# Supplemental Material for: "Asymptotics of representation learning in finite Bayesian neural networks" 

Jacob A. Zavatone-Veth ${ }^{1,2}$, Abdulkadir Canatar ${ }^{1,2}$, Benjamin S. Ruben ${ }^{3}$, Cengiz Pehlevan ${ }^{2,4}$<br>${ }^{1}$ Department of Physics, ${ }^{2}$ Center for Brain Science, ${ }^{3}$ Biophysics Graduate Program,<br>${ }^{4}$ John A. Paulson School of Engineering and Applied Sciences Harvard University Cambridge, MA 02138<br>\{jzavatoneveth, canatara, benruben\}@g.harvard.edu<br>cpehlevan@seas.harvard.edu

## Contents

A Preliminary technical results ..... S2
A. 1 Isserlis' theorem for Gaussian moments ..... S2
A. 2 Neumann series for matrix inverses near the identity ..... S2
A. 3 Series expansion of the log-determinant near the identity ..... S2
B Perturbation theory for wide Bayesian neural networks with linear readout ..... S3
B. 1 Integrating out the readout layer ..... S4
B. 2 Perturbative expansion ..... S5
C Explicit covariance computations in deep linear networks ..... S6
C. 1 Fully-connected linear networks ..... S7
C. 2 Convolutional linear networks ..... S8
D Direct computation of the average hidden layer kernels of a deep linear MLP ..... S11
D. 1 The cumulant generating function of learned features for a MLP ..... S11
D. 2 General form of the perturbative layer integrals for a deep linear network ..... S12
D. 3 Perturbative computation of the partition function of a deep linear network ..... S16
D. 4 Computing the average hidden layer kernels of a deep linear network ..... S18
E Average kernels in a deep feedforward linear network with skip connections ..... S19
E. 1 Perturbative computation of the partition function ..... S19
E. 2 Computing the average hidden layer kernels ..... S22
F Comparison to the results of Aitchison (2020) and Li \& Sompolinsky (2020) ..... S23
F. 1 Comparison to the results of Aitchison (2020) ..... S23
F. 2 Comparison to the results of Li and Sompolinsky (2020) ..... S24
G Predictor statistics and generalization in deep linear networks ..... S25
G. 1 Predictor statistics ..... S26
G. 2 Bias-variance decompositions and the low-temperature limit ..... S26
G. 3 Effects of alternative regularization temperature-dependence ..... S27
H Derivation of the average kernels for a depth-two network ..... S28
I Numerical methods ..... S29
Checklist ..... S31

## A Preliminary technical results

In this appendix, we review useful technical results upon which our calculations rely.

## A. 1 Isserlis' theorem for Gaussian moments

Let $\left(x_{1}, x_{2}, \ldots, x_{n}\right)$ be a zero-mean Gaussian random vector. Then, Isserlis' theorem [1] states that

$$
\mathbb{E}\left[x_{1} x_{2} \cdots x_{n}\right]= \begin{cases}\sum_{p \in P_{n}^{2}} \prod_{(i, j) \in p} \operatorname{cov}\left(x_{i}, x_{j}\right) & n \text { even }  \tag{A.1}\\ 0 & n \text { odd }\end{cases}
$$

where the sum is over all pairings $p$ of $\{1,2, \ldots, n\}$ and the product is over all pairs contained in $p$. In particular, for $n=4$, we have

$$
\begin{equation*}
\mathbb{E}\left[x_{1} x_{2} x_{3} x_{4}\right]=\operatorname{cov}\left(x_{1}, x_{2}\right) \operatorname{cov}\left(x_{3}, x_{4}\right)+\operatorname{cov}\left(x_{1}, x_{3}\right) \operatorname{cov}\left(x_{2}, x_{4}\right)+\operatorname{cov}\left(x_{1}, x_{4}\right) \operatorname{cov}\left(x_{2}, x_{3}\right) \tag{A.2}
\end{equation*}
$$

In physics, Isserlis' theorem is often known as Wick's probability theorem [2].

