical mean field theory to compute the NTK’s dynamics. We showed that, in the rich regime, FA learning rules do indeed align the network’s gradient vectors to their pseudo-gradients and that this alignment improves with γ_0 . We show that initial correlation ρ between forward and backward pass weights alters the inductive bias of FA in both lazy and rich regimes. In the rich regime, larger ρ networks have smaller eNTK evolution. Overall, our study is a step towards understanding learned representations in neural networks, and the quest to reverse-engineer learning rules from observations of evolving neural representations during learning in the brain.

Many open problems remain unresolved with the present work. We currently have only implemented our theory in MLPs. An implementation in CNNs could explain some of the observed advantages of partial initial alignment in ρ -FA (Xiao et al., 2018; Moskovitz et al., 2018; Bartunov et al., 2018; Refinetti et al., 2021). In addition, our framework is sufficiently flexible to propose and test new learning rules by providing new $\tilde{\mathbf{g}}_\mu^\ell(t)$ formulas. Our DMFT gives a recipe to compute their initial kernels, function dynamics and analyze their learned representations. The generalization performance of these learning rules at varying γ_0 is yet to be explored. Lastly, our DMFT is numerically expensive for large datasets and training intervals, making it difficult to scale up to realistic datasets. Future work could provide theoretical convergence guarantees for our DMFT solver.

REFERENCES

- Elisabeth Agoritsas, Giulio Biroli, Pierfrancesco Urbani, and Francesco Zamponi. Out-of-equilibrium dynamical mean-field equations for the perceptron model. *Journal of Physics A: Mathematical and Theoretical*, 51(8):085002, 2018.
- Alexander Atanasov, Blake Bordelon, and Cengiz Pehlevan. Neural networks as kernel learners: The silent alignment effect. In *International Conference on Learning Representations, 2022*. URL <https://openreview.net/forum?id=1NvflqAdoom>.
- Aristide Baratin, Thomas George, César Laurent, R Devon Hjelm, Guillaume Lajoie, Pascal Vincent, and Simon Lacoste-Julien. Implicit regularization via neural feature alignment. In Arindam Banerjee and Kenji Fukumizu (eds.), *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, volume 130 of *Proceedings of Machine Learning Research*, pp. 2269–2277. PMLR, 13–15 Apr 2021. URL <https://proceedings.mlr.press/v130/baratin21a.html>.
- Sergey Bartunov, Adam Santoro, Blake Richards, Luke Marris, Geoffrey E Hinton, and Timothy Lillicrap. Assessing the scalability of biologically-motivated deep learning algorithms and architectures. *Advances in neural information processing systems*, 31, 2018.
- Carl M Bender, Steven Orszag, and Steven A Orszag. *Advanced mathematical methods for scientists and engineers I: Asymptotic methods and perturbation theory*, volume 1. Springer Science & Business Media, 1999.
- Akhilan Boopathy and Ila Fiete. How to train your wide neural network without backprop: An input-weight alignment perspective. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (eds.), *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 2178–2205. PMLR, 17–23 Jul 2022. URL <https://proceedings.mlr.press/v162/boopathy22a.html>.
- Blake Bordelon and Cengiz Pehlevan. Self-consistent dynamical field theory of kernel evolution in wide neural networks. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (eds.), *Advances in Neural Information Processing Systems, 2022*. URL <https://openreview.net/forum?id=sipwrPCrIS>.
- David Budden, Adam Marblestone, Eren Sezener, Tor Lattimore, Gregory Wayne, and Joel Veness. Gaussian gated linear networks. *Advances in Neural Information Processing Systems*, 33:16508–16519, 2020.
- Yinan Cao, Christopher Summerfield, and Andrew Saxe. Characterizing emergent representations in a space of candidate learning rules for deep networks. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 8660–8670. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/file/6275d7071d005260ab9d0766d6df1145-Paper.pdf>.
- Michael Celentano, Chen Cheng, and Andrea Montanari. The high-dimensional asymptotics of first order methods with random data. *arXiv preprint arXiv:2112.07572*, 2021.
- Lenaic Chizat, Edouard Oyallon, and Francis Bach. On lazy training in differentiable programming. *Advances in Neural Information Processing Systems*, 32, 2019.
- Francis Crick. The recent excitement about neural networks. *Nature*, 337(6203):129–132, 1989.
- A Crisanti and H Sompolinsky. Path integral approach to random neural networks. *Physical Review E*, 98(6):062120, 2018.
- Stéphane d’Ascoli, Levent Sagun, Giulio Biroli, and Joan Bruna. Finding the needle in the haystack with convolutions: on the benefits of architectural bias. *Advances in Neural Information Processing Systems*, 32, 2019.
- Jonathan Fiat, Eran Malach, and Shai Shalev-Shwartz. Decoupling gating from linearity. *arXiv preprint arXiv:1906.05032*, 2019.

- Charlotte Frenkel, Martin Lefebvre, and David Bol. Learning without feedback: direct random target projection as a feedback-alignment algorithm with layerwise feedforward training. *arXiv preprint arXiv:1909.01311*, 10, 2019.
- Shivam Garg and Santosh Vempala. How and when random feedback works: A case study of low-rank matrix factorization. In Gustau Camps-Valls, Francisco J. R. Ruiz, and Isabel Valera (eds.), *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, volume 151 of *Proceedings of Machine Learning Research*, pp. 4070–4108. PMLR, 28–30 Mar 2022. URL <https://proceedings.mlr.press/v151/garg22a.html>.
- Mario Geiger, Leonardo Petrini, and Matthieu Wyart. Landscape and training regimes in deep learning. *Physics Reports*, 924:1–18, 2021.
- Cedric Gerbelot, Emanuele Troiani, Francesca Mignacco, Florent Krzakala, and Lenka Zdeborova. Rigorous dynamical mean field theory for stochastic gradient descent methods, 2022. URL <https://arxiv.org/abs/2210.06591>.
- Gabriel Goh. Why momentum really works. *Distill*, 2(4):e6, 2017.
- Ian Goodfellow, Yoshua Bengio, Aaron Courville, and Yoshua Bengio. *Deep learning*, volume 1. MIT Press, 2016.
- Donald O. Hebb. *The organization of behavior: A neuropsychological theory*. Wiley, New York, June 1949. ISBN 0-8058-4300-0.
- Jiri Hron, Yasaman Bahri, Roman Novak, Jeffrey Pennington, and Jascha Sohl-Dickstein. Exact posterior distributions of wide bayesian neural networks. *arXiv preprint arXiv:2006.10541*, 2020.
- Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. *Advances in neural information processing systems*, 31, 2018.
- Mehran Kardar. *Statistical physics of fields*. Cambridge University Press, 2007.
- Nikolaus Kriegeskorte and Xue-Xin Wei. Neural tuning and representational geometry. *Nature Reviews Neuroscience*, 22(11):703–718, 2021.
- Kamesh Krishnamurthy, Tankut Can, and David J Schwab. Theory of gating in recurrent neural networks. *Physical Review X*, 12(1):011011, 2022.
- Julien Launay, Iacopo Poli, François Boniface, and Florent Krzakala. Direct feedback alignment scales to modern deep learning tasks and architectures. *Advances in neural information processing systems*, 33:9346–9360, 2020.
- Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553):436–444, 2015.
- Jaehoon Lee, Jascha Sohl-dickstein, Jeffrey Pennington, Roman Novak, Sam Schoenholz, and Yasaman Bahri. Deep neural networks as gaussian processes. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=B1EA-M-0Z>.
- Jaehoon Lee, Lechao Xiao, Samuel Schoenholz, Yasaman Bahri, Roman Novak, Jascha Sohl-Dickstein, and Jeffrey Pennington. Wide neural networks of any depth evolve as linear models under gradient descent. *Advances in neural information processing systems*, 32, 2019.
- Jaehoon Lee, Samuel Schoenholz, Jeffrey Pennington, Ben Adlam, Lechao Xiao, Roman Novak, and Jascha Sohl-Dickstein. Finite versus infinite neural networks: an empirical study. *Advances in Neural Information Processing Systems*, 33:15156–15172, 2020.
- Qianli Liao, Joel Leibo, and Tomaso Poggio. How important is weight symmetry in backpropagation? In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30, 2016.
- Timothy P Lillicrap, Daniel Cownden, Douglas B Tweed, and Colin J Akerman. Random synaptic feedback weights support error backpropagation for deep learning. *Nature communications*, 7(1): 1–10, 2016.

- Yizhang Lou, Chris E Mingard, and Soufiane Hayou. Feature learning and signal propagation in deep neural networks. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (eds.), *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 14248–14282. PMLR, 17–23 Jul 2022. URL <https://proceedings.mlr.press/v162/lou22a.html>.
- Alessandro Manacorda, Grégory Schehr, and Francesco Zamponi. Numerical solution of the dynamical mean field theory of infinite-dimensional equilibrium liquids. *The Journal of chemical physics*, 152(16):164506, 2020.
- Paul Cecil Martin, ED Siggia, and HA Rose. Statistical dynamics of classical systems. *Physical Review A*, 8(1):423, 1973.
- Alexander G.D.G. Matthews, Jiri Hron, Mark Rowland, Richard E. Turner, and Zoubin Ghahramani. Gaussian process behaviour in wide deep neural networks. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=H1-nGgWC->.
- Song Mei, Andrea Montanari, and Phan-Minh Nguyen. A mean field view of the landscape of two-layer neural networks. *Proceedings of the National Academy of Sciences*, 115(33):E7665–E7671, 2018.
- Francesca Mignacco, Florent Krzakala, Pierfrancesco Urbani, and Lenka Zdeborová. Dynamical mean-field theory for stochastic gradient descent in gaussian mixture classification. *Advances in Neural Information Processing Systems*, 33:9540–9550, 2020.
- Theodore H Moskovitz, Ashok Litwin-Kumar, and LF Abbott. Feedback alignment in deep convolutional networks. *arXiv preprint arXiv:1812.06488*, 2018.
- Arild Nøkland. Direct feedback alignment provides learning in deep neural networks. *Advances in neural information processing systems*, 29, 2016.
- Jasper Poort, Adil G Khan, Marius Pachitariu, Abdellatif Nemri, Ivana Orsolich, Julija Krupic, Marius Bauza, Maneesh Sahani, Georg B Keller, Thomas D Mrsic-Flogel, et al. Learning enhances sensory and multiple non-sensory representations in primary visual cortex. *Neuron*, 86(6):1478–1490, 2015.
- Maria Refinetti, Stéphane d’Ascoli, Ruben Ohana, and Sebastian Goldt. Align, then memorise: the dynamics of learning with feedback alignment. In *International Conference on Machine Learning*, pp. 8925–8935. PMLR, 2021.
- Andrew Saxe, Shagun Sodhani, and Sam Jay Lewallen. The neural race reduction: Dynamics of abstraction in gated networks. In *International Conference on Machine Learning*, pp. 19287–19309. PMLR, 2022.
- Andrew M Saxe, James L McClelland, and Surya Ganguli. Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. *arXiv preprint arXiv:1312.6120*, 2013.
- Joseph W Schumacher, Matthew K McCann, Katherine J Maximov, and David Fitzpatrick. Selective enhancement of neural coding in v1 underlies fine-discrimination learning in tree shrew. *Current Biology*, 32(15):3245–3260, 2022.
- Eren Sezener, Agnieszka Grabska-Barwińska, Dimitar Kostadinov, Maxime Beau, Sanjukta Krishnagopal, David Budden, Marcus Hutter, Joel Veness, Matthew Botvinick, Claudia Clopath, et al. A rapid and efficient learning rule for biological neural circuits. *BioRxiv*, 2021.
- Haozhe Shan and Blake Bordelon. A theory of neural tangent kernel alignment and its influence on training. *arXiv e-prints*, pp. arXiv–2105, 2021.
- Ganlin Song, Ruitu Xu, and John Lafferty. Convergence and alignment of gradient descent with random backpropagation weights. *Advances in Neural Information Processing Systems*, 34:19888–19898, 2021.

Joel Veness, Tor Lattimore, David Budden, Avishkar Bhoopchand, Christopher Mattern, Agnieszka Grabska-Barwinska, Eren Sezener, Jianan Wang, Peter Toth, Simon Schmitt, et al. Gated linear networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 10015–10023, 2021.

James CR Whittington and Rafal Bogacz. Theories of error back-propagation in the brain. *Trends in cognitive sciences*, 23(3):235–250, 2019.

Bernard Widrow and Marcian E Hoff. Adaptive switching circuits. Technical report, Stanford Univ Ca Stanford Electronics Labs, 1960.

Will Xiao, Honglin Chen, Qianli Liao, and Tomaso Poggio. Biologically-plausible learning algorithms can scale to large datasets. *arXiv preprint arXiv:1811.03567*, 2018.

Greg Yang and Edward J. Hu. Tensor programs iv: Feature learning in infinite-width neural networks. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 11727–11737. PMLR, 18–24 Jul 2021. URL <https://proceedings.mlr.press/v139/yang21c.html>.

APPENDIX

A ALGORITHM TO SOLVE NONLINEAR DMFT EQUATIONS

Algorithm 1: Alternating Monte-Carlo Solution to Saddle Point Equations

Data: \mathbf{K}^x, \mathbf{y} , Initial Guesses $\{\Phi^\ell, \mathbf{G}^\ell, \tilde{\mathbf{G}}^\ell, \tilde{\tilde{\mathbf{G}}}^\ell\}_{\ell=1}^L, \{\mathbf{A}^\ell, \mathbf{B}^\ell, \mathbf{C}^\ell, \mathbf{D}^\ell\}_{\ell=1}^{L-1}$, Sample count \mathcal{S} , Update Speed β

Result: Network predictions through training $f_\mu(t)$, correlation functions $\{\Phi^\ell, \mathbf{G}^\ell, \tilde{\mathbf{G}}^\ell, \tilde{\tilde{\mathbf{G}}}^\ell\}_{\ell=1}^L$, response functions $\{\mathbf{A}^\ell, \mathbf{B}^\ell, \mathbf{C}^\ell, \mathbf{D}^\ell\}_{\ell=1}^{L-1}$,

```

1  $\Phi^0 = \mathbf{K}^x \otimes \mathbf{1}\mathbf{1}^\top, \mathbf{G}^{L+1} = \mathbf{1}\mathbf{1}^\top;$ 
2 while Kernels Not Converged do
3   From  $\{\Phi^\ell, \mathbf{G}^\ell\}$  compute  $\mathbf{K}^{NTK}(t, t)$  and solve  $\frac{d}{dt}f_\mu(t) = \sum_\alpha \Delta_\alpha(t)K_{\mu\alpha}^{NTK}(t, t);$ 
4    $\ell = 1;$ 
5   while  $\ell < L + 1$  do
6     Draw  $\mathcal{S}$  samples  $\{u_{\mu,n}^\ell(t)\}_{n=1}^{\mathcal{S}} \sim \mathcal{GP}(0, \Phi^{\ell-1}),$ 
7      $\{r_{\mu,n}^\ell(t), v_{\mu,n}^\ell(t)\}_{n=1}^{\mathcal{S}} \sim \mathcal{GP}\left(0, \begin{bmatrix} \mathbf{G}^{\ell+1} & \tilde{\mathbf{G}}^{\ell+1} \\ \tilde{\tilde{\mathbf{G}}}^{\ell+1\top} & \tilde{\tilde{\mathbf{G}}}^{\ell+1} \end{bmatrix}\right);$ 
8     Solve equation 5 for each sample to get  $\{h_{\mu,n}^\ell(t), z_{\mu,n}^\ell(t), \tilde{g}_{\mu,n}^\ell(t)\}_{n=1}^{\mathcal{S}};$ 
9     Use learning rule (Table 1) to compute  $\{\tilde{g}_{\mu,n}^\ell(t)\}_{n=1}^{\mathcal{S}};$ 
10    Compute new correlation function  $\{\Phi^\ell, \mathbf{G}^\ell, \tilde{\mathbf{G}}^\ell, \tilde{\tilde{\mathbf{G}}}^\ell\}$  estimates:
11     $\Phi_{\mu\nu}^{\ell,new}(t, s) = \frac{1}{\mathcal{S}} \sum_{n \in [\mathcal{S}]} \phi(h_{\mu,n}^\ell(t))\phi(h_{\nu,n}^\ell(s)),$ 
12     $\mathbf{G}_{\mu\nu}^{\ell,new}(t, s) = \frac{1}{\mathcal{S}} \sum_{n \in [\mathcal{S}]} g_{\mu,n}^\ell(t)g_{\nu,n}^\ell(s),$ 
13     $\tilde{\mathbf{G}}_{\mu\nu}^{\ell,new}(t, s) = \frac{1}{\mathcal{S}} \sum_{n \in [\mathcal{S}]} g_{\mu,n}^\ell(t)\tilde{g}_{\nu,n}^\ell(s),$ 
14     $\tilde{\tilde{\mathbf{G}}}_{\mu\nu}^{\ell,new}(t, s) = \frac{1}{\mathcal{S}} \sum_{n \in [\mathcal{S}]} \tilde{g}_{\mu,n}^\ell(t)\tilde{g}_{\nu,n}^\ell(s);$ 
15    Solve for Jacobians on each sample  $\frac{\partial\phi(\mathbf{h}_n^\ell)}{\partial\mathbf{r}_n^{\ell\top}}, \frac{\partial\phi(\mathbf{h}_n^\ell)}{\partial\mathbf{v}_n^{\ell\top}}, \frac{\partial g_n^\ell}{\partial\mathbf{u}_n^{\ell\top}}, \frac{\partial\tilde{g}_n^\ell}{\partial\mathbf{u}_n^{\ell\top}};$ 
16    Compute new response functions  $\{\mathbf{A}^\ell, \mathbf{B}^{\ell-1}, \mathbf{C}^\ell, \mathbf{D}^{\ell-1}\}$  estimates:
17     $\mathbf{A}^{\ell,new} = \frac{1}{\mathcal{S}} \sum_{n \in [\mathcal{S}]} \frac{\partial\phi(\mathbf{h}_n^\ell)}{\partial\mathbf{r}_n^{\ell\top}}, \mathbf{B}^{\ell-1,new} = \frac{1}{\mathcal{S}} \sum_{n \in [\mathcal{S}]} \frac{\partial g_n^\ell}{\partial\mathbf{u}_n^{\ell\top}};$ 
18     $\mathbf{C}^{\ell,new} = \frac{1}{\mathcal{S}} \sum_{n \in [\mathcal{S}]} \frac{\partial\phi(\mathbf{h}_n^\ell)}{\partial\mathbf{v}_n^{\ell\top}}, \mathbf{D}^{\ell-1,new} = \frac{1}{\mathcal{S}} \sum_{n \in [\mathcal{S}]} \frac{\partial\tilde{g}_n^\ell}{\partial\mathbf{u}_n^{\ell\top}};$ 
19     $\ell \leftarrow \ell + 1;$ 
20  end
21   $\ell = 1;$ 
22  while  $\ell < L + 1$  do
23    Update correlation functions
24     $\Phi^\ell \leftarrow (1 - \beta)\Phi^\ell + \beta\Phi^{\ell,new}, \mathbf{G}^\ell \leftarrow (1 - \beta)\mathbf{G}^\ell + \beta\mathbf{G}^{\ell,new};$ 
25     $\tilde{\mathbf{G}}^\ell \leftarrow (1 - \beta)\tilde{\mathbf{G}}^\ell + \beta\tilde{\mathbf{G}}^{\ell,new}, \tilde{\tilde{\mathbf{G}}}^\ell \leftarrow (1 - \beta)\tilde{\tilde{\mathbf{G}}}^\ell + \beta\tilde{\tilde{\mathbf{G}}}^{\ell,new};$ 
26    if  $\ell < L$  then
27      Update response functions
28       $\mathbf{A}^\ell \leftarrow (1 - \beta)\mathbf{A}^\ell + \beta\mathbf{A}^{\ell,new}, \mathbf{B}^\ell \leftarrow (1 - \beta)\mathbf{B}^\ell + \beta\mathbf{B}^{\ell,new}$ 
29       $\mathbf{C}^\ell \leftarrow (1 - \beta)\mathbf{C}^\ell + \beta\mathbf{C}^{\ell,new}, \mathbf{D}^\ell \leftarrow (1 - \beta)\mathbf{D}^\ell + \beta\mathbf{D}^{\ell,new}$ 
30    end
31     $\ell \leftarrow \ell + 1$ 
32  end
33 return  $\{f_\mu(t)\}_{\mu=1}^P, \{\Phi^\ell, \mathbf{G}^\ell, \tilde{\mathbf{G}}^\ell, \tilde{\tilde{\mathbf{G}}}^\ell\}_{\ell=1}^L, \{\mathbf{A}^\ell, \mathbf{B}^\ell, \mathbf{C}^\ell, \mathbf{D}^\ell\}_{\ell=1}^{L-1}$ 