## A. 2 Neumann series for matrix inverses near the identity

The Neumann series is the generalization of the geometric series to bounded linear operators, including square matrices. In particular, let $A$ be a $p \times p$ square matrix. Then, we have

$$
\begin{equation*}
\left(I_{p}-A\right)^{-1}=\sum_{k=0}^{\infty} A^{k} \tag{A.3}
\end{equation*}
$$

provided that the series converges in the operator norm [3]. We will use this result without concern for rigorous convergence conditions, as we are interested only in asymptotic expansions.

## A. 3 Series expansion of the log-determinant near the identity

Let $A$ be a $p \times p$ square matrix, and let $t$ be a small parameter. Then, we have

$$
\begin{equation*}
\log \operatorname{det}\left(I_{p}+t A\right)=\sum_{k=1}^{\infty} \frac{(-1)^{k+1}}{k} \operatorname{tr}\left(A^{k}\right) t^{k} \tag{A.4}
\end{equation*}
$$

assuming that the series converges. We will not concern ourselves with rigorous convergence conditions, as we will use this expansion formally.
This result follows from the fact that

$$
\begin{equation*}
\frac{\partial^{k}}{\partial t^{k}} \log \operatorname{det}\left(I_{p}+t A\right)=(-1)^{k+1}(k-1)!\operatorname{tr}\left(\left(I_{p}+t A\right)^{-k} A^{k}\right) \quad(k=1,2, \ldots) \tag{A.5}
\end{equation*}
$$

The base case $k=1$ is given by Jacobi's formula [3]:

$$
\begin{equation*}
\frac{\partial}{\partial t} \log \operatorname{det}\left(I_{p}+t A\right)=\operatorname{tr}\left(\left(I_{p}+t A\right)^{-1} A\right) \tag{A.6}
\end{equation*}
$$

Then, using the identity

$$
\begin{equation*}
\frac{\partial}{\partial t}\left(I_{p}+t A\right)^{-1}=-\left(I_{p}+t A\right)^{-1} A\left(I_{p}+t A\right)^{-1} \tag{A.7}
\end{equation*}
$$

and the fact that $A$ commutes with $\left(I_{p}+t A\right)^{-1}$, we find that the claim holds by induction. As $\left.\log \operatorname{det}\left(I_{p}+t A\right)\right|_{t=0}=0$, this implies the desired Maclaurin series.

## B Perturbation theory for wide Bayesian neural networks with linear readout

In this appendix, we derive Conjecture 1. As outlined in the main text, we consider a depth- $d$ neural network $\mathbf{f}: \mathbb{R}^{n_{0}} \rightarrow \mathbb{R}^{n_{d}}$ with linear readout, written as

$$
\begin{equation*}
\mathbf{f}\left(\mathbf{x} ; W^{d}, \mathcal{W}\right)=\frac{1}{\sqrt{n_{d-1}}} W^{(d)} \boldsymbol{\psi}(\mathbf{x} ; \mathcal{W}) \tag{B.1}
\end{equation*}
$$

in terms of the hidden layer feature map $\psi(\cdot ; \mathcal{W}): \mathbb{R}^{n_{0}} \rightarrow \mathbb{R}^{n_{d-1}}$. The full set of trainable parameters is then $\Theta=\left\{W^{(d)}, \mathcal{W}\right\}$, where $\mathcal{W}$ is the set of feature map parameters. We assume isotropic Gaussian priors over these parameters, with, for instance,

$$
\begin{equation*}
W_{i j}^{(d)} \underset{\text { i.i.d. }}{\sim} \mathcal{N}\left(0, \sigma_{d}^{2}\right) \tag{B.2}
\end{equation*}
$$

We fix an arbitrary training dataset $\mathcal{D}=\left\{\left(\mathbf{x}_{\mu}, \mathbf{y}_{\mu}\right)\right\}_{\mu=1}^{p}$ of $p$ examples, and use a Gaussian likelihood $p(\mathcal{D} \mid \Theta) \propto \exp (-\beta E)$, where

$$
\begin{equation*}
E(\Theta ; \mathcal{D})=\frac{1}{2} \sum_{\mu=1}^{p}\left\|\mathbf{f}\left(\mathbf{x}_{\mu}\right)-\mathbf{y}_{\mu}\right\|^{2} \tag{B.3}
\end{equation*}
$$