```

The sample-and-solve procedure we developed and describe below for nonlinear networks is based on numerical recipes used in the dynamical mean field simulations in computational physics Manacorda et al. (2020) and is similar to recent work in the GD case Bordelon & Pehlevan (2022). The basic principle is to leverage the fact that, conditional on order parameters, we can easily draw

samples $\{u_\mu^\ell(t), r_\mu^\ell(t), \zeta_\mu^\ell(t), \tilde{\zeta}_\mu^\ell(t)\}$ from their appropriate GPs. From these sampled fields, we can identify the kernel order parameters by simple estimation of the appropriate moments. The algorithm is provided in Algorithm 1.

The parameter β controls recency weighting of the samples obtained at each iteration. If $\beta = 1$, then the rank of the kernel estimates is limited to the number of samples \mathcal{S} used in a single iteration, but with $\beta < 1$ smaller sample sizes \mathcal{S} can be used to still obtain accurate results. We used $\beta = 0.6$ in our deep network experiments.

B DERIVATION OF DMFT EQUATIONS

In this section, we derive the DMFT description of infinite network dynamics. The path integral theory we develop is based on the Martin-Siggia-Rose-De Dominicis-Janssen (MSRDJ) framework Martin et al. (1973). A useful review of this technique applied to random recurrent networks can be found here Crisanti & Sompolinsky (2018). This framework was recently extended for deep learning with GD in (Bordelon & Pehlevan, 2022).

B.1 WRITING EVOLUTION EQUATIONS IN FEATURE SPACE

First, we will express all of the learning dynamics in terms of preactivation features $\mathbf{h}_\mu^\ell(t) = \frac{1}{\sqrt{N}} \mathbf{W}^\ell(t) \phi(\mathbf{h}_\mu^\ell(t))$, pre-gradient features $\mathbf{z}_\mu^\ell(t) = \frac{1}{\sqrt{N}} \mathbf{W}^\ell(t)^\top \mathbf{g}^{\ell+1}$ and pseudogradient features $\tilde{\mathbf{g}}_\mu^\ell(t)$. Since we would like to understand typical behavior over random initializations of weights $\boldsymbol{\theta}(0) = \{\mathbf{W}^0(0), \mathbf{W}^1(0), \dots, \mathbf{W}^L(0)\}$, we want to isolate the dependence of our evolution equations by $\mathbf{W}^\ell(0)$. We achieve this separation by using our learning dynamics for $\mathbf{W}^\ell(t)$

$$\mathbf{W}^\ell(t) = \mathbf{W}^\ell(0) + \frac{\gamma_0}{\sqrt{N}} \int_0^t ds \sum_{\mu=1}^P \Delta_\mu(s) \tilde{\mathbf{g}}_\mu^{\ell+1}(s) \phi(\mathbf{h}_\mu^\ell(s))^\top. \quad (9)$$

The inclusion of the prefactor $\frac{\gamma_0}{\sqrt{N}}$ in the weight dynamics ensures that $\frac{d}{dt} f = O_{\gamma_0, N}(1)$ and $\frac{d}{dt} h^\ell = O_{\gamma_0, N}(\gamma_0)$ at initialization (Chizat et al., 2019; Bordelon & Pehlevan, 2022). Using the forward and backward pass equations, we find the following evolution equations for our feature vectors

$$\begin{aligned} \mathbf{h}_\mu^\ell(t) &= \boldsymbol{\chi}_\mu^\ell(t) + \gamma_0 \int_0^t ds \sum_{\nu=1}^P \Delta_\nu(s) \tilde{\mathbf{g}}_\mu^{\ell+1}(s) \Phi_{\mu\nu}^{\ell-1}(t, s), \quad \boldsymbol{\chi}_\mu^\ell(t) = \frac{1}{\sqrt{N}} \mathbf{W}^\ell(0) \phi(\mathbf{h}_\mu^\ell(t)) \\ \mathbf{z}_\mu^\ell(t) &= \boldsymbol{\xi}_\mu^\ell(t) + \gamma_0 \int_0^t ds \sum_{\nu=1}^P \Delta_\nu(s) \phi(\mathbf{h}_\mu^\ell(s)) \tilde{G}_{\mu\nu}^{\ell+1}(t, s), \quad \boldsymbol{\xi}_\mu^\ell(t) = \frac{1}{\sqrt{N}} \mathbf{W}^\ell(0)^\top \mathbf{g}_\mu^{\ell+1}(t), \end{aligned} \quad (10)$$

where we introduced the following feature and gradient/pseudo-gradient kernels

$$\Phi_{\mu\nu}^\ell(t, s) = \frac{1}{N} \phi(\mathbf{h}_\mu^\ell(t)) \cdot \phi(\mathbf{h}_\nu^\ell(s)), \quad \tilde{G}_{\mu\nu}^\ell(t, s) = \frac{1}{N} \mathbf{g}_\mu^\ell(t) \cdot \tilde{\mathbf{g}}_\nu^\ell(s). \quad (11)$$

The particular learning rule defines the definition of the pseudo-gradient $\tilde{\mathbf{g}}_\mu^\ell(t)$. We note that, for all learning rules considered, the pseudogradient $\tilde{g}_{i,\mu}^\ell(t)$ is a function of the fields $\{h_{i,\mu}^\ell(t), z_{i,\mu}^\ell(t), m_{i,\mu}^\ell(t), \zeta_{i,\mu}^\ell(t), \tilde{\zeta}_{i,\mu}^\ell(t)\}_{\mu \in [P], t \in \mathbb{R}_+}$, conditional on the value of the kernels $\{\Phi^\ell, \tilde{G}^\ell\}$. The additional fields have definitions

$$\zeta_\mu^\ell(t) = \frac{1}{\sqrt{N}} \mathbf{W}^\ell(0)^\top \tilde{\mathbf{g}}_\mu^{\ell+1}(t), \quad \tilde{\zeta}_\mu^{\ell+1}(t) = \frac{1}{\sqrt{N}} \tilde{\mathbf{W}}^{\ell+1} \tilde{\mathbf{g}}_\mu^{\ell+1}(t) \quad (12)$$

and are specifically required for ρ -FA with $\rho > 0$ since $\tilde{\mathbf{g}}_\mu^\ell(t) = \rho \dot{\phi}(\mathbf{h}_\mu^\ell(t)) \odot \zeta_\mu^\ell(t) + \sqrt{1 - \rho^2} \dot{\phi}(\mathbf{h}_\mu^\ell(t)) \odot \tilde{\zeta}_\mu^\ell(t)$. The fields $\mathbf{m}_\mu^\ell = \frac{1}{\sqrt{D}} \mathbf{M}^\ell \mathbf{x}_\mu$ are required for GLNs.

All of the necessary fields $\{\mathbf{h}_\mu^\ell(t), \mathbf{z}_\mu^\ell(t), \tilde{\mathbf{g}}_\mu^\ell(t)\}$ are thus causal functions of the stochastic fields $\{\boldsymbol{\chi}_\mu^\ell(t), \boldsymbol{\xi}_\mu^\ell(t), \mathbf{m}_\mu^\ell, \zeta_\mu^\ell(t), \tilde{\zeta}_\mu^\ell(t)\}$ and the kernels $\{\Phi^\ell, \tilde{G}^\ell\}$. It thus suffices to characterize the distribution of these latter objects over random initialization of $\boldsymbol{\theta}(0)$ in the $N \rightarrow \infty$ limit, which we study in the next section.

B.2 MOMENT GENERATING FUNCTIONAL

We will now attempt to characterize the probability density of the random fields

$$\begin{aligned}\chi_\mu^{\ell+1}(t) &= \frac{1}{\sqrt{N}} \mathbf{W}^\ell(0) \phi(\mathbf{h}_\mu^\ell(t)), \quad \xi_\mu^\ell(t) = \frac{1}{\sqrt{N}} \mathbf{W}^\ell(0)^\top \mathbf{g}_\mu^{\ell+1}(t), \quad \mathbf{m}_\mu^\ell = \frac{1}{\sqrt{D}} \mathbf{M}^\ell \mathbf{x}_\mu \\ \zeta_\mu^\ell(t) &= \frac{1}{\sqrt{N}} \mathbf{W}^\ell(0)^\top \tilde{\mathbf{g}}_\mu^{\ell+1}(t), \quad \tilde{\zeta}_\mu^\ell(t) = \frac{1}{\sqrt{N}} \tilde{\mathbf{W}}^{\ell\top} \tilde{\mathbf{g}}_\mu^{\ell+1}(t).\end{aligned}\quad (13)$$

It is readily apparent that the fields \mathbf{m}_μ^ℓ are independent of the others and have a Gaussian distribution over random Gaussian \mathbf{M}^ℓ . These fields, therefore do not can be handled independently from the others, which are statistically coupled through the initial conditions. We will thus characterize the moment generating functional of the remaining fields $\{\chi_\mu^\ell(t), \xi_\mu^\ell(t), \zeta_\mu^\ell(t), \tilde{\zeta}_\mu^\ell(t)\}$ over random initial condition and random backward pass weights

$$\begin{aligned}Z[\{\mathbf{j}_\mu^\ell(t), \mathbf{k}_\mu^\ell(t), \mathbf{n}_\mu^\ell(t), \mathbf{p}_\mu^\ell(t)\}] \\ = \mathbb{E}_{\theta(0), \{\tilde{\mathbf{W}}^\ell\}} \exp\left(\sum_{\mu=1}^P \int_0^\infty dt \left[\mathbf{j}_\mu^\ell(t) \cdot \chi_\mu^\ell(t) + \mathbf{k}_\mu^\ell(t) \cdot \xi_\mu^\ell(t) + \mathbf{n}_\mu^\ell(t) \cdot \zeta_\mu^\ell(t) + \mathbf{p}_\mu^\ell(t) \cdot \tilde{\zeta}_\mu^\ell(t)\right]\right)\end{aligned}\quad (14)$$

where $\chi^\ell, \xi, \zeta, \tilde{\zeta}$ are regarded as functions of $\theta(0), \{\tilde{\mathbf{W}}^\ell\}$. Arbitrary moments of these random variables can be computed by differentiation of Z near zero source. For example, a two-point correlation function can be obtained as

$$\left\langle \chi_{i,\mu}^\ell(t) \zeta_{i',\nu}^{\ell'}(s) \right\rangle = \lim_{\mathbf{j}, \mathbf{k}, \mathbf{n}, \mathbf{p} \rightarrow 0} \frac{\delta}{\delta j_{i,\mu}^\ell(t)} \frac{\delta}{\delta n_{i',\nu}^{\ell'}(s)} Z[\{\mathbf{j}_\mu^\ell(t), \mathbf{k}_\mu^\ell(t), \mathbf{n}_\mu^\ell(t), \mathbf{p}_\mu^\ell(t)\}]. \quad (15)$$

More generally, we let $\boldsymbol{\mu} = (i, \mu, t)$ be a tuple containing the neuron, time, and sample index for an entry of one of these fields so that $\chi_\mu^\ell = \chi_{i,\mu}^\ell(t)$. Further, we let $\mathcal{N}_{\chi^\ell}, \mathcal{N}_{\xi^\ell}, \mathcal{N}_{\zeta^\ell}, \mathcal{N}_{\tilde{\zeta}^\ell}$ be index sets which contain sample and time indices as well as neuron indices $\mathcal{N}_\chi = \{\mu_1^\chi, \dots, |\mu|_{\mathcal{N}_\chi}^\chi\}$ for all of the indices we wish to compute an average over. Then arbitrary moments can be computed with the formula

$$\begin{aligned}\left\langle \prod_\ell \left[\prod_{\mu \in \mathcal{N}_{\chi^\ell}} \chi_\mu^\ell \prod_{\nu \in \mathcal{N}_{\xi^\ell}} \xi_\nu^\ell \prod_{\alpha \in \mathcal{N}_{\zeta^\ell}} \zeta_\alpha^\ell \prod_{\beta \in \mathcal{N}_{\tilde{\zeta}^\ell}} \tilde{\zeta}_\beta^\ell \right] \right\rangle \\ = \lim_{\mathbf{j}, \mathbf{k}, \mathbf{n}, \mathbf{p} \rightarrow 0} \prod_\ell \left[\prod_{\mu \in \mathcal{N}_{\chi^\ell}} \frac{\delta}{\delta j_\mu^\ell} \prod_{\nu \in \mathcal{N}_{\xi^\ell}} \frac{\delta}{\delta k_\nu^\ell} \prod_{\alpha \in \mathcal{N}_{\zeta^\ell}} \frac{\delta}{\delta n_\alpha^\ell} \prod_{\beta \in \mathcal{N}_{\tilde{\zeta}^\ell}} \frac{\delta}{\delta p_\beta^\ell} \right] Z[\{\mathbf{j}_\mu^\ell(t), \mathbf{k}_\mu^\ell(t), \mathbf{n}_\mu^\ell(t), \mathbf{p}_\mu^\ell(t)\}].\end{aligned}\quad (16)$$

We now to study this moment generating functional Z in the large width $N \rightarrow \infty$ limit.

B.3 PATH INTEGRAL FORMULATION AND INTEGRATION OVER WEIGHTS

To enable the average over the weights, we multiply Z by an integral representation of unity that enforces the relationship between $\chi_\mu^{\ell+1}(t), \mathbf{W}^\ell(0), \phi(\mathbf{h}_\mu^\ell(t))$

$$\begin{aligned}1 &= \int_{\mathbb{R}^N} d\chi_\mu^{\ell+1}(t) \delta\left(\chi_\mu^{\ell+1}(t) - \frac{1}{\sqrt{N}} \mathbf{W}^\ell(0) \phi(\mathbf{h}_\mu^\ell(t))\right) \\ &= \int_{\mathbb{R}^N} \int_{\mathbb{R}^N} \frac{d\chi_\mu^{\ell+1}(t) d\hat{\chi}_\mu^{\ell+1}(t)}{(2\pi)^N} \exp\left(i\hat{\chi}_\mu^{\ell+1}(t) \cdot \left[\chi_\mu^{\ell+1}(t) - \frac{1}{\sqrt{N}} \mathbf{W}^\ell(0) \phi(\mathbf{h}_\mu^\ell(t))\right]\right).\end{aligned}\quad (17)$$

In the second line, we used the Fourier representation of the Dirac-Delta function for each of the N neuron indices $\delta(r) = \int_{-\infty}^{\infty} \frac{d\hat{r}}{2\pi} \exp(i\hat{r}r)$. We repeat this procedure for the other fields $\xi_\mu^\ell(t), \zeta_\mu^\ell(t), \tilde{\zeta}_\mu^\ell(t)$ at each time t and each sample μ . After inserting these delta functions, we find

the following form of the moment generating functional

$$\begin{aligned}
Z &= \int \prod_{\ell, \mu} \frac{d\chi_{\mu}^{\ell}(t) d\hat{\chi}_{\mu}^{\ell}(t)}{(2\pi)^N} \frac{d\xi_{\mu}^{\ell}(t) d\hat{\xi}_{\mu}^{\ell}(t)}{(2\pi)^N} \frac{d\zeta_{\mu}^{\ell}(t) d\hat{\zeta}_{\mu}^{\ell}(t)}{(2\pi)^N} \frac{d\tilde{\zeta}_{\mu}^{\ell}(t) d\hat{\tilde{\zeta}}_{\mu}^{\ell}(t)}{(2\pi)^N} \\
&\times \exp \left(\int_0^{\infty} dt \sum_{\ell, \mu} \left[\chi_{\mu}^{\ell}(t) \cdot (\mathbf{j}_{\mu}^{\ell}(t) + i\hat{\chi}_{\mu}^{\ell}(t)) + \xi_{\mu}^{\ell}(t) \cdot (\mathbf{k}_{\mu}^{\ell}(t) + i\hat{\xi}_{\mu}^{\ell}(t)) \right] \right) \\
&\times \exp \left(\int_0^{\infty} dt \sum_{\ell, \mu} \left[\zeta_{\mu}^{\ell}(t) \cdot (\mathbf{n}_{\mu}^{\ell}(t) + i\hat{\zeta}_{\mu}^{\ell}(t)) + \tilde{\zeta}_{\mu}^{\ell}(t) \cdot (\mathbf{p}_{\mu}^{\ell}(t) + i\hat{\tilde{\zeta}}_{\mu}^{\ell}(t)) \right] \right) \\
&\times \prod_{\ell} \mathbb{E}_{\mathbf{W}^{\ell}(0)} \exp \left(-\frac{i}{\sqrt{N}} \text{Tr} \mathbf{W}^{\ell}(0)^{\top} \left[\int dt \sum_{\mu} \hat{\chi}_{\mu}^{\ell+1}(t) \phi(\mathbf{h}_{\mu}^{\ell}(t))^{\top} + \mathbf{g}_{\mu}^{\ell+1}(t) \hat{\xi}_{\mu}^{\ell}(t)^{\top} \right] \right) \\
&\times \exp \left(-\frac{i}{\sqrt{N}} \mathbf{W}^{\ell}(0)^{\top} \left[\int dt \sum_{\mu} \tilde{\mathbf{g}}_{\mu}^{\ell+1}(t) \zeta_{\mu}^{\ell}(t)^{\top} \right] \right) \\
&\times \prod_{\ell} \mathbb{E}_{\tilde{\mathbf{W}}^{\ell}} \exp \left(-\frac{i}{\sqrt{N}} \text{Tr} \tilde{\mathbf{W}}^{\ell \top} \left[\int dt \sum_{\mu} \tilde{\mathbf{g}}_{\mu}^{\ell+1}(t) \hat{\tilde{\zeta}}_{\mu}^{\ell}(t)^{\top} \right] \right). \tag{18}
\end{aligned}$$

We see that we often have simultaneous integrals over time t and sums over samples μ so we will again adopt a shorthand notation for indices $\boldsymbol{\mu} = (\mu, t)$ and define a summation convention $\sum_{\boldsymbol{\mu}} a_{\boldsymbol{\mu}} b_{\boldsymbol{\mu}} = \int_0^{\infty} dt \sum_{\mu=1}^P a_{\boldsymbol{\mu}}(t) b_{\boldsymbol{\mu}}(t)$. To perform the averages over weights, we note that for a standard normal variable W_{ij} , that $\mathbb{E}_{W_{ij}} \exp(iW_{ij}a_i b_j) = \exp(-\frac{1}{2}a_i^2 b_j^2)$. Using this fact for each of the weight matrix averages, we have

$$\begin{aligned}
&\mathbb{E}_{\mathbf{W}^{\ell}(0)} \exp \left(-\frac{i}{\sqrt{N}} \text{Tr} \mathbf{W}^{\ell}(0)^{\top} \left[\sum_{\boldsymbol{\mu}} \hat{\chi}_{\boldsymbol{\mu}}^{\ell+1} \phi(\mathbf{h}_{\boldsymbol{\mu}}^{\ell})^{\top} + \mathbf{g}_{\boldsymbol{\mu}}^{\ell+1} \hat{\xi}_{\boldsymbol{\mu}}^{\ell \top} + \tilde{\mathbf{g}}_{\boldsymbol{\mu}}^{\ell+1} \hat{\tilde{\zeta}}_{\boldsymbol{\mu}}^{\ell \top} \right] \right) \\
&= \exp \left(-\frac{1}{2N} \left| \sum_{\boldsymbol{\mu}} \hat{\chi}_{\boldsymbol{\mu}}^{\ell+1} \phi(\mathbf{h}_{\boldsymbol{\mu}}^{\ell})^{\top} + \mathbf{g}_{\boldsymbol{\mu}}^{\ell+1} \hat{\xi}_{\boldsymbol{\mu}}^{\ell \top} + \tilde{\mathbf{g}}_{\boldsymbol{\mu}}^{\ell+1} \hat{\tilde{\zeta}}_{\boldsymbol{\mu}}^{\ell \top} \right|_F^2 \right) \\
&= \exp \left(-\frac{1}{2} \sum_{\boldsymbol{\mu}, \boldsymbol{\nu}} \left[\hat{\chi}_{\boldsymbol{\mu}}^{\ell+1} \cdot \hat{\chi}_{\boldsymbol{\nu}}^{\ell+1} \Phi_{\boldsymbol{\mu}, \boldsymbol{\nu}}^{\ell} + \hat{\xi}_{\boldsymbol{\mu}}^{\ell} \cdot \hat{\xi}_{\boldsymbol{\nu}}^{\ell} G_{\boldsymbol{\mu}, \boldsymbol{\nu}}^{\ell+1} + \hat{\zeta}_{\boldsymbol{\mu}}^{\ell} \cdot \hat{\zeta}_{\boldsymbol{\nu}}^{\ell} \tilde{G}_{\boldsymbol{\mu}, \boldsymbol{\nu}}^{\ell+1} + \hat{\xi}_{\boldsymbol{\mu}}^{\ell} \cdot \hat{\zeta}_{\boldsymbol{\nu}}^{\ell} \tilde{G}_{\boldsymbol{\mu}, \boldsymbol{\nu}}^{\ell+1} \right] \right) \\
&\times \exp \left(-i \sum_{\boldsymbol{\mu}, \boldsymbol{\nu}} \left[\hat{\chi}_{\boldsymbol{\mu}}^{\ell+1} \cdot \mathbf{g}_{\boldsymbol{\nu}}^{\ell+1} A_{\boldsymbol{\mu}, \boldsymbol{\nu}}^{\ell} + \hat{\chi}_{\boldsymbol{\mu}}^{\ell+1} \cdot \tilde{\mathbf{g}}_{\boldsymbol{\nu}}^{\ell+1} C_{\boldsymbol{\mu}, \boldsymbol{\nu}}^{\ell} \right] \right). \tag{19}
\end{aligned}$$

In the above, we introduced a collection of order parameters $\{\Phi, G, \tilde{G}, \tilde{\tilde{G}}, A, C\}$, which will correspond to correlation and response functions of our DMFT. These are defined as

$$\begin{aligned}
\Phi_{\boldsymbol{\mu}, \boldsymbol{\nu}}^{\ell} &= \frac{1}{N} \phi(\mathbf{h}_{\boldsymbol{\mu}}^{\ell}) \cdot \phi(\mathbf{h}_{\boldsymbol{\nu}}^{\ell}), \quad G_{\boldsymbol{\mu}, \boldsymbol{\nu}}^{\ell} = \frac{1}{N} \mathbf{g}_{\boldsymbol{\mu}}^{\ell} \cdot \mathbf{g}_{\boldsymbol{\nu}}^{\ell}, \quad \tilde{G}_{\boldsymbol{\mu}, \boldsymbol{\nu}}^{\ell} = \frac{1}{N} \mathbf{g}_{\boldsymbol{\mu}}^{\ell} \cdot \tilde{\mathbf{g}}_{\boldsymbol{\nu}}^{\ell} \\
\tilde{\tilde{G}}_{\boldsymbol{\mu}, \boldsymbol{\nu}}^{\ell+1} &= \frac{1}{N} \tilde{\mathbf{g}}_{\boldsymbol{\mu}}^{\ell+1} \cdot \tilde{\mathbf{g}}_{\boldsymbol{\nu}}^{\ell+1}, \quad iA_{\boldsymbol{\mu}, \boldsymbol{\nu}}^{\ell} = \frac{1}{N} \phi(\mathbf{h}_{\boldsymbol{\mu}}^{\ell}) \cdot \hat{\xi}_{\boldsymbol{\nu}}^{\ell}, \quad iC_{\boldsymbol{\mu}, \boldsymbol{\nu}}^{\ell} = \frac{1}{N} \phi(\mathbf{h}_{\boldsymbol{\mu}}^{\ell}) \cdot \hat{\zeta}_{\boldsymbol{\nu}}^{\ell}. \tag{20}
\end{aligned}$$

We perform a similar average over $\tilde{\mathbf{W}}^{\ell}$ can be obtained directly

$$\mathbb{E}_{\tilde{\mathbf{W}}^{\ell}} \exp \left(-\frac{i}{\sqrt{N}} \text{Tr} \tilde{\mathbf{W}}^{\ell \top} \left[\sum_{\boldsymbol{\mu}} \tilde{\mathbf{g}}_{\boldsymbol{\mu}}^{\ell+1} \hat{\tilde{\zeta}}_{\boldsymbol{\mu}}^{\ell \top} \right] \right) = \exp \left(-\frac{1}{2} \sum_{\boldsymbol{\mu}, \boldsymbol{\nu}} \hat{\zeta}_{\boldsymbol{\mu}}^{\ell} \cdot \hat{\zeta}_{\boldsymbol{\nu}}^{\ell} \tilde{\tilde{G}}_{\boldsymbol{\mu}, \boldsymbol{\nu}}^{\ell+1} \right). \tag{21}$$

Now that we have defined our collection of order parameters, we enforce their definitions with Dirac-Delta functions by multiplying by one. For example,

$$\begin{aligned}
1 &= N \int d\Phi_{\boldsymbol{\mu}, \boldsymbol{\nu}}^{\ell} \delta(N\Phi_{\boldsymbol{\mu}, \boldsymbol{\nu}}^{\ell} - \phi(\mathbf{h}_{\boldsymbol{\mu}}^{\ell}) \cdot \phi(\mathbf{h}_{\boldsymbol{\nu}}^{\ell})) \\
&= \int \frac{d\Phi_{\boldsymbol{\mu}, \boldsymbol{\nu}}^{\ell} d\hat{\Phi}_{\boldsymbol{\mu}, \boldsymbol{\nu}}^{\ell}}{2\pi N^{-1}} \exp \left(N\hat{\Phi}_{\boldsymbol{\mu}, \boldsymbol{\nu}}^{\ell} \Phi_{\boldsymbol{\mu}, \boldsymbol{\nu}}^{\ell} - \hat{\Phi}_{\boldsymbol{\mu}, \boldsymbol{\nu}}^{\ell} \phi(\mathbf{h}_{\boldsymbol{\mu}}^{\ell}) \cdot \phi(\mathbf{h}_{\boldsymbol{\nu}}^{\ell}) \right). \tag{22}
\end{aligned}$$

We enforce these definitions for all order parameters $\{\Phi_{\mu\nu}^\ell, G_{\mu\nu}^\ell, \tilde{G}_{\mu\nu}^\ell, \hat{G}_{\mu\nu}^\ell, A_{\mu,\nu}^\ell, C_{\mu,\nu}^\ell\}$. We let the corresponding Fourier duals for each of these order parameters be $\{\hat{\Phi}_{\mu\nu}^\ell, \hat{G}_{\mu\nu}^\ell, \hat{\tilde{G}}_{\mu\nu}^\ell, \hat{\hat{G}}_{\mu\nu}^\ell, -B_{\mu,\nu}^\ell, -D_{\mu,\nu}^\ell\}$. In the next section we show the resulting formula for the moment generating function and take the $N \rightarrow \infty$ limit to derive our DMFT equations.

B.4 DMFT ACTION

After inserting the Dirac-Delta functions to enforce the definitions of the order parameters, we derive the following moment generating functional in terms of $\mathbf{q} = \{\Phi, \hat{\Phi}, G, \hat{G}, \tilde{G}, \hat{\tilde{G}}, \hat{\hat{G}}, A, B, C, D, j, k, n, p\}$

$$Z = \int \prod_{\ell, \mu, \nu} \frac{d\Phi_{\mu\nu}^\ell d\hat{\Phi}_{\mu\nu}^\ell}{2\pi N^{-1}} \frac{dG_{\mu\nu}^\ell d\hat{G}_{\mu\nu}^\ell}{2\pi N^{-1}} \frac{d\tilde{G}_{\mu\nu}^\ell d\hat{\tilde{G}}_{\mu\nu}^\ell}{2\pi N^{-1}} \frac{d\hat{\hat{G}}_{\mu\nu}^\ell d\tilde{\hat{\hat{G}}}_{\mu\nu}^\ell}{2\pi N^{-1}} \frac{dA_{\mu\nu}^\ell dB_{\mu\nu}^\ell}{2\pi N^{-1}} \frac{dC_{\mu\nu}^\ell dD_{\mu\nu}^\ell}{2\pi N^{-1}} \exp(NS[\mathbf{q}])$$

where $S[\mathbf{q}]$ is the $O_N(1)$ DMFT action which takes the form

$$S[\mathbf{q}] = \sum_{\ell, \mu, \nu} \left[\Phi_{\mu,\nu}^\ell \hat{\Phi}_{\mu,\nu}^\ell + G_{\mu\nu}^\ell \hat{G}_{\mu\nu}^\ell + \tilde{G}_{\mu\nu}^\ell \hat{\tilde{G}}_{\mu\nu}^\ell + \hat{\tilde{G}}_{\mu\nu}^\ell \tilde{\hat{\hat{G}}}_{\mu\nu}^\ell - A_{\mu\nu}^\ell B_{\mu\nu}^\ell - C_{\mu\nu}^\ell D_{\mu\nu}^\ell \right] + \frac{1}{N} \sum_{i=1}^N \sum_{\ell=1}^L \ln \mathcal{Z}_i^\ell[\mathbf{q}]. \quad (23)$$

The single-site moment generating functionals (MGF) \mathcal{Z}_i^ℓ involve only the integrals with sources $\{j_i^\ell, k_i^\ell, n_i^\ell, p_i^\ell\}$ for neuron $i \in [N]$ in layer ℓ . For a given set of order parameters \mathbf{q} at zero source, these functionals become identical across all neuron sites i . Concretely, for any $\ell \in [L], i \in [N]$, the single site MGF takes the form

$$\begin{aligned} \mathcal{Z}_i^\ell = & \int \prod_{\mu} \frac{d\chi_\mu^\ell d\hat{\chi}_\mu^\ell}{2\pi} \frac{d\xi_\mu^\ell d\hat{\xi}_\mu^\ell}{2\pi} \frac{d\zeta_\mu^\ell d\hat{\zeta}_\mu^\ell}{2\pi} \frac{d\tilde{\zeta}_\mu^\ell d\hat{\tilde{\zeta}}_\mu^\ell}{2\pi} \quad (24) \\ & \exp \left(-\frac{1}{2} \sum_{\mu, \nu} \left[\hat{\chi}_\mu^{\ell+1} \hat{\chi}_\nu^{\ell+1} \Phi_{\mu,\nu}^\ell + \hat{\xi}_\mu^\ell \hat{\xi}_\nu^\ell G_{\mu\nu}^{\ell+1} + \hat{\zeta}_\mu^\ell \hat{\zeta}_\nu^\ell \tilde{G}_{\mu\nu}^{\ell+1} \right] \right) \\ & \exp \left(-\frac{1}{2} \sum_{\mu, \nu} \left[\hat{\zeta}_\mu^\ell \hat{\zeta}_\nu^\ell \tilde{G}_{\mu,\nu}^{\ell+1} + 2\hat{\xi}_\mu^\ell \hat{\xi}_\nu^\ell \tilde{G}_{\mu,\nu}^{\ell+1} \right] \right) \\ & \exp \left(-\sum_{\mu, \nu} \left[\hat{\Phi}_{\mu\nu}^\ell \phi(h_\mu^\ell) \phi(h_\nu^\ell) + \hat{G}_{\mu\nu}^\ell g_\mu^\ell g_\nu^\ell + \hat{\tilde{G}}_{\mu\nu}^\ell g_\mu^\ell \tilde{g}_\nu^\ell + \hat{\hat{G}}_{\mu\nu}^\ell \tilde{g}_\mu^\ell \tilde{g}_\nu^\ell \right] \right) \\ & \exp \left(-i \sum_{\mu, \nu} \left[\hat{\chi}_\mu^{\ell+1} g_\nu^{\ell+1} A_{\mu\nu}^\ell + \hat{\chi}_\mu^{\ell+1} \tilde{g}_\nu^{\ell+1} C_{\mu\nu}^\ell + \phi(h_\mu^\ell) \hat{\xi}_\nu^\ell B_{\mu\nu}^\ell + \phi(h_\mu^\ell) \hat{\zeta}_\nu^\ell D_{\mu\nu}^\ell \right] \right) \\ & \exp \left(\sum_{\mu} \left[\chi_\mu^\ell (j_{i,\mu}^\ell + i\hat{\chi}_\mu^\ell) + \xi_\mu^\ell (k_{i,\mu}^\ell + i\hat{\xi}_\mu^\ell) + \zeta_\mu^\ell (n_{i,\mu}^\ell + i\hat{\zeta}_\mu^\ell) + \tilde{\zeta}_\mu^\ell (p_{i,\mu}^\ell + i\hat{\tilde{\zeta}}_\mu^\ell) \right] \right). \end{aligned}$$

As promised, the only terms in \mathcal{Z}_i which vary over site index i are the sources $\{j, k, n, p\}$. To simplify our later saddle point equations, we will abstract the notation for the single site MGF, letting

$$\mathcal{Z}_i^\ell = \int \prod_{\mu} \frac{d\chi_\mu d\hat{\chi}_\mu}{2\pi} \frac{d\xi_\mu d\hat{\xi}_\mu}{2\pi} \frac{d\zeta_\mu d\hat{\zeta}_\mu}{2\pi} \frac{d\tilde{\zeta}_\mu d\hat{\tilde{\zeta}}_\mu}{2\pi} \exp \left(-\mathcal{H}_i^\ell[\chi, \hat{\chi}, \xi, \hat{\xi}, \zeta, \hat{\zeta}, \tilde{\zeta}, \hat{\tilde{\zeta}}] \right) \quad (25)$$

where \mathcal{H}_i^ℓ is the single site effective Hamiltonian for neuron i and layer ℓ . Note that at zero source, \mathcal{H}_i^ℓ are identical for all $i \in [N]$.

B.5 SADDLE POINT EQUATIONS

Letting the full collection of concatenated order parameters \mathbf{q} be indexed by b . We now take the $N \rightarrow \infty$ limit, using the method of steepest descent

$$Z = \int \prod_b \frac{\sqrt{N} dq_b}{\sqrt{2\pi}} \exp(NS[\mathbf{q}]) \sim \exp(NS[\mathbf{q}^*]) , \nabla S[\mathbf{q}]|_{\mathbf{q}^*} = 0 , N \rightarrow \infty. \quad (26)$$

We see that the integral over \mathbf{q} is exponentially dominated by the saddle point where $\nabla S[\mathbf{q}] = 0$. We thus need to solve these saddle point equations for the \mathbf{q}^* . To do this, we need to introduce some notation. Let $O(\chi, \hat{\chi}, \xi, \hat{\xi}, \zeta, \hat{\zeta}, \tilde{\zeta}, \hat{\tilde{\zeta}})$ be an arbitrary function of the single site stochastic processes. We define the ℓ -th layer i -th single site average, denoted by $\langle O(\chi, \hat{\chi}, \xi, \hat{\xi}, \zeta, \hat{\zeta}, \tilde{\zeta}, \hat{\tilde{\zeta}}) \rangle_{\ell, i}$ as

$$\begin{aligned} \langle O(\chi, \hat{\chi}, \xi, \hat{\xi}, \zeta, \hat{\zeta}, \tilde{\zeta}, \hat{\tilde{\zeta}}) \rangle_{\ell, i} &= \frac{1}{Z_i^\ell} \int \prod_{\mu} \frac{d\chi_{\mu} d\hat{\chi}_{\mu}}{2\pi} \frac{d\xi_{\mu} d\hat{\xi}_{\mu}}{2\pi} \frac{d\zeta_{\mu} d\hat{\zeta}_{\mu}}{2\pi} \frac{d\tilde{\zeta}_{\mu} d\hat{\tilde{\zeta}}_{\mu}}{2\pi} \\ &\quad \exp\left(-\mathcal{H}_i^\ell[\chi, \hat{\chi}, \xi, \hat{\xi}, \zeta, \hat{\zeta}, \tilde{\zeta}, \hat{\tilde{\zeta}}]\right) O(\chi, \hat{\chi}, \xi, \hat{\xi}, \zeta, \hat{\zeta}, \tilde{\zeta}, \hat{\tilde{\zeta}}) \end{aligned} \quad (27)$$

which can be interpreted as an average over the Gibbs measure defined by energy \mathcal{H}_i^ℓ . With this notation, we now set about computing the saddle point equations which define the primal order parameters $\{\Phi, G, \tilde{G}, \hat{\tilde{G}}\}$.

$$\begin{aligned} \frac{\partial S}{\partial \Phi_{\mu\nu}^\ell} &= \Phi_{\mu\nu}^\ell - \frac{1}{N} \sum_{i=1}^N \langle \phi(h_{\mu}^\ell) \phi(h_{\nu}^\ell) \rangle_{\ell, i} = 0 \\ \frac{\partial S}{\partial G_{\mu\nu}^\ell} &= G_{\mu\nu}^\ell - \frac{1}{N} \sum_{i=1}^N \langle g_{\mu}^\ell g_{\nu}^\ell \rangle_{\ell, i} = 0 \\ \frac{\partial S}{\partial \tilde{G}_{\mu\nu}^\ell} &= \tilde{G}_{\mu\nu}^\ell - \frac{1}{N} \sum_{i=1}^N \langle g_{\mu}^\ell \tilde{g}_{\nu}^\ell \rangle_{\ell, i} = 0 \\ \frac{\partial S}{\partial \hat{\tilde{G}}_{\mu\nu}^\ell} &= \hat{\tilde{G}}_{\mu\nu}^\ell - \frac{1}{N} \sum_{i=1}^N \langle \tilde{g}_{\mu}^\ell \hat{g}_{\nu}^\ell \rangle_{\ell, i} = 0 \end{aligned}$$

We further compute the saddle point equations for the dual order parameters

$$\begin{aligned}
\frac{\partial S}{\partial \Phi_{\mu\nu}^\ell} &= \hat{\Phi}_{\mu\nu}^\ell - \frac{1}{2N} \sum_{i=1}^N \langle \hat{\chi}_\mu^{\ell+1} \hat{\chi}_\nu^{\ell+1} \rangle_{\ell+1,i} = 0 \\
\frac{\partial S}{\partial G_{\mu\nu}^\ell} &= \hat{G}_{\mu\nu}^\ell - \frac{1}{2N} \sum_{i=1}^N \langle \hat{\xi}_\mu^{\ell-1} \hat{\xi}_\nu^{\ell-1} \rangle_{\ell-1,i} = 0 \\
\frac{\partial S}{\partial \tilde{G}_{\mu\nu}^\ell} &= \hat{G}_{\mu\nu}^\ell - \frac{1}{N} \sum_{i=1}^N \langle \hat{\xi}_\mu^{\ell-1} \hat{\xi}_\nu^{\ell-1} \rangle_{\ell-1,i} = 0 \\
\frac{\partial S}{\partial \tilde{\tilde{G}}_{\mu\nu}^\ell} &= \hat{G}_{\mu\nu}^\ell - \frac{1}{2N} \sum_{i=1}^N \langle [\hat{\xi}_\mu^{\ell-1} \hat{\xi}_\nu^{\ell-1} + \hat{\zeta}_\mu^{\ell-1} \hat{\zeta}_\nu^{\ell-1}] \rangle_{\ell-1,i} = 0 \\
\frac{\partial S}{\partial A_{\mu\nu}^\ell} &= -B_{\mu\nu}^\ell - \frac{i}{N} \sum_{i=1}^N \langle \hat{\chi}_\mu^{\ell+1} g_\nu^{\ell+1} \rangle_{\ell+1,i} = 0 \\
\frac{\partial S}{\partial B_{\mu\nu}^\ell} &= -A_{\mu\nu}^\ell - \frac{i}{N} \sum_{i=1}^N \langle \phi(h_\mu^\ell) \hat{\xi}_\nu^\ell \rangle_{\ell,i} = 0 \\
\frac{\partial S}{\partial C_{\mu\nu}^\ell} &= -D_{\mu\nu}^\ell - \frac{i}{N} \sum_{i=1}^N \langle \hat{\chi}_\mu^{\ell+1} \tilde{g}_\nu^{\ell+1} \rangle_{\ell+1,i} = 0 \\
\frac{\partial S}{\partial D_{\mu\nu}^\ell} &= -C_{\mu\nu}^\ell - \frac{i}{N} \sum_{i=1}^N \langle \phi(h_\mu^\ell) \hat{\zeta}_\nu^\ell \rangle_{\ell,i} = 0.
\end{aligned} \tag{28}$$

The correlation functions involving real variables $\{h, g, \tilde{g}\}$ have a straightforward interpretation. However, it is not immediately clear what to do with terms involving the dual fields $\{\hat{\chi}, \hat{\xi}, \hat{\zeta}\}$. As a starting example, let's consider one of the terms for $B_{\mu\nu}^{\ell-1}$, namely $-i \langle \hat{\chi}_\nu^\ell g_\nu^\ell \rangle$. We make progress by inserting another fictitious source term u_μ^ℓ and differentiating near zero source

$$-i \langle \hat{\chi}_\nu^\ell g_\nu^\ell \rangle_i = \lim_{\{u_\mu\} \rightarrow 0} \frac{\partial}{\partial u_\nu^\ell} \left\langle g_\nu^\ell \exp \left(-i \sum_{\nu'} u_{\nu'} \hat{\chi}_{\nu'}^\ell \right) \right\rangle_i. \tag{29}$$

Introducing a vectorization notation $\mathbf{u}^\ell = \text{Vec}\{u_\mu^\ell\}_\mu$, $\hat{\chi}^\ell = \text{Vec}\{\hat{\chi}_\mu^\ell\}_\mu$ and $\Phi^{\ell-1} = \text{Mat}\{\Phi_{\mu,\nu}^{\ell-1}\}_{\mu,\nu}$, we can perform the internal integrals over $\hat{\chi}^\ell$

$$\begin{aligned}
& \int \prod_\mu \frac{d\hat{\chi}_\mu^\ell}{\sqrt{2\pi}} \exp \left(-\frac{1}{2} \hat{\chi}^{\ell\top} \Phi^{\ell-1} \hat{\chi}^\ell + i \hat{\chi}^\ell \cdot (\boldsymbol{\chi}^\ell - \mathbf{u}^\ell - \mathbf{A}^{\ell-1} \mathbf{g}^\ell - \mathbf{C}^{\ell-1} \tilde{\mathbf{g}}^\ell) \right) \\
&= \exp \left(-\frac{1}{2} (\boldsymbol{\chi}^\ell - \mathbf{u}^\ell - \mathbf{A}^{\ell-1} \mathbf{g}^\ell - \mathbf{C}^{\ell-1} \tilde{\mathbf{g}}^\ell) [\Phi^{\ell-1}]^{-1} (\boldsymbol{\chi}^\ell - \mathbf{u}^\ell - \mathbf{A}^{\ell-1} \mathbf{g}^\ell - \mathbf{C}^{\ell-1} \tilde{\mathbf{g}}^\ell) \right) \\
& \quad \exp \left(-\frac{1}{2} \ln \det \Phi^{\ell-1} \right).
\end{aligned} \tag{30}$$

We thus need to compute a derivative of the above function with respect to \mathbf{u}^ℓ at $\mathbf{u}^\ell = 0$, which gives

$$-i \langle \hat{\chi}^\ell g^{\ell\top} \rangle_i = [\Phi^{\ell-1}]^{-1} \langle (\boldsymbol{\chi}^\ell - \mathbf{A}^{\ell-1} \mathbf{g}^\ell - \mathbf{C}^{\ell-1} \tilde{\mathbf{g}}^\ell) g^{\ell\top} \rangle_i. \tag{31}$$

From the above reasoning, we can also easily obtain $\hat{\Phi}^{\ell-1}$ using

$$\begin{aligned}
\langle \hat{\chi}^\ell \hat{\chi}^{\ell\top} \rangle &= -\frac{\partial^2}{\partial \mathbf{u}^\ell \partial \mathbf{u}^{\ell\top}} \Big|_{\mathbf{u}=0} \langle \exp(-i \mathbf{u}^\ell \cdot \hat{\chi}^\ell) \rangle \\
&= -\int d\chi^\ell \dots \frac{\partial^2}{\partial \mathbf{u}^\ell \partial \mathbf{u}^{\ell\top}} \Big|_{\mathbf{u}=0} \\
&\times \exp\left(-\frac{1}{2}(\chi^\ell - \mathbf{u}^\ell - \mathbf{A}^{\ell-1} \mathbf{g}^\ell - \mathbf{C}^{\ell-1} \tilde{\mathbf{g}}^\ell) [\Phi^{\ell-1}]^{-1} (\chi^\ell - \mathbf{u}^\ell - \mathbf{A}^{\ell-1} \mathbf{g}^\ell - \mathbf{C}^{\ell-1} \tilde{\mathbf{g}}^\ell) - \dots\right) \\
&= [\Phi^{\ell-1}]^{-1} - [\Phi^{\ell-1}]^{-1} \langle (\chi^\ell - \mathbf{A}^{\ell-1} \mathbf{g}^\ell - \mathbf{C}^{\ell-1} \tilde{\mathbf{g}}^\ell) (\chi^\ell - \mathbf{A}^{\ell-1} \mathbf{g}^\ell - \mathbf{C}^{\ell-1} \tilde{\mathbf{g}}^\ell)^\top \rangle [\Phi^{\ell-1}]^{-1}.
\end{aligned} \tag{32}$$

Performing a similar analysis, we insert source fields $\mathbf{r}_+ = \begin{bmatrix} \mathbf{r}^\ell \\ \mathbf{v}^\ell \end{bmatrix}$ for $\hat{\xi}_+^\ell = \begin{bmatrix} \hat{\xi}^\ell \\ \hat{\zeta}^\ell \end{bmatrix}$ and define $\mathbf{G}_+^{\ell+1} = \begin{bmatrix} \mathbf{G}^{\ell+1} & \tilde{\mathbf{G}}^{\ell+1} \\ \tilde{\mathbf{G}}^{\ell+1\top} & \tilde{\tilde{\mathbf{G}}}^{\ell+1} \end{bmatrix}$, and $\mathbf{B}_+^\ell = [\mathbf{B}^\ell \quad \mathbf{D}^\ell]$ and then we can compute the necessary averages using the same technique

$$\begin{aligned}
-i \langle \phi(\mathbf{h}^\ell) \hat{\xi}_+^{\ell\top} \rangle &= \frac{\partial}{\partial \mathbf{r}_+^\ell} \Big|_{\mathbf{r}_+^\ell=0} \langle \phi(\mathbf{h}^\ell) \exp(-i \mathbf{r}_+^\ell \cdot \hat{\xi}_+^\ell) \rangle = \langle \phi(\mathbf{h}^\ell) (\xi_+^\ell - \mathbf{B}_+^{\ell\top} \phi(\mathbf{h}^\ell))^\top \rangle \\
\langle \hat{\xi}_+^\ell \hat{\xi}_+^{\ell\top} \rangle &= -\frac{\partial^2}{\partial \mathbf{r}_+^\ell \partial \mathbf{r}_+^{\ell\top}} \Big|_{\mathbf{r}_+^\ell=0} \langle \exp(-i \mathbf{r}_+^\ell \cdot \hat{\xi}_+^\ell) \rangle \\
&= [\mathbf{G}_+^{\ell+1}]^{-1} - [\mathbf{G}_+^{\ell+1}]^{-1} \langle (\xi_+^\ell - \mathbf{B}_+^{\ell\top} \phi(\mathbf{h}^\ell)) (\xi_+^\ell - \mathbf{B}_+^{\ell\top} \phi(\mathbf{h}^\ell))^\top \rangle [\mathbf{G}_+^{\ell+1}]^{-1}.
\end{aligned} \tag{33}$$

We now have formulas for all the necessary averages entirely in terms of the primal fields $\{\chi, \xi, \zeta, \tilde{\zeta}\}$.

B.6 LINEARIZING WITH THE HUBBARD TRICK

Now, using the fact that in the $N \rightarrow \infty$ limit \mathbf{q} concentrates around \mathbf{q}^* , we will simplify our single site stochastic processes so we can obtain a final formula for $\{A, B, C, D, \hat{\Phi}, \hat{G}, \tilde{\hat{G}}, \tilde{\tilde{\hat{G}}}\}$. To do so, we utilize the Hubbard-Stratanovich identity

$$\exp\left(-\frac{\sigma^2}{2} k^2\right) = \int \frac{du}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{1}{2\sigma^2} u^2 - iku\right), \tag{34}$$

which is merely a consequence of the Fourier transform of the Gaussian distribution. This is often referred to as ‘‘linearizing’’ the action since the a term quadratic in k was replaced with an average of an action which is linear in k . In our setting, we perform this trick on a collection of variables which appear in the quadratic forms of our single site MGFs \mathcal{Z}_i^ℓ . For example, for the $\hat{\chi}^{\ell+1}$ fields, we have

$$\exp\left(-\frac{1}{2} \sum_{\mu\nu} \hat{\chi}_\mu^{\ell+1} \hat{\chi}_\nu^{\ell+1} \Phi_{\mu\nu}^\ell\right) = \left\langle \exp\left(-i \sum_{\mu} \hat{\chi}_\mu^{\ell+1} u_\mu^{\ell+1}\right) \right\rangle_{\{u_\mu^{\ell+1}\} \sim \mathcal{N}(0, \Phi^\ell)}. \tag{35}$$

Similarly, we perform a joint decomposition for the $\{\hat{\xi}^\ell, \hat{\zeta}^\ell\}$ fields which gives

$$\begin{aligned}
&\exp\left(-\frac{1}{2} \sum_{\mu\nu} \left[\hat{\xi}_\mu^\ell \hat{\xi}_\nu^\ell G_{\mu,\nu}^{\ell+1} + 2\hat{\xi}_\mu^\ell \hat{\zeta}_\nu^\ell \tilde{G}_{\mu\nu}^{\ell+1} + \hat{\zeta}_\mu^\ell \hat{\zeta}_\nu^\ell \tilde{\tilde{G}}_{\mu\nu}^{\ell+1} \right]\right) \\
&= \left\langle \exp\left(-i \sum_{\mu} [r_\mu^\ell \hat{\xi}_\mu^\ell + v_\mu^\ell \hat{\zeta}_\mu^\ell]\right) \right\rangle_{\{r_\mu^\ell, v_\mu^\ell\} \sim \mathcal{N}(0, \mathbf{G}_+^{\ell+1})}, \quad \mathbf{G}_+^{\ell+1} = \begin{bmatrix} \mathbf{G}^{\ell+1} & \tilde{\mathbf{G}}^{\ell+1} \\ \tilde{\mathbf{G}}^{\ell+1} & \tilde{\tilde{\mathbf{G}}}^{\ell+1} \end{bmatrix}.
\end{aligned} \tag{36}$$

We thus see that the Gaussian sources $\{r_\mu^\ell\}_\mu$ and $\{v_\mu^\ell\}_\mu$ are mean zero with correlation given by $\Sigma^{\ell+1}$. Now that we have linearized the quadratic components involving each of the dual fields

$\{\hat{\chi}, \hat{\xi}, \hat{\zeta}\}$, we now perform integration over these variables, giving

$$\begin{aligned}
& \int \prod_{\mu} \frac{d\hat{\chi}_{\mu}^{\ell}}{2\pi} \exp \left(i \sum_{\mu} \hat{\chi}_{\mu}^{\ell} \left(\chi_{\mu}^{\ell} - u_{\mu}^{\ell} - \sum_{\nu} A_{\mu\nu}^{\ell-1} g_{\nu}^{\ell} - \sum_{\nu} C_{\mu\nu}^{\ell-1} \tilde{g}_{\nu}^{\ell} \right) \right) \\
&= \prod_{\mu} \delta \left(\chi_{\mu}^{\ell} - u_{\mu}^{\ell} - \sum_{\nu} A_{\mu\nu}^{\ell-1} g_{\nu}^{\ell} - \sum_{\nu} C_{\mu\nu}^{\ell-1} \tilde{g}_{\nu}^{\ell} \right) \\
& \int \prod_{\mu} \frac{d\hat{\xi}_{\mu}^{\ell}}{2\pi} \exp \left(i \sum_{\mu} \hat{\xi}_{\mu}^{\ell} \left(\xi_{\mu}^{\ell} - r_{\mu}^{\ell} - \sum_{\nu} B_{\nu\mu}^{\ell} \phi(h_{\nu}^{\ell}) \right) \right) = \prod_{\mu} \delta \left(\xi_{\mu}^{\ell} - r_{\mu}^{\ell} - \sum_{\nu} B_{\nu\mu}^{\ell} \phi(h_{\nu}^{\ell}) \right) \\
& \int \prod_{\mu} \frac{d\hat{\zeta}_{\mu}^{\ell}}{2\pi} \exp \left(i \sum_{\mu} \hat{\zeta}_{\mu}^{\ell} \left(\zeta_{\mu}^{\ell} - v_{\mu}^{\ell} - \sum_{\nu} D_{\nu\mu}^{\ell} \phi(h_{\nu}^{\ell}) \right) \right) = \prod_{\mu} \delta \left(\zeta_{\mu}^{\ell} - v_{\mu}^{\ell} - \sum_{\nu} D_{\nu\mu}^{\ell} \phi(h_{\nu}^{\ell}) \right).
\end{aligned} \tag{37}$$

This reveals the following set of identities

$$\begin{aligned}
\chi_{\mu}^{\ell} &= u_{\mu}^{\ell} + \sum_{\nu} A_{\mu\nu}^{\ell-1} g_{\nu}^{\ell} + \sum_{\nu} C_{\mu\nu}^{\ell-1} \tilde{g}_{\nu}^{\ell} \\
\xi_{\mu}^{\ell} &= r_{\mu}^{\ell} + \sum_{\nu} B_{\nu\mu}^{\ell} \phi(h_{\nu}^{\ell}), \quad \zeta_{\mu}^{\ell} = v_{\mu}^{\ell} + \sum_{\nu} D_{\nu\mu}^{\ell-1} \phi(h_{\nu}^{\ell}).
\end{aligned} \tag{38}$$

Since we know by construction that $\mathbf{u}^{\ell} = \boldsymbol{\chi}^{\ell} - \mathbf{A}^{\ell-1} \mathbf{g}^{\ell} - \mathbf{C}^{\ell-1} \tilde{\mathbf{g}}^{\ell}$ is a zero mean Gaussian with covariance $\boldsymbol{\Phi}^{\ell-1}$, we can simplify our expressions for $\mathbf{B}^{\ell-1}$ and $\hat{\boldsymbol{\Phi}}^{\ell-1}$ using Stein's Lemma

$$\begin{aligned}
\mathbf{B}^{\ell-1} &= \frac{1}{N} \sum_{i=1}^N [\boldsymbol{\Phi}^{\ell-1}]^{-1} \langle \mathbf{u}^{\ell} \mathbf{g}^{\ell\top} \rangle_i = \frac{1}{N} \sum_{i=1}^N \left\langle \frac{\partial \mathbf{g}^{\ell\top}}{\partial \mathbf{u}^{\ell}} \right\rangle_i \\
\mathbf{D}^{\ell-1} &= \frac{1}{N} \sum_{i=1}^N [\boldsymbol{\Phi}^{\ell-1}]^{-1} \langle \mathbf{u}^{\ell} \tilde{\mathbf{g}}^{\ell\top} \rangle_i = \frac{1}{N} \sum_{i=1}^N \left\langle \frac{\partial \tilde{\mathbf{g}}^{\ell\top}}{\partial \mathbf{u}^{\ell}} \right\rangle_i \\
\hat{\boldsymbol{\Phi}}^{\ell-1} &= \frac{1}{2} [\boldsymbol{\Phi}^{\ell-1}]^{-1} - \frac{1}{2N} \sum_{i=1}^N [\boldsymbol{\Phi}^{\ell-1}]^{-1} \langle \mathbf{u}^{\ell} \mathbf{u}^{\ell\top} \rangle_i [\boldsymbol{\Phi}^{\ell-1}]^{-1} = 0.
\end{aligned} \tag{39}$$

Similarly, using the Gaussianity of $\mathbf{r}^{\ell}, \boldsymbol{\zeta}^{\ell}, \hat{\boldsymbol{\zeta}}^{\ell}$ we have

$$\mathbf{A}^{\ell} = \frac{1}{N} \sum_{i=1}^N \left\langle \frac{\partial \phi(\mathbf{h}^{\ell})}{\partial \mathbf{r}^{\ell\top}} \right\rangle_i, \quad \mathbf{C}^{\ell} = \frac{1}{N} \sum_{i=1}^N \left\langle \frac{\partial \phi(\mathbf{h}^{\ell})}{\partial \mathbf{v}^{\ell\top}} \right\rangle_i, \quad \hat{\mathbf{G}}^{\ell+1} = \hat{\tilde{\mathbf{G}}}^{\ell+1} = \hat{\tilde{\tilde{\mathbf{G}}}}^{\ell+1} = 0.$$

B.7 FINAL DMFT EQUATIONS

We now take the limit of zero source $\mathbf{j}^{\ell}, \mathbf{k}^{\ell}, \mathbf{n}^{\ell}, \mathbf{p}^{\ell} \rightarrow 0$. In this limit, all single site averages $\langle \cdot \rangle_i$ become identical so we can simplify the expressions for the order parameters. To ‘‘symmetrize’’ the equations we will also make the substitution $\mathbf{B} \rightarrow \mathbf{B}^{\top}, \mathbf{D} \rightarrow \mathbf{D}^{\top}$. Next, we also rescale all of the response functions $\{A^{\ell}, B^{\ell}, C^{\ell}, D^{\ell}\}$ by γ_0^{-1} so that they are $O_{\gamma_0}(1)$ at small γ_0 . This gives us the following set of equations for the order parameters

$$\begin{aligned}
\Phi_{\mu\nu}^{\ell}(t, s) &= \langle \phi(h_{\mu}^{\ell}(t)) \phi(h_{\nu}^{\ell}(s)) \rangle, \quad G_{\mu\nu}^{\ell}(t, s) = \langle g_{\mu}^{\ell}(t) g_{\nu}^{\ell}(s) \rangle, \quad \tilde{G}_{\mu\nu}^{\ell}(t, s) = \langle g_{\mu}^{\ell}(t) \tilde{g}_{\nu}^{\ell}(s) \rangle \\
\tilde{\tilde{G}}_{\mu\nu}^{\ell}(t, s) &= \langle \tilde{g}_{\mu}^{\ell}(t) \tilde{g}_{\nu}^{\ell}(s) \rangle, \quad A_{\mu\nu}^{\ell}(t, s) = \gamma_0^{-1} \left\langle \frac{\delta \phi(h_{\mu}^{\ell}(t))}{\delta r_{\nu}^{\ell}(s)} \right\rangle, \quad C_{\mu\nu}^{\ell}(t, s) = \gamma_0^{-1} \left\langle \frac{\delta \phi(h_{\mu}^{\ell}(t))}{\delta v_{\nu}^{\ell}(s)} \right\rangle \\
B_{\mu\nu}^{\ell}(t, s) &= \gamma_0^{-1} \left\langle \frac{\delta g_{\mu}^{\ell+1}(t)}{\delta u_{\nu}^{\ell+1}(s)} \right\rangle, \quad D_{\mu\nu}^{\ell}(t, s) = \gamma_0^{-1} \left\langle \frac{\delta \tilde{g}_{\mu}^{\ell+1}(t)}{\delta u_{\nu}^{\ell+1}(s)} \right\rangle.
\end{aligned}$$

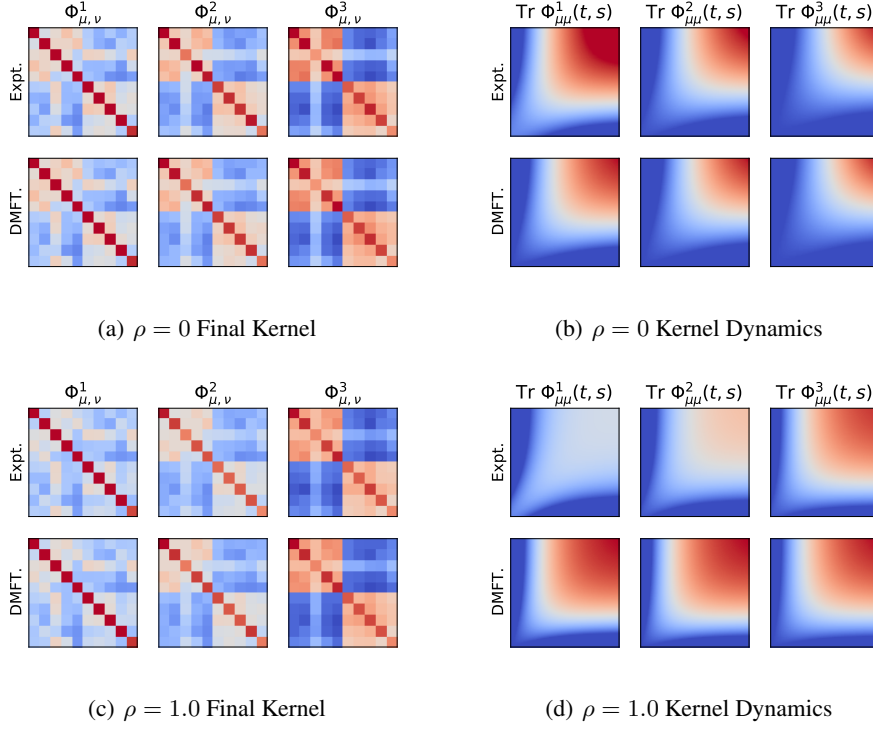


Figure 6: Feature kernels Φ^ℓ and their dynamics predicted by solving full set of saddle point equations equation 40 for ρ -FA in depth 3 tanh network with $\gamma_0 = 1.0$. Solving deep nonlinear ρ -FA requires sampling the full triplet of Gaussian sources $\{u^\ell, r^\ell, v^\ell\}$ for each layer and computing all four response functions $\{A^\ell, B^\ell, C^\ell, D^\ell\}$. DMFT theoretical predictions are compared to a width $N = 3000$ neural network.

For the fields $h_\mu^\ell(t)$, $z_\mu^\ell(t)$, $\tilde{z}_\mu^\ell(t)$, we have the following equations

$$\begin{aligned}
 h_\mu^\ell(t) &= u_\mu^\ell(t) + \gamma_0 \int_0^t ds \sum_{\nu=1}^P [A_{\mu\nu}^{\ell-1}(t, s) g_\nu^\ell(s) + C_{\mu\nu}^{\ell-1}(t, s) \tilde{g}_\mu^\ell(s) + \Delta_\nu(s) \Phi_{\mu\nu}^{\ell-1}(t, s) \tilde{g}_\nu^\ell(s)] \\
 z_\mu^\ell(t) &= r_\mu^\ell(t) + \gamma_0 \int_0^t ds \sum_{\nu=1}^P [B_{\mu\nu}^\ell(t, s) \phi(h_\nu^\ell(s)) + \Delta_\nu(s) \tilde{G}_{\mu\nu}^{\ell+1}(t, s) \phi(h_\nu^\ell(s))] \\
 \tilde{g}_\mu^\ell(t) &= \begin{cases} \dot{\phi}(h_\mu^\ell(t)) z_\mu^\ell(t) & \text{GD} \\ \dot{\phi}(h_\mu^\ell(t)) [\sqrt{1 - \rho^2} \tilde{\zeta}_\mu^\ell(t) + \rho v_\mu^\ell(t) + \rho \gamma_0 \int_0^t ds \sum_{\nu=1}^P D_{\mu\nu}^\ell(t, s) \phi(h_\nu^\ell(s))] & \rho\text{-FA} \\ \dot{\phi}(h_\mu^\ell(t)) \tilde{z}^\ell, \tilde{z}^\ell \sim \mathcal{N}(0, 1) & \text{DFA} \\ \dot{\phi}(m_\mu^\ell(t)) z_\mu^\ell(t) & \text{GLN} \\ \Delta_\mu(t) \phi(h_\mu^\ell(t)) & \text{Hebb} \end{cases} \\
 \{u_\mu^\ell(t)\} &\sim \mathcal{GP}(0, \Phi^{\ell-1}), \{r_\mu^\ell(t), v_\mu^\ell(t)\} \sim \mathcal{GP}(0, \mathbf{G}_+^{\ell+1}), \{\tilde{\zeta}_\mu^\ell(t)\} \sim \mathcal{GP}(0, \tilde{\mathbf{G}}^{\ell+1}) \\
 \mathbf{G}_+^{\ell+1} &= \begin{bmatrix} \mathbf{G}^{\ell+1} & \tilde{\mathbf{G}}^{\ell+1} \\ \tilde{\mathbf{G}}^{\ell+1, \top} & \tilde{\mathbf{G}}^{\ell+1} \end{bmatrix}. \tag{40}
 \end{aligned}$$

C EXTENSION TO OTHER ARCHITECTURES AND OPTIMIZERS

In this section, we consider the effect of changing architectural details (multiple output channels and convolutional structure) and also optimization choices (momentum, regularization).

C.1 MULTIPLE OUTPUT CLASSES

Similar to pre-existing work on the GD case (Bordelon & Pehlevan, 2022), our new generalized DMFT can be easily extended to C output channels, provided the number of channels C is not simultaneously taken to infinity with network width N . We note that the outputs of the network are now vectors $\mathbf{f}_\mu \in \mathbb{R}^C$ and that each eNTK entry is now a $C \times C$ matrix $\mathbf{K}_{\mu\nu}(t, s) \in \mathbb{R}^{C \times C}$. The relevant true gradient fields are vectors $\mathbf{g}_{c,\mu}^\ell = \frac{\partial f_{c,\mu}}{\partial \mathbf{h}_\mu^\ell}$. We construct pseudo-gradients $\tilde{\mathbf{g}}_{c,\mu}^\ell$ as before using each of our learning rules. The gradient-pseudogradient kernel $\mathbf{G}_{\mu\nu}^\ell \in \mathbb{R}^{C \times C}$ is $\tilde{\mathbf{G}}_{c,c',\mu\nu}^\ell = \frac{1}{N} \mathbf{g}_{c,\mu}^\ell \cdot \mathbf{g}_{c',\nu}^\ell$. The eNTK $\mathbf{K}_{\mu\nu} = \sum_\ell \tilde{\mathbf{G}}_{\mu\nu}^{\ell+1} \Phi_{\mu\nu}^\ell$ can be used to derive the function dynamics

$$\frac{\partial \mathbf{f}_\mu}{\partial t} = \sum_\nu \mathbf{K}_{\mu\nu} \Delta_\nu, \quad \Delta_\mu = -\frac{\partial \mathcal{L}}{\partial \mathbf{f}_\mu}. \quad (41)$$

At infinite width $N \rightarrow \infty$, the field dynamics for $h_\mu^\ell(t) \in \mathbb{R}$, $\mathbf{g}_\mu^\ell(t) \in \mathbb{R}^C$ satisfy

$$\begin{aligned} h_\mu^\ell(t) &= u_\mu^\ell(t) + \gamma_0 \int_0^t ds \sum_{\nu=1}^P [\tilde{\mathbf{g}}_\nu^\ell(s) \cdot \Delta_\nu(s) \Phi_{\mu\nu}^{\ell-1}(t, s) + \mathbf{A}_{\mu\nu}^{\ell-1}(t, s) \cdot \mathbf{g}_\nu^\ell(s) + \mathbf{C}_{\mu\nu}^{\ell-1} \cdot \tilde{\mathbf{g}}_\nu^\ell(s)] \\ \mathbf{z}_\mu^\ell(t) &= \mathbf{r}_\mu^\ell(t) + \gamma_0 \int_0^t ds \sum_{\nu=1}^P \phi(h_\nu^\ell(s)) [\tilde{\mathbf{G}}_{\mu\nu}^{\ell+1}(t, s) \Delta_\nu(s) + \mathbf{B}_{\mu\nu}^\ell(t, s)], \end{aligned} \quad (42)$$

where $\mathbf{A}_{\mu\nu}^\ell(t, s) = \gamma_0 \left\langle \frac{\delta \phi(h_\mu^\ell(t))}{\delta \mathbf{r}_\nu^\ell(s)} \right\rangle \in \mathbb{R}^C$, $\mathbf{A}_{\mu\nu}^\ell(t, s) = \gamma_0 \left\langle \frac{\delta \phi(h_\mu^\ell(t))}{\delta \mathbf{v}_\nu^\ell(s)} \right\rangle \in \mathbb{R}^C$, and $\mathbf{B}_{\mu\alpha}(t, s) = \frac{\partial \mathbf{g}_\mu^{\ell+1}(t)}{\partial u_\alpha^\ell(s)} \in \mathbb{R}^C$. The feature kernels are the same as before $\Phi_{\mu\nu}^\ell(t, s) = \langle \phi(h_\mu^\ell(t)) \phi(h_\nu^\ell(s)) \rangle$ while the gradient-pseudogradient kernel is $\tilde{\mathbf{G}}_{\mu\nu}^\ell(t, s) = \langle \mathbf{g}_\mu^\ell(t) \tilde{\mathbf{g}}_\nu^{\ell\top}(s) \rangle \in \mathbb{R}^{C \times C}$. The pseudogradient fields $\tilde{\mathbf{g}}^\ell$ are defined analogously for each learning rule as in the single class setting.

$$\tilde{\mathbf{g}}_\mu^\ell(t) = \begin{cases} \dot{\phi}(h_\mu^\ell(t)) \mathbf{z}_\mu^\ell(t) & \text{GD} \\ \dot{\phi}(h_\mu^\ell(t)) \left[\rho \mathbf{v}_\mu^\ell(t) + \sqrt{1 - \rho^2} \tilde{\zeta}_\mu^\ell(t) + \rho \gamma_0 \int_0^t ds \sum_{\nu=1}^P \mathbf{D}_\nu^\ell(t, s) \phi(h_\nu^\ell(s)) \right] & \text{\(\rho\)-FA} \\ \dot{\phi}(h_\mu^\ell(t)) \tilde{\mathbf{z}}^\ell, \quad \tilde{\mathbf{z}}^\ell \sim \mathcal{N}(0, \mathbf{I}) & \text{DFA} \\ \dot{\phi}(m_\mu^\ell(t)) \mathbf{z}_\mu^\ell(t) & \text{GLN} \\ \mathbf{1} \Delta_\mu(t) \phi(h_\mu^\ell(t)) & \text{Hebb} \end{cases} \quad (43)$$

C.2 CNN

The DMFT described for each of these learning rules can also be extended to CNNs with infinitely many channels. Following the work of Bordelon & Pehlevan (2022) Appendix G on the GD DMFT limit for CNNs, we let $W_{ij,a}^\ell$ represent the value of the filter at spatial displacement \mathbf{a} from the center of the filter, which maps relates activity at channel j of layer ℓ to channel i of layer $\ell + 1$. The fields $h_{\mu,i,a}^\ell$ satisfy the recursion

$$h_{\mu,i,a}^{\ell+1} = \frac{1}{\sqrt{N}} \sum_{j=1}^N \sum_{\mathbf{b} \in \mathcal{S}^\ell} W_{ij,b}^\ell \phi(h_{\mu,j,a+\mathbf{b}}^\ell), \quad i \in [N], \quad (44)$$

where \mathcal{S}^ℓ is the spatial receptive field at layer ℓ . For example, a $(2k+1) \times (2k+1)$ convolution will have $\mathcal{S}^\ell = \{(i, j) \in \mathbb{Z}^2 : -k \leq i \leq k, -k \leq j \leq k\}$. The output function is obtained from the last layer is defined as $f_\mu = \frac{1}{\gamma_0 N} \sum_{i=1}^N \sum_{\mathbf{a}} w_{i,a}^L \phi(h_{\mu,i,a}^L)$. The true gradient fields have the same definition as before $\mathbf{g}_{\mu,a}^\ell = \gamma_0 N \frac{\partial f_\mu}{\partial \mathbf{h}_{\mu,a}^\ell} \in \mathbb{R}^N$, which as before enjoy the following recursion

$$\mathbf{g}_{\mu,a}^\ell = \gamma_0 N \sum_{\mathbf{b}} \frac{\partial f_\mu}{\partial \mathbf{h}_{\mu,b}^{\ell+1}} \cdot \frac{\partial \mathbf{h}_{\mu,b}^{\ell+1}}{\partial \mathbf{h}_{\mu,a}^\ell} = \dot{\phi}(h_{\mu,a}^\ell) \odot \left[\frac{1}{\sqrt{N}} \sum_{j=1}^N \sum_{\mathbf{b} \in \mathcal{S}^\ell} \mathbf{W}_b^{\ell\top} \mathbf{g}_{\mu,a-\mathbf{b}}^{\ell+1} \right]. \quad (45)$$

We consider the following learning dynamics for the filters

$$\frac{d}{dt} \mathbf{W}_b^\ell = \frac{\gamma_0}{\sqrt{N}} \sum_{\mu,c} \Delta_\mu \tilde{\mathbf{g}}_{\mu,c}^{\ell+1} \phi(\mathbf{h}_{\mu,c+b}^\ell)^\top \quad (46)$$

where as before $\tilde{\mathbf{g}}^\ell$ is determined by the learning rule. The relevant kernel order parameters now have spatial indices. For instance the feature kernel at each layer has form $\Phi_{\mu,\nu,ab}^\ell = \frac{1}{N} \phi(\mathbf{h}_{\mu,a}^\ell(t)) \cdot \phi(\mathbf{h}_{\nu,b}^\ell(s))$. At the infinite width $N \rightarrow \infty$, the order parameters and field dynamics have the form

$$\begin{aligned} h_{\mu,a}^\ell(t) &= u_{\mu,a}^\ell(t) + \gamma_0 \int_0^t ds \sum_{\nu,b,c} \Delta_\nu(s) \Phi_{\mu\nu,a+b,b+c}^{\ell-1}(t,s) \tilde{g}_{\nu,c}^\ell(s) \\ &+ \gamma_0 \int_0^t ds \sum_{\nu,b} [A_{\mu\nu,ab}^{\ell-1}(t,s) g_{\nu,b}^\ell(s) + C_{\mu\nu,ab}^{\ell-1}(t,s) \tilde{g}_{\nu,b}^\ell(s)] \\ z_{\mu,a}^\ell(t) &= r_{\mu,a}^\ell(t) + \gamma_0 \int_0^t ds \sum_{\nu,b,c} \tilde{G}_{\mu\nu,a-b,c-b}^{\ell+1}(t,s) \phi(h_{\nu,c}^\ell(s)) \\ &+ \gamma_0 \int_0^t ds \sum_{\nu,b} B_{\mu\nu,ab}^\ell(t,s) \phi(h_{\nu,b}^\ell(s)) \end{aligned} \quad (47)$$

$$(48)$$

where correlation and response functions have the usual definitions

$$\begin{aligned} \Phi_{\mu\alpha,ab}^\ell(t,s) &= \langle \phi(h_{\mu a}^\ell(t)) \phi(h_{\alpha b}^\ell(s)) \rangle, \quad G_{\mu\alpha,ab}^\ell(t,s) = \langle g_{\mu a}^\ell(t) g_{\alpha b}^\ell(s) \rangle, \quad \tilde{G}_{\mu\nu,ab}^\ell(t,s) = \langle g_{\mu a}^\ell(t) \tilde{g}_{\alpha b}^\ell(s) \rangle \\ A_{\mu\alpha,ab}^\ell(t,s) &= \frac{1}{\gamma_0} \left\langle \frac{\delta \phi(h_{\mu a}^\ell(t))}{\delta r_{\alpha b}^\ell(s)} \right\rangle, \quad B_{\mu\alpha,ab}^\ell(t,s) = \frac{1}{\gamma_0} \left\langle \frac{\delta g_{\mu a}^{\ell+1}(t)}{\delta u_{\alpha b}^{\ell+1}(s)} \right\rangle. \end{aligned} \quad (49)$$

C.3 L2 REGULARIZATION (WEIGHT DECAY)

L2 regularization on the weights \mathbf{W}^ℓ (weight decay) can also be modeled within DMFT. We start by looking at the weight dynamics

$$\begin{aligned} \frac{d}{dt} \mathbf{W}^\ell &= \frac{\gamma_0}{\sqrt{N}} \sum_{\mu=1}^P \Delta_\mu \tilde{\mathbf{g}}_\mu^{\ell+1} \phi(\mathbf{h}_\mu^\ell)^\top - \lambda \mathbf{W}^\ell \\ \implies \mathbf{W}^\ell(t) &= e^{-\lambda t} \mathbf{W}^\ell(0) + \frac{\gamma_0}{\sqrt{N}} \int_0^t ds e^{-\lambda(t-s)} \sum_{\mu} \Delta_\mu(s) \tilde{\mathbf{g}}_\mu^{\ell+1}(s) \phi(\mathbf{h}_\mu^\ell(s))^\top \end{aligned} \quad (50)$$

In the second line we used an integrating factor $e^{\lambda t}$. We can thus arrive at the following feature dynamics in the DMFT limit

$$\begin{aligned} h_\mu^\ell(t) &= e^{-\lambda t} \chi_\mu^\ell(t) + \gamma_0 \int_0^t ds e^{-\lambda(t-s)} \sum_{\nu} \Delta_\nu(s) \Phi_{\mu\nu}^{\ell-1}(s) \tilde{g}_\nu^\ell(s) \\ z_\mu^\ell(t) &= e^{-\lambda t} \xi_\mu^\ell(t) + \gamma_0 \int_0^t ds e^{-\lambda(t-s)} \sum_{\nu} \Delta_\nu(s) \tilde{G}_{\mu\nu}^{\ell+1}(s) \phi(h_\nu^\ell(s)). \end{aligned}$$

The \tilde{g} dynamics are also modified appropriately with factors of $e^{-\lambda t}$ and $e^{-\lambda(t-s)}$ for each of our learning rules. We see that the contribution from the initial conditions χ, ξ are suppressed at late times while the feature learning update which is $O(\gamma_0/\lambda)$ in the first layer dominates scale of the final features.

C.4 MOMENTUM

Momentum uses a low-pass filtered version of the gradients to update the weights (Goh, 2017). A continuous time limit of momentum dynamics on the trainable parameters $\{\mathbf{W}^\ell\}$ would give the

following differential equations

$$\begin{aligned}\frac{\partial}{\partial t} \mathbf{W}^\ell(t) &= \mathbf{Q}^\ell(t) \\ \tau \frac{d}{dt} \mathbf{Q}^\ell(t) &= -\mathbf{Q}^\ell + \frac{\gamma_0}{\sqrt{N}} \sum_{\mu} \Delta_{\mu}(t) \tilde{\mathbf{g}}_{\mu}^{\ell+1}(t) \phi(\mathbf{h}_{\mu}^{\ell}(t))^{\top}.\end{aligned}\quad (51)$$

We write the expression this way so that the small time constant $\tau \rightarrow 0$ limit corresponds to classic gradient descent. Integration of the $\mathbf{Q}^\ell(t)$ dynamics gives the following integral expression for \mathbf{W}^ℓ

$$\mathbf{W}^\ell(t) = \mathbf{W}^\ell(0) + \frac{\gamma_0}{\sqrt{N}\tau} \int_0^t dt' \int_0^{t'} dt'' e^{-(t'-t'')/\tau} \sum_{\mu} \Delta_{\mu}(t'') \tilde{\mathbf{g}}_{\mu}^{\ell+1}(t'') \phi(\mathbf{h}_{\mu}^{\ell}(t''))^{\top}.\quad (52)$$

These weight dynamics give rise to the following field evolution

$$\begin{aligned}h_{\mu}^{\ell+1}(t) &= \chi_{\mu}^{\ell+1}(t) + \frac{\gamma_0}{\tau} \int_0^t dt' \int_0^{t'} dt'' e^{-(t'-t'')/\tau} \sum_{\nu} \Delta_{\nu}(t'') \tilde{\mathbf{g}}_{\nu}^{\ell+1}(t'') \Phi_{\mu\nu}^{\ell}(t, t'') \\ z_{\mu}^{\ell}(t) &= \xi_{\mu}^{\ell}(t) + \frac{\gamma_0}{\tau} \int_0^t dt' \int_0^{t'} dt'' e^{-(t'-t'')/\tau} \sum_{\nu} dt'' \Delta_{\alpha}(t'') \tilde{G}_{\mu\nu}^{\ell+1}(t, t'') \phi(h_{\nu}^{\ell}(t'')).\end{aligned}\quad (53)$$

We see that in the $\tau \rightarrow 0$ limit, the t'' integral is dominated by the contribution at $t'' \sim t'$ recovering usual gradient descent dynamics. For $\tau \gg 0$, we see that the integral accumulates additional contributions from the past values of fields and kernels.

D LAZY LIMITS

In this section we discuss the lazy $\gamma_0 \rightarrow 0$ limit. In this limit we see that $h^{\ell}(t) = u^{\ell}(t)$ and $z^{\ell}(t) = r^{\ell}(t)$ for all time t . Since the input data gram matrix $\Phi_{\mu\nu}^0 = \frac{1}{D} \mathbf{x}_{\mu} \cdot \mathbf{x}_{\nu}$ is a constant in time the sources in the first hidden layer u_{μ}^1 are constant in time. Consequently, the first layer feature kernel is constant in time since

$$\Phi_{\mu\nu}^1(t, s) = \langle \phi(h_{\mu}^1(t)) \phi(h_{\nu}^1(s)) \rangle = \langle \phi(u_{\mu}^1) \phi(u_{\nu}^1) \rangle_{\mathbf{u}^1 \sim \mathcal{N}(0, \Phi^0)}.\quad (54)$$

Now, we see that this argument can proceed inductively. Since Φ^1 is time-independent, the second layer fields $\mathbf{h}^2 = \mathbf{u}^2 \sim \mathcal{N}(0, \Phi^1)$ are also constant in time, implying Φ^2 is constant in time. This argument is repeated for all layer $\ell \in [L]$. Similarly, we can analyze the backward pass fields \mathbf{z}^{ℓ} . Since $\mathbf{z}^L \sim \mathcal{N}(0, \mathbf{G}^{L+1})$ are constant, then \mathbf{z}^{ℓ} are time-independent for all ℓ . It thus suffices to compute the static kernels $\{\Phi^{\ell}, \mathbf{G}^{\ell}, \tilde{\mathbf{G}}^{\ell}\}$ at initialization

$$\begin{aligned}\Phi^{\ell} &= \langle \phi(\mathbf{u}^{\ell}) \phi(\mathbf{u}^{\ell})^{\top} \rangle_{\mathbf{u}^{\ell} \sim \mathcal{N}(0, \Phi^{\ell-1})} \\ \mathbf{G}^{\ell} &= \left\langle [\dot{\phi}(\mathbf{u}^{\ell}) \odot \mathbf{r}^{\ell}] [\dot{\phi}(\mathbf{u}^{\ell}) \odot \mathbf{r}^{\ell}]^{\top} \right\rangle_{\mathbf{u}^{\ell} \sim \mathcal{N}(0, \Phi^{\ell-1}), \mathbf{r}^{\ell} \sim \mathcal{N}(0, \mathbf{G}^{\ell+1})} \\ &= \mathbf{G}^{\ell+1} \odot \dot{\Phi}^{\ell}, \quad \dot{\Phi}^{\ell} = \left\langle \dot{\phi}(\mathbf{u}^{\ell}) \dot{\phi}(\mathbf{u}^{\ell})^{\top} \right\rangle_{\mathbf{u}^{\ell} \sim \mathcal{N}(0, \Phi^{\ell-1})}.\end{aligned}\quad (55)$$

where in the last line we utilized the independence of $\mathbf{u}^{\ell}, \mathbf{r}^{\ell}$. These above equations give a forward pass recursion for the Φ^{ℓ} kernels and the backward pass recursion for \mathbf{G}^{ℓ} . Lastly, depending on the learning rule, we arrive at the following definitions for $\tilde{\mathbf{G}}^{\ell}$ for $\ell \in \{1, \dots, L\}$

$$\tilde{\mathbf{G}}^{\ell} = \left\langle [\dot{\phi}(\mathbf{u}^{\ell}) \odot \mathbf{r}^{\ell}] \tilde{\mathbf{g}}^{\ell \top} \right\rangle = \begin{cases} \mathbf{G}^{\ell+1} \odot \dot{\Phi}^{\ell} & \text{GD} \\ \rho \tilde{\mathbf{G}}^{\ell+1} \odot \dot{\Phi}^{\ell} & \rho\text{-FA} \\ 0 & \text{DFA, Hebb} \\ \mathbf{G}^{\ell+1} \odot \left\langle \dot{\phi}(\mathbf{m}^{\ell}) \dot{\phi}(\mathbf{m}^{\ell})^{\top} \right\rangle & \text{GLN} \end{cases}\quad (56)$$

Using these results for $\tilde{\mathbf{G}}$, we can compute the initial eNTK $\mathbf{K} = \sum_{\ell=0}^L \tilde{\mathbf{G}}^{\ell+1} \odot \Phi^{\ell}$ which governs prediction dynamics.

D.1 LAZY LIMIT PERFORMANCES ON REALISTIC TASKS

We note that, while the DMFT equations on P datapoints and T timesteps require $O(P^3T^3)$ time complexity to solve in the rich regime, the lazy limit gives neural network predictions in $O(P^3)$ time, since the predictor can be obtained by solving a linear system of P equations. The performance of these lazy limit kernels on realistic tasks would match the performances reported by Lee et al. (2020). Specifically, GD and $\rho = 1$ FA would match the test accuracy reported for “infinite width GD”, while $\rho = 0$ FA, DFA, and Hebbian rules would match “infinite width Bayesian” networks in Figure 1 of Lee et al. (2020).

E DEEP LINEAR NETWORKS

In deep linear networks, the DMFT equations close without needing any numerical sampling procedure, as was shown in prior work on the GD case (Yang & Hu, 2021; Bordelon & Pehlevan, 2022). The key observation is that for all of the following learning rules, the fields $\{\mathbf{h}, \mathbf{g}, \tilde{\mathbf{g}}\}$ are linear combinations of the Gaussian sources $\{u, r, v\}$, and are thus Gaussian themselves. Concretely, we introduce a vector notation $\mathbf{h}^\ell = \text{Vec}\{h_\mu^\ell(t)\}$ and $\mathbf{g}^\ell = \text{Vec}\{g_\mu^\ell(t)\}$, etc. We have in each layer

$$\begin{aligned}\mathbf{h}^\ell &= \mathbf{R}_{h,u}\mathbf{u}^\ell + \mathbf{R}_{h,r}\mathbf{r}^\ell + \mathbf{R}_{h,v}\mathbf{v}^\ell + \mathbf{R}_{h,\zeta}\tilde{\boldsymbol{\zeta}}^\ell \\ \mathbf{g}^\ell &= \mathbf{R}_{g,u}\mathbf{u}^\ell + \mathbf{R}_{g,r}\mathbf{r}^\ell + \mathbf{R}_{g,v}\mathbf{v}^\ell + \mathbf{R}_{g,\zeta}\tilde{\boldsymbol{\zeta}}^\ell \\ \tilde{\mathbf{g}}^\ell &= \mathbf{R}_{\tilde{g},u}\mathbf{u}^\ell + \mathbf{R}_{\tilde{g},r}\mathbf{r}^\ell + \mathbf{R}_{\tilde{g},v}\mathbf{v}^\ell + \mathbf{R}_{\tilde{g},\zeta}\tilde{\boldsymbol{\zeta}}^\ell\end{aligned}$$

where the matrices \mathbf{R} depend on the learning rule and the data. The necessary kernels $\mathbf{H}^\ell = \langle \mathbf{h}^\ell \mathbf{h}^{\ell\top} \rangle$ can thus be closed algebraically since all of the correlation statistics of the sources $\{u, r, v\}$ have known two-point correlation statistics.

E.1 LINEAR NETWORK TRAINED WITH GD

The \mathbf{R} matrices for GD were provided in (Bordelon & Pehlevan, 2022). We start by noting the following DMFT equations for $\mathbf{h}^\ell, \mathbf{g}^\ell$

$$\mathbf{h}^\ell = \mathbf{u}^\ell + \gamma_0(\mathbf{A}^{\ell-1} + \mathbf{H}_\Delta^{\ell-1})\mathbf{g}^\ell, \quad \mathbf{g}^\ell = \mathbf{r}^\ell + \gamma_0(\mathbf{B}^\ell + \mathbf{G}_\Delta^{\ell+1})\mathbf{h}^\ell \quad (57)$$

where $[\mathbf{H}_\Delta^{\ell-1}]_{\mu\nu,ts} = H_{\mu\nu}^\ell(t,s)\Delta_\nu(s)$. Isolating the dependence of these equations on \mathbf{u} and \mathbf{r} , we have

$$\begin{aligned}[\mathbf{I} - \gamma_0^2(\mathbf{A}^{\ell-1} + \mathbf{H}_\Delta^{\ell-1})(\mathbf{B}^\ell + \mathbf{G}_\Delta^{\ell+1})]\mathbf{h}^\ell &= \mathbf{u}^\ell + \gamma_0^2(\mathbf{A}^{\ell-1} + \mathbf{H}_\Delta^{\ell-1})(\mathbf{B}^\ell + \mathbf{G}_\Delta^{\ell+1})\mathbf{r}^\ell \\ [\mathbf{I} - \gamma_0^2(\mathbf{B}^\ell + \mathbf{G}_\Delta^{\ell+1})(\mathbf{A}^{\ell-1} + \mathbf{H}_\Delta^{\ell-1})]\mathbf{g}^\ell &= \mathbf{r}^\ell + \gamma_0^2(\mathbf{B}^\ell + \mathbf{G}_\Delta^{\ell+1})(\mathbf{A}^{\ell-1} + \mathbf{H}_\Delta^{\ell-1})\mathbf{r}^\ell.\end{aligned} \quad (58)$$

These equations can easily be closed for \mathbf{H}^ℓ and \mathbf{G}^ℓ .

E.2 ρ -ALIGNED FEEDBACK ALIGNMENT

In ρ -FA we define the following pseudo-gradient fields

$$\tilde{\mathbf{g}}^\ell = \sqrt{1 - \rho^2}\tilde{\boldsymbol{\zeta}}^\ell + \rho\mathbf{v}^\ell + \rho\gamma_0\mathbf{D}^\ell\mathbf{h}^\ell \quad (59)$$

Next, we note that, at initialization, the $\tilde{\mathbf{G}}^\ell$ can be computed recursively

$$\tilde{\mathbf{G}}^\ell = \rho\tilde{\mathbf{G}}^{\ell+1} \quad (60)$$

We note that $\frac{\partial}{\partial r^1}\mathbf{h}^1 = 0$ which implies $\mathbf{A}^1 = 0$. Similarly we have $\frac{\partial}{\partial r^2}\mathbf{h}^2 = 0$. Thus $\mathbf{A}^2 = 0$. Proceeding inductively, we find $\mathbf{A}^\ell = 0$. Similarly, we note that $\frac{\partial \tilde{\mathbf{g}}^L}{\partial \mathbf{u}^L} = 0$ so $\mathbf{D}^{L-1} = 0$. Inductively, we have $\mathbf{D}^\ell = 0$ for all ℓ . Using these facts, we thus find the following equations

$$\mathbf{h}^\ell = \mathbf{u}^\ell + \gamma_0(\mathbf{C}^{\ell-1} + \mathbf{H}_\Delta^{\ell-1})\tilde{\mathbf{g}}^\ell \quad (61)$$

$$\mathbf{g}^\ell = \mathbf{r}^\ell + \gamma_0(\mathbf{B}^\ell + \mathbf{G}_\Delta^{\ell+1})\mathbf{h}^\ell \quad (62)$$

$$\tilde{\mathbf{g}}^\ell = \sqrt{1 - \rho^2}\tilde{\boldsymbol{\zeta}}^\ell + \rho\mathbf{v}^\ell \quad (63)$$

We can close these equations for \mathbf{H}^ℓ and $\tilde{\mathbf{G}}^\ell$

$$\begin{aligned}\mathbf{H}^\ell &= \mathbf{H}^{\ell-1} + \gamma_0^2 (\mathbf{C}^{\ell-1} + \mathbf{H}_{\Delta}^{\ell-1}) \tilde{\mathbf{G}}^{\ell+1} (\mathbf{C}^{\ell-1} + \mathbf{H}_{\Delta}^{\ell-1})^\top \\ \tilde{\mathbf{G}}^\ell &= \rho \tilde{\mathbf{G}}^{\ell+1} + \gamma_0^2 (\mathbf{B}^\ell + \mathbf{G}_{\Delta}^{\ell+1}) (\mathbf{C}^{\ell-1} + \mathbf{H}_{\Delta}^{\ell-1}) \tilde{\mathbf{G}}^{\ell+1} \\ \mathbf{C}^\ell &= \rho (\mathbf{C}^{\ell-1} + \mathbf{H}_{\Delta}^{\ell-1}), \quad \mathbf{B}^\ell = \mathbf{B}^{\ell+1} + \mathbf{G}_{\Delta}^{\ell+2}.\end{aligned}\quad (64)$$

The matrices $\tilde{\mathbf{G}}^\ell = \mathbf{1}\mathbf{1}^\top$ are all rank one. Thus it suffices to compute the vectors $\mathbf{c}^\ell = (\mathbf{C}^{\ell-1} + \mathbf{H}_{\Delta}^{\ell-1}) \mathbf{1}$. Further, it suffices to consider $\mathbf{d}^\ell = \tilde{\mathbf{G}}^\ell \mathbf{1} / |\mathbf{1}|^2$. With this formalism we have

$$\mathbf{H}^\ell = \mathbf{H}^{\ell-1} + \gamma_0^2 \mathbf{c}^\ell \mathbf{c}^{\ell\top}, \quad \mathbf{d}^\ell = \rho \mathbf{d}^{\ell+1} + \gamma_0^2 (\mathbf{B}^\ell + \mathbf{G}_{\Delta}^{\ell+1}) \mathbf{c}^\ell. \quad (65)$$

The analysis for DFA and Hebb rules is very similar.

F EXACTLY SOLVEABLE 2 LAYER LINEAR MODEL

F.1 GRADIENT FLOW

Based on the prior results from (Bordelon & Pehlevan, 2022), the $H_y = \mathbf{y}^\top \mathbf{H} \mathbf{y} / |\mathbf{y}|^2$ dynamics for GD are coupled to the dynamics for the error $\Delta(t) = \frac{1}{|\mathbf{y}|} \mathbf{y} \cdot \mathbf{\Delta}(t)$ have the form

$$\frac{d}{dt} H_y(t) = 2\gamma_0^2 (y - \Delta) \Delta, \quad \frac{d}{dt} \Delta = -2H_y \Delta. \quad (66)$$

These dynamics have the conservation law $\frac{d}{dt} H_y^2 = \gamma_0^2 \frac{d}{dt} (y - \Delta)^2$. Integrating this conservation law from time 0 to time t , we find $H_y(t)^2 = 1 + \gamma_0^2 (y - \Delta(t))^2$. We can therefore solve a single ODE for $\Delta(t)$, giving the following simplified dynamics

$$\frac{d}{dt} \Delta = -2\sqrt{1 + \gamma_0^2 (y - \Delta)^2} \Delta, \quad H_y = \sqrt{1 + \gamma_0^2 (y - \Delta)^2}. \quad (67)$$

These dynamics interpolate between exponential convergence (at small γ_0) and a logistic convergence (at large γ_0) of $\Delta(t)$ to zero. Since $\Delta \rightarrow 0$ at late time, the final value of the kernel alignment is $H_y = \sqrt{1 + \gamma_0^2 y^2}$.

F.2 ρ -ALIGNED FA

For the two layer linear network, the ρ -FA field dynamics are

$$\frac{d}{dt} h_\mu(t) = \gamma_0 \sum_\nu \tilde{g}_\nu(t) \Delta_\nu(s) K_{\mu\nu}^x, \quad \frac{d}{dt} g_\mu(t) = \gamma_0 \sum_\nu \Delta_\nu(t) h_\nu(t). \quad (68)$$

FA we have $\tilde{g}_\mu(t) = \tilde{g} \sim \mathcal{N}(0, 1)$ which is a constant standard normal. We let $a_\mu(t) = \langle \tilde{g} h_\mu(t) \rangle$. The dynamics for $H_{\mu\nu}$ and a_μ are coupled

$$\begin{aligned}\frac{d}{dt} H_{\mu\nu} &= \gamma_0 a_\nu(t) \sum_\nu \Delta_\nu(t) K_{\mu\nu}^x + \gamma_0 a_\mu(t) \sum_\nu \Delta_\nu(t) K_{\mu\nu}^x \\ \frac{d}{dt} a_\mu(t) &= \gamma_0 \sum_\nu \Delta_\nu K_{\mu\nu}^x \\ \frac{d}{dt} \tilde{G}(t) &= \gamma_0 \sum_\mu \Delta_\mu(t) a_\mu(t), \quad \frac{d}{dt} \Delta_\mu(t) = -\sum_\nu [H_{\mu\nu}(t) + \tilde{G}(t) K_{\mu\nu}^x] \Delta_\nu(t).\end{aligned}\quad (69)$$

Whitening the dataset $\mathbf{K}^x = \mathbf{I}$ and projecting all dynamics on $\hat{\mathbf{y}}$ subspace gives the reduced dynamics

$$\frac{d}{dt} H = 2\gamma_0 a \Delta, \quad \frac{d}{dt} a = \gamma_0 \Delta, \quad \frac{d}{dt} \tilde{G} = \gamma_0 \Delta a, \quad \frac{d}{dt} \Delta = -[H + \tilde{G}] \Delta. \quad (70)$$

From these dynamics we identify the following set of conservation laws

$$\begin{aligned} 2\frac{d}{dt}\tilde{G} &= \frac{d}{dt}a^2 = \frac{d}{dt}H \\ \implies 2\tilde{G} - 2\rho &= a^2 = H - 1. \end{aligned} \quad (71)$$

Writing everything in terms of Δ, a we have

$$\frac{d}{dt}a = \gamma_0\Delta, \quad \frac{d}{dt}\Delta = -\left[\frac{3}{2}a^2 + (1+\rho)\right]\Delta = -\gamma_0^{-1}\frac{d}{dt}\left[\frac{1}{2}a^3 + (1+\rho)a\right]$$

Integrating both sides of this equation from 0 to t gives $\Delta = y - \gamma_0^{-1}\left[\frac{1}{2}a^3 + (1+\rho)a\right]$. Thus, the a dynamics now one dimensional, giving

$$\frac{d}{dt}a = \gamma_0 y - \frac{1}{2}a^3 - (1+\rho)a. \quad (72)$$

When run from initial condition $a = 0$, this will converge to the smallest positive root of the cubic equation $\frac{1}{2}a^3 + (1+\rho)a = \gamma_0 y$. This implies that, for small γ_0 we have $a \sim \frac{\gamma_0 y}{1+\rho}$ so that $\Delta H = 2\Delta\tilde{G} \sim \frac{\gamma_0^2 y^2}{(1+\rho)^2}$ and so that larger initial alignment ρ leads to smaller changes in the feature kernel and pseudo-gradient alignment kernel. At large $\gamma_0 y$, we have that $a \sim (2\gamma_0 y)^{1/3}$ so that $\Delta H = 2\Delta\tilde{G} \sim (2\gamma_0 y)^{2/3}$.

F.3 HEBB

For the Hebb rule, $\tilde{G}_\mu = \langle gh_\mu \rangle \Delta_\mu = \gamma_0 f_\mu \Delta_\mu = \gamma_0 (y_\mu - \Delta_\mu) \Delta_\mu$. Under the whitening assumption $K_{\mu\nu}^x = \delta_{\mu\nu}$, the dynamics decouples over samples

$$\frac{d}{dt}H_{\mu,\mu} = 2\gamma_0 H_{\mu\mu} \Delta_\mu^2, \quad \frac{d}{dt}\Delta_\mu = -[H_{\mu\mu} + \gamma_0 (y_\mu - \Delta_\mu) \Delta_\mu] \Delta_\mu. \quad (73)$$

We see that $H_{\mu\mu}$ strictly increases. The possible fixed points for Δ_μ are $\Delta_\mu = 0$ or $\Delta_\mu = \frac{1}{2}\left[y_\mu \pm \sqrt{y_\mu^2 + \gamma_0^{-1}H_{\mu\mu}}\right]$. One of these roots shares a sign with y_μ and has larger absolute value. The other root has the opposite sign from y_μ . From the initial condition $\Delta_\mu = y_\mu$ and $H_{\mu\mu} = 1$, Δ_μ is initially approaching decreasing in absolute value so that $|\Delta_\mu| \in (0, |y_\mu|)$ and will have the same sign as y_μ . In this regime $\frac{d}{dt}|\Delta_\mu| < 0$. Thus, the system will eventually reach the fixed point at $\Delta_\mu = 0$, rather than increasing in magnitude to the root which shares a sign with y_μ or continuing to the root with the opposite sign as y_μ .

G DISCUSSION OF MODIFIED HEBB RULE

We chose to modify the traditional Hebb rule to include a weighing of each example by its instantaneous error. In this section we discuss this choice and provide a brief discussion of alternatives

- *Traditional Hebb Learning*: $\frac{d}{dt}\mathbf{W}^\ell \propto \sum_\mu \phi(\mathbf{h}^{\ell+1})\phi(\mathbf{h}^\ell)^\top$. In the absence of regularization or normalization, this learning rule will continue to update the weights even once the task is fully learned, leading to divergences at infinite time $t \rightarrow \infty$.
- *Single Power of the Error*: $\frac{d}{dt}\mathbf{W}^\ell \propto \sum_\mu \Delta_\mu \phi(\mathbf{h}^{\ell+1})\phi(\mathbf{h}^\ell)^\top$. While this rule may naively appear plausible, it can only learn training points with positive target values y_μ in a linear network if $\gamma_0 > 0$. Further this rule only gives Hebbian updates when $\Delta_\mu > 0$.
- *Two Powers of the Error*: $\frac{d}{dt}\mathbf{W}^\ell \propto \sum_\mu \Delta_\mu^2 \phi(\mathbf{h}^{\ell+1})\phi(\mathbf{h}^\ell)^\top$. This was our error modified Hebb rule. We note that the update always has the correct sign for a Hebbian update and the updates stop when the network converges to zero error, preventing divergence of the features at late time.

H FINITE SIZE EFFECTS

We can reason about the fluctuations of \mathbf{q} around the saddle point \mathbf{q}^* at large but finite N using a Taylor expansion of the DMFT action S around the saddle point. This argument will show that

at large but finite N , we can treat q as fluctuating over initializations with mean q^* and variance $O(N^{-1})$. We will first illustrate the mechanics of this computation of an arbitrary observable with a scalar example before applying this to the DMFT.

H.1 SCALAR EXAMPLE

Suppose we have a scalar variable q with a distribution defined by Gibbs measure $\frac{e^{-NS[q]}}{\int dq e^{-NS[q]}}$ for action S . We consider averaging some arbitrary observable $\mathcal{O}(q)$ over this distribution

$$\langle \mathcal{O}(q) \rangle = \frac{\int dq \exp(-NS[q]) \mathcal{O}(q)}{\int dq \exp(-NS[q])}. \quad (74)$$

We Taylor expand S around its saddle point q^* giving $S[q] = S[q^*] + \frac{1}{2}S''[q^*](q - q^*)^2 + \sum_{k=3}^{\infty} S^{(k)}[q^*](q - q^*)^k$. This gives

$$\langle \mathcal{O}(q) \rangle = \frac{\int dq \exp(-N[\frac{1}{2}S''[q^*](q - q^*)^2 - \sum_{k=3}^{\infty} S^{(k)}[q^*](q - q^*)^k]) \mathcal{O}(q)}{\int dq \exp(-N[\frac{1}{2}S''[q^*](q - q^*)^2 - \sum_{k=3}^{\infty} S^{(k)}[q^*](q - q^*)^k])}. \quad (75)$$

The $\exp(NS[q^*])$ terms canceled in both numerator and denominator. We let the variable $q - q^* = \frac{1}{\sqrt{N}}\delta$. After this change of variable, we have

$$\langle \mathcal{O}(q) \rangle = \frac{\int d\delta \exp(-\frac{1}{2}S''[q^*]\delta^2 - \sum_{k=3}^{\infty} N^{1-k/2}S^{(k)}[q^*]\delta^k) \mathcal{O}(q^* + N^{-1/2}\delta)}{\int d\delta \exp(-\frac{1}{2}S''[q^*]\delta^2 - \sum_{k=3}^{\infty} N^{1-k/2}S^{(k)}[q^*]\delta^k)}. \quad (76)$$

We note that all the higher order derivatives ($k \geq 3$) are suppressed by at least $N^{-1/2}$ compared to the quadratic term. Letting $U = \sum_{k=3}^{\infty} N^{1-k/2}S^{(k)}[q^*]\delta^k$ represent the *perturbed potential*, we can Taylor expand the exponential around the Gaussian unperturbed potential $\exp(-\frac{1}{2}\delta^2 S''[q^*])$. We let $\langle \mathcal{O}(\delta) \rangle_0 = \mathbb{E}_{q \sim \mathcal{N}(0, S''[q^*]^{-1})} \mathcal{O}(\delta)$ represent an average over this unperturbed potential

$$\langle \mathcal{O}(q) \rangle = \frac{\int d\delta \exp(-\frac{1}{2}S''[q^*]\delta^2) [1 - U + \frac{1}{2}U^2 + \dots] \mathcal{O}(q)}{\int d\delta \exp(-\frac{1}{2}S''[q^*]\delta^2) [1 - U + \frac{1}{2}U^2 + \dots]} \quad (77)$$

$$= \frac{\langle \mathcal{O}(q) \rangle_0 - \langle \mathcal{O}(q)U \rangle_0 + \frac{1}{2} \langle \mathcal{O}(q)U^2 \rangle_0 + \dots}{1 - \langle U \rangle_0 + \frac{1}{2} \langle U^2 \rangle_0 + \dots}. \quad (78)$$

Truncating each series in numerator and denominator at a certain order in $1/N$ gives a Pade-Approximant to the full observable average (Bender et al., 1999). Alternatively, this can be expressed in terms of a cumulant expansion (Kardar, 2007)

$$\langle \mathcal{O} \rangle = \sum_{k=0}^{\infty} \frac{(-1)^k}{k!} \langle \mathcal{O}U^k \rangle_0^c, \quad (79)$$

where $\langle \mathcal{O}U^k \rangle_0^c$ are the *connected* correlations, or alternatively the *cumulants*. The first two connected correlations have the form

$$\begin{aligned} \langle \mathcal{O}U \rangle_0^c &= \langle \mathcal{O}U \rangle_0 - \langle \mathcal{O} \rangle \langle U \rangle_0 \\ \langle \mathcal{O}U^2 \rangle_0^c &= \langle \mathcal{O}U^2 \rangle_0 - 2 \langle \mathcal{O}U \rangle_0 \langle U \rangle_0 - \langle \mathcal{O} \rangle_0 \langle U^2 \rangle_0 + 2 \langle \mathcal{O} \rangle \langle U \rangle_0^2. \end{aligned} \quad (80)$$

Using Stein's lemma, we can now attempt to extract the leading $O(N^{-1})$ behavior from each of these terms. First, we will note the following useful identity which follows from Stein's Lemma

$$\langle \mathcal{O}(q)\delta^k \rangle = N^{k/2} \langle \mathcal{O}(q)(q - q^*)^k \rangle \quad (81)$$

$$= N^{k/2-1} [S'']^{-1} [(k-1) \langle \mathcal{O}(q)(q - q^*)^{k-2} \rangle + \langle \mathcal{O}'(q)(q - q^*)^{k-1} \rangle]$$

$$= (k-1) [S'']^{-1} \langle \mathcal{O}(q)\delta^{k-2} \rangle + \frac{1}{\sqrt{N}} [S'']^{-1} \langle \mathcal{O}'(q)\delta^{k-1} \rangle. \quad (82)$$

Using these this fact, we can find the first few correlation functions of interest

$$\begin{aligned}\langle \mathcal{O}(q)U \rangle &= 3N^{-1}S^{(3)}[S'']^{-2} \langle \mathcal{O}'(q) \rangle_0 + 3N^{-1}S^{(4)}[S'']^{-2} \langle \mathcal{O}(q) \rangle_0 + O(N^{-2}) \\ \langle \mathcal{O}(q)U^2 \rangle &= 15N^{-1}[S^{(3)}]^2[S'']^{-3} \langle \mathcal{O}(q) \rangle + O(N^{-1}).\end{aligned}\quad (83)$$

Thus, the leading order Pade-Approximant has the form

$$\langle \mathcal{O}(q) \rangle = \frac{\langle \mathcal{O} \rangle_0 - \frac{3}{N}S^{(3)}[S'']^{-2} \langle \mathcal{O}'(q) \rangle_0 - \frac{3}{N}S^{(4)}[S'']^{-2} \langle \mathcal{O}(q) \rangle_0 + \frac{15}{2N}[S^{(3)}]^2[S'']^{-3} \langle \mathcal{O}(q) \rangle}{1 - \frac{3}{N}S^{(3)}[S'']^{-2} - \frac{3}{N}S^{(4)}[S'']^{-2} + \frac{15}{2N}[S^{(3)}]^2[S'']^{-3}}.\quad (84)$$

H.1.1 DMFT ACTION EXPANSION

The logic of the previous section can be extended to our DMFT. We first redefine the action as its negation $S \rightarrow -S$ to simplify the argument. Concretely, this action $S[\mathbf{q}]$ defines a Gibbs measure over the order parameters \mathbf{q} which we can use to compute observable averages

$$\langle \mathcal{O}(\mathbf{q}) \rangle = \frac{\int \exp(-NS[\mathbf{q}]) \mathcal{O}(\mathbf{q})}{\int \exp(-NS[\mathbf{q}])}\quad (85)$$

As before, one can Taylor expand the action around the saddle point \mathbf{q}^*

$$S[\mathbf{q}] \sim S[\mathbf{q}^*] + \frac{1}{2}(\mathbf{q} - \mathbf{q}^*)\nabla^2 S[\mathbf{q}^*](\mathbf{q} - \mathbf{q}^*) + \dots\quad (86)$$

As before, the linear term vanishes since $\nabla_{\mathbf{q}} S[\mathbf{q}^*] = 0$ at the saddle point \mathbf{q}^* . We again change variables to $\boldsymbol{\delta} = \sqrt{N}(\mathbf{q} - \mathbf{q}^*)$ and express the average as

$$\begin{aligned}\langle \mathcal{O} \rangle &= \frac{\int d\boldsymbol{\delta} \exp(-\frac{1}{2}\boldsymbol{\delta}^\top \nabla^2 S \boldsymbol{\delta} + U(\boldsymbol{\delta})) \mathcal{O}(\boldsymbol{\delta})}{\int d\boldsymbol{\delta} \exp(-\frac{1}{2}\boldsymbol{\delta}^\top \nabla^2 S \boldsymbol{\delta} + U(\boldsymbol{\delta}))} \\ &= \sum_{k=0}^{\infty} \frac{(-1)^k}{k!} \langle \mathcal{O}U^k \rangle_0^c\end{aligned}\quad (87)$$

where $\langle \cdot \rangle_0$ denotes a Gaussian average over $\mathbf{q} \sim \mathcal{N}(\mathbf{q}^*, \frac{1}{N}[\nabla^2 S]^{-1})$.

H.2 HESSIAN COMPONENTS OF DMFT ACTION

To gain insight into the Hessian, we will first restrict our attention to the subset of Hessian entries related to $\Phi^\ell, \hat{\Phi}^\ell$. We again adopt a multi-index notation $\boldsymbol{\mu} = (\mu, \nu, t, s)$ so that $\Phi_{\boldsymbol{\mu}}^\ell = \Phi_{\mu\nu}^\ell(t, s)$

$$\begin{aligned}\frac{\partial^2 S}{\partial \Phi_{\boldsymbol{\mu}}^\ell \partial \Phi_{\boldsymbol{\mu}'}^{\ell'}} &= 0 \\ \frac{\partial^2 S}{\partial \Phi_{\boldsymbol{\mu}}^\ell \partial \hat{\Phi}_{\boldsymbol{\mu}'}^{\ell'}} &= \delta_{\ell, \ell'} \delta_{\boldsymbol{\mu}, \boldsymbol{\mu}'} - \delta_{\ell', \ell+1} \frac{\partial}{\partial \Phi_{\boldsymbol{\mu}}^\ell} \Phi_{\boldsymbol{\mu}'}^{\ell+1} \\ \frac{\partial^2 S}{\partial \hat{\Phi}_{\boldsymbol{\mu}}^\ell \partial \hat{\Phi}_{\boldsymbol{\mu}'}^{\ell'}} &= \delta_{\ell, \ell'} [\langle \phi(h_{\boldsymbol{\mu}}^\ell(t)) \phi(h_{\boldsymbol{\nu}}^\ell(s)) \phi(h_{\boldsymbol{\mu}'}^{\ell'}(t')) \phi(h_{\boldsymbol{\nu}'}^{\ell'}(s')) \rangle - \Phi_{\boldsymbol{\mu}}^\ell \Phi_{\boldsymbol{\mu}'}^{\ell'}].\end{aligned}$$

The first equation follows from the fact that $\hat{\chi}$ has vanishing moments due to the normalization of the probability distribution induced by \mathcal{Z}^ℓ . Similarly, for the $\mathbf{G}, \hat{\mathbf{G}}$ kernels we have

$$\begin{aligned}\frac{\partial^2 S}{\partial G_{\boldsymbol{\mu}}^\ell \partial G_{\boldsymbol{\mu}'}^{\ell'}} &= 0 \\ \frac{\partial^2 S}{\partial G_{\boldsymbol{\mu}}^\ell \partial \hat{G}_{\boldsymbol{\mu}'}^{\ell'}} &= \delta_{\ell, \ell'} \delta_{\boldsymbol{\mu}, \boldsymbol{\mu}'} - \delta_{\ell', \ell-1} \frac{\partial}{\partial G_{\boldsymbol{\mu}}^\ell} G_{\boldsymbol{\mu}'}^{\ell-1} \\ \frac{\partial^2 S}{\partial \hat{\Phi}_{\boldsymbol{\mu}}^\ell \partial \hat{\Phi}_{\boldsymbol{\mu}'}^{\ell'}} &= \delta_{\ell, \ell'} [\langle g_{\boldsymbol{\mu}}^\ell(t) g_{\boldsymbol{\nu}}^\ell(s) g_{\boldsymbol{\mu}'}^{\ell'}(t') g_{\boldsymbol{\nu}'}^{\ell'}(s') \rangle - G_{\boldsymbol{\mu}}^\ell G_{\boldsymbol{\mu}'}^{\ell'}].\end{aligned}$$

Proceeding in a similar manner, we can compute all off-diagonal components such as $\frac{\partial^2 S}{\partial \Phi \partial \hat{\mathbf{G}}}$ and $\frac{\partial^2 S}{\partial \hat{\Phi} \partial \hat{\mathbf{G}}}$. Once all entries are computed, one can seek an inverse of the Hessian to obtain the covariance of the order parameters.

H.3 SINGLE SAMPLE NEXT-TO-LEADING ORDER PERTURBATION THEORY

In order to obtain exact analytical expressions, we will consider L -hidden layer ReLU and linear neural networks in the lazy regime trained on a single sample with $K^x = \frac{|\mathbf{x}|^2}{D} = 1$. To ensure preservation of norm, we will use $\phi(h) = \sqrt{2}h\Theta(h)$ for ReLU and $\phi(h) = h$ for linear networks. First, we note that in either case, the infinite width saddle point equations give

$$\begin{aligned}\Phi^\ell &= \langle \phi(h)^2 \rangle_{h \sim \mathcal{N}(0, \Phi^{\ell-1})} = \Phi^{\ell-1}, \quad \Phi^0 = 1 \\ G^\ell &= \langle \dot{\phi}(h)^2 z^2 \rangle_{h, z} = G^{\ell+1}, \quad G^{L+1} = 1 \\ \implies \Phi^\ell &= 1, \quad G^\ell = 1, \quad \forall \ell \in [1, \dots, L].\end{aligned}\tag{88}$$

At large but finite width, the kernels therefore fluctuate around this typical mean value of $\Phi^\ell = 1$ and $G^\ell = 1$. We now compute the necessary ingredients to invert the Hessian

$$\begin{aligned}V_\phi &= \langle \phi(h)^4 \rangle - \Phi^2 = \begin{cases} 5 & \text{ReLU} \\ 2 & \text{Linear} \end{cases} \\ V_g &= \langle \dot{\phi}(h)^4 z^4 \rangle - G^2 = \begin{cases} 5 & \text{ReLU} \\ 2 & \text{Linear} \end{cases}.\end{aligned}\tag{89}$$

Next, we compute the sensitivity of each layer's kernel to the previous layer

$$\frac{\partial}{\partial \Phi^\ell} \Phi^{\ell+1} = 1, \quad \frac{\partial}{\partial G^{\ell+1}} G^\ell = 1.\tag{90}$$

First, let's analyze the marginal covariance statistics for $\Phi = \text{Vec}\{\Phi^\ell\}_{\ell=1}^L$ and $\hat{\Phi} = \text{Vec}\{\hat{\Phi}^\ell\}_{\ell=1}^L$. We note that the DMFT action has Hessian components

$$\mathbf{H}_\Phi = \begin{bmatrix} \nabla_{\Phi}^2 S & \nabla_{\Phi \hat{\Phi}}^2 S \\ \nabla_{\hat{\Phi} \Phi}^2 S & \nabla_{\hat{\Phi}}^2 S \end{bmatrix} = \begin{bmatrix} 0 & \mathbf{U} \\ \mathbf{U}^\top & V_\phi \mathbf{I} \end{bmatrix}, \quad \mathbf{U} = \begin{bmatrix} 1 & -1 & 0 & \dots & 0 & 0 \\ 0 & 1 & -1 & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \dots & 1 & -1 \\ 0 & 0 & 0 & \dots & 0 & 1 \end{bmatrix}.\tag{91}$$

We seek a (physical) inverse \mathbf{C} which has vanishing lower diagonal entry, indicating zero variance in the dual order parameters $\hat{\Phi}$. This gives us the following linear equations

$$\begin{aligned}\mathbf{H}_\Phi \mathbf{C} &= \begin{bmatrix} 0 & \mathbf{U} \\ \mathbf{U}^\top & V_\phi \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{C}_{11} & \mathbf{C}_{12} \\ \mathbf{C}_{12}^\top & 0 \end{bmatrix} = \begin{bmatrix} \mathbf{I} & 0 \\ 0 & \mathbf{I} \end{bmatrix} \\ \implies \mathbf{U} \mathbf{C}_{12}^\top &= \mathbf{I}, \quad \mathbf{U}^\top \mathbf{C}_{11} + V_\phi \mathbf{C}_{12}^\top = 0.\end{aligned}\tag{92}$$

The relevant entry is $\mathbf{C}_{11} = -V_\phi [\mathbf{U}^\top]^{-1} \mathbf{U}^{-1}$. This matrix has the form

$$\mathbf{C}_{11} = -V_\phi \begin{bmatrix} 1 & 0 & 0 & \dots & 0 \\ 1 & 1 & 0 & \dots & 0 \\ 1 & 1 & 1 & \dots & 0 \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ 1 & 1 & 1 & \dots & 1 \end{bmatrix} \begin{bmatrix} 1 & 1 & 1 & \dots & 1 \\ 0 & 1 & 1 & \dots & 1 \\ 0 & 0 & 1 & \dots & 1 \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & 1 \end{bmatrix} = -V_\phi \begin{bmatrix} 1 & 1 & 1 & \dots & 1 \\ 1 & 2 & 2 & \dots & 2 \\ 1 & 2 & 3 & \dots & 3 \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ 1 & 2 & 3 & \dots & L \end{bmatrix}.\tag{93}$$

Using the fact that the covariance is the negative of the Hessian inverse multiplied by $1/N$, we have the following covariance structure for $\{\Phi^\ell\}$

$$\text{Cov}(\Phi^\ell, \Phi^{\ell'}) = \frac{1}{N} V_\phi \min\{\ell, \ell'\}.\tag{94}$$

This result can be interpreted as the covariance of Brownian motion. Following an identical argument, we find

$$\text{Cov}(G^\ell, G^{\ell'}) = \frac{1}{N} V_g \min\{L+1-\ell, L+1-\ell'\}.\tag{95}$$

We verify these scalings against experiments below in Figure 7.

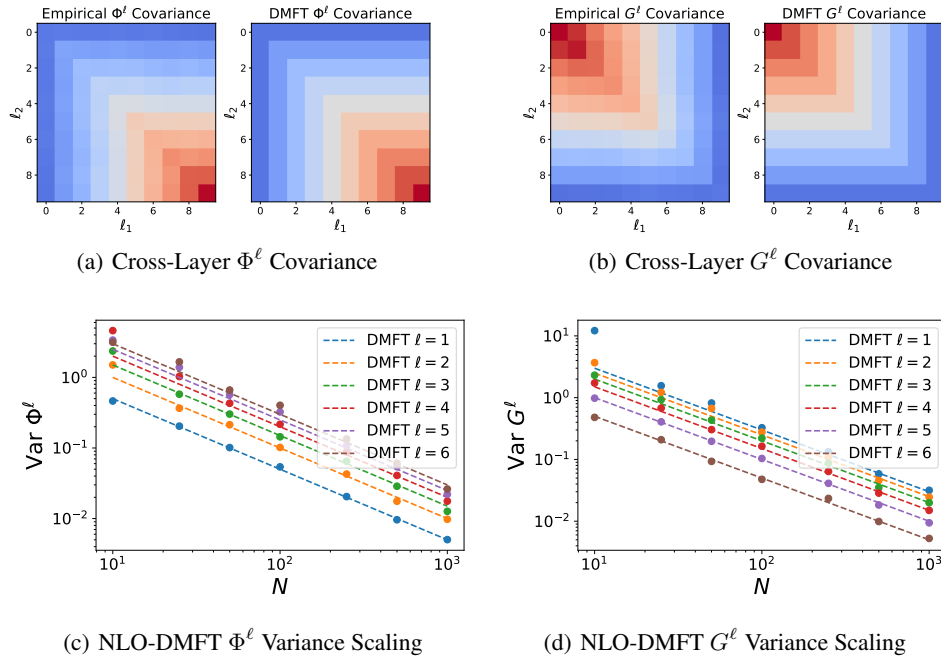


Figure 7: Verification of kernel fluctuations through next-to-leading-order (NLO) perturbation theory within DMFT formalism. (a) The cross layer covariance structure of $\{\Phi^\ell\}$ in a $L = 10$ hidden layer ReLU MLP. The empirical covariance was estimated by initializing a large number (500) of random networks and computing their Φ^ℓ kernels. We see that variance for Φ^ℓ increases as ℓ increases. (b) The cross-layer covariance structure of $\{G^\ell\}$. The variance of G^ℓ is larger for smaller ℓ . (c) The predicted variance of Φ^ℓ for different layer ℓ and widths N . All layers have variance scaling as $1/N$, consistent with NLO perturbation theory. (d) The scaling of G^ℓ variance.