is a quadratic cost. We then introduce the Bayes posterior

$$
\begin{equation*}
p(\Theta \mid \mathcal{D})=\frac{p(\mathcal{D} \mid \Theta) p(\Theta)}{p(\mathcal{D})} \tag{B.4}
\end{equation*}
$$

averages with respect to this distribution will be denoted by $\langle\cdot\rangle$.
We define the postactivation feature map kernel

$$
\begin{equation*}
K^{(d-1)}\left(\mathbf{x}, \mathbf{x}^{\prime}\right) \equiv n_{d-1}^{-1} \boldsymbol{\psi}(\mathbf{x}, \mathcal{W}) \cdot \boldsymbol{\psi}\left(\mathbf{x}^{\prime}, \mathcal{W}\right) \tag{B.5}
\end{equation*}
$$

and write $\left[K^{(d-1)}\right]_{\mu \nu} \equiv K^{(d-1)}\left(\mathbf{x}_{\mu}, \mathbf{x}_{\nu}\right)$ for the kernel evaluated on the training set. For brevity, we will frequently abbreviate $K \equiv K^{(d-1)}$ throughout this appendix.
We denote expectation by $\mathbb{E}$, and prior expectation by $\mathbb{E}_{\mathcal{W}}$. We also introduce the joint cumulant operator $\mathbb{K}$ and its prior counterpart $\mathbb{K}_{\mathcal{W}}$. We will only require the second and third joint cumulants, which, for random variables $A, B$, and $C$, are given as

$$
\begin{equation*}
\mathbb{K}(A, B)=\mathbb{E}[(A-\mathbb{E} A)(B-\mathbb{E} B)] \tag{B.6}
\end{equation*}
$$

and

$$
\begin{equation*}
\mathbb{K}(A, B, C)=\mathbb{E}[(A-\mathbb{E} A)(B-\mathbb{E} B)(C-\mathbb{E} C)] \tag{B.7}
\end{equation*}
$$

respectively.
Our starting point is the partition function $Z$ of the Bayes posterior (B.4) for the network (B.1), including a source term for the (generically matrix-valued) observable $O$ :

$$
\begin{equation*}
Z(J)=\mathbb{E}_{W^{(d)}} \mathbb{E}_{\mathcal{W}} \exp \left(-\beta E+\operatorname{tr}\left(J^{\top} O\right)\right) \tag{B.8}
\end{equation*}
$$

where $\mathcal{W}$ denotes all of the parameters except for the readout weight matrix $W^{(d)}$ and expectation is taken with respect to the Gaussian prior. The logarithm of the partition function is the posterior cumulant generating function of the observable $O$, with

$$
\begin{equation*}
\langle O\rangle=\left.\frac{\delta \log Z}{\delta J}\right|_{J=0} \tag{B.9}
\end{equation*}
$$

and covariance

$$
\begin{equation*}
\operatorname{cov}\left(O_{\rho \gamma}, O_{\omega \chi}\right)=\left.\frac{\partial^{2} \log Z}{\partial J_{\rho \gamma} \partial J_{\omega \chi}}\right|_{J=0} \tag{B.10}
\end{equation*}
$$

## B. 1 Integrating out the readout layer

We first show that the readout layer can be integrated out exactly. As the source term is independent of $W^{(d)}$, Fubini's theorem yields

$$
\begin{equation*}
Z=\mathbb{E}_{\mathcal{W}}\left[\exp \left(\operatorname{tr}\left(J^{\top} O\right)\right) \mathbb{E}_{W^{(d)}} \exp (-\beta E)\right] \tag{B.11}
\end{equation*}
$$

The expectation over $W^{d}$ is a Gaussian integral, hence it is easy to evaluate exactly:

$$
\begin{align*}
& \mathbb{E}_{W^{(d)}} \exp (-\beta E)  \tag{B.12}\\
& =\mathbb{E}_{W^{(d)}} \exp \left(-\frac{1}{2} \beta \sum_{\mu=1}^{p}\left\|\frac{1}{\sqrt{n_{d-1}}} W^{(d)} \boldsymbol{\psi}_{\mu}-\mathbf{y}_{\mu}\right\|^{2}\right)  \tag{B.13}\\
& =\exp \left(-\frac{1}{2} \beta \operatorname{tr}\left(Y^{\top} Y\right)\right) \\
& \quad \times \prod_{j=1}^{n_{d}}\left[\int \frac{d \mathbf{w}_{j}}{\left(2 \pi \sigma_{d}^{2}\right)^{n_{d-1} / 2}} \exp \left(-\frac{1}{2} \mathbf{w}_{j}^{\top}\left(\sigma_{d}^{-2} I_{n}+\frac{\beta}{n_{d-1}} \Psi^{\top} \Psi\right) \mathbf{w}_{j}+\frac{\beta}{\sqrt{n_{d-1}}}\left(Y^{\top} \Psi\right)_{j} \cdot \mathbf{w}_{j}\right)\right]  \tag{B.14}\\
& =\operatorname{det}\left(I_{n}+\frac{\beta \sigma_{d}^{2}}{n_{d-1}} \Psi^{\top} \Psi\right)^{-n_{d} / 2} \\
& \quad \times \exp \left(\frac{1}{2} \frac{\beta^{2} \sigma_{d}^{2}}{n_{d-1}} \operatorname{tr}\left[Y^{\top} \Psi\left(I_{n}+\frac{\beta \sigma_{d}^{2}}{n_{d-1}} \Psi^{\top} \Psi\right)^{-1} \Psi^{\top} Y\right]-\frac{1}{2} \beta \operatorname{tr}\left(Y^{\top} Y\right)\right), \tag{B.15}
\end{align*}
$$

where we abbreviate $\boldsymbol{\psi}_{\mu} \equiv \boldsymbol{\psi}\left(\mathbf{x}_{\mu} ; \mathcal{W}\right)$ and introduce the matrices $\Psi_{\mu j} \equiv \psi_{\mu, j}$ and $Y_{\mu j} \equiv y_{\mu, j}$. Here, we have used the fact that the matrix $I_{n}+\left(\beta \sigma_{d}^{2} / n_{d-1}\right) \Psi^{\top} \Psi$ is invertible at any finite temperature. By the Weinstein-Aronszajn identity [3],

$$
\begin{equation*}
\operatorname{det}\left(I_{n}+\frac{\beta \sigma_{d}^{2}}{n_{d-1}} \Psi^{\top} \Psi\right)=\operatorname{det}\left(I_{p}+\frac{\beta \sigma_{d}^{2}}{n_{d-1}} \Psi \Psi^{\top}\right)=\operatorname{det}\left(I_{p}+\beta \sigma_{d}^{2} K\right) \tag{B.16}
\end{equation*}
$$

where we introduce the (non-constant) kernel matrix

$$
\begin{equation*}
K=K^{(d-1)} \equiv \frac{1}{n_{d-1}} \Psi \Psi^{\top} \tag{B.17}
\end{equation*}
$$

as mentioned above, we abbreviate $K \equiv K^{(d-1)}$ for brevity. By the push-through identity [3],

$$
\begin{equation*}
\frac{1}{n_{d-1}} \Psi\left(I_{n}+\frac{\beta \sigma_{d}^{2}}{n_{d-1}} \Psi^{\top} \Psi\right)^{-1} \Psi^{\top}=\left(I_{p}+\frac{\beta \sigma_{d}^{2}}{n_{d-1}} \Psi \Psi^{\top}\right)^{-1} \frac{1}{n_{d-1}} \Psi \Psi^{\top}=\left(I_{p}+\beta \sigma_{d}^{2} K\right)^{-1} K \tag{B.18}
\end{equation*}
$$

hence, using the cyclic property of the trace,

$$
\begin{align*}
& \frac{1}{2} \frac{\beta^{2} \sigma_{d}^{2}}{n_{d-1}} \operatorname{tr}\left[Y^{\top} \Psi\left(I_{n}+\frac{\beta \sigma_{d}^{2}}{n_{d-1}} \Psi^{\top} \Psi\right)^{-1} \Psi^{\top} Y\right]-\frac{1}{2} \beta \operatorname{tr}\left(Y^{\top} Y\right) \\
& \quad=\frac{1}{2} \beta n_{d} \operatorname{tr}\left[\left(\beta \sigma_{d}^{2}\left(I_{p}+\beta \sigma_{d}^{2} K\right)^{-1} K-I_{p}\right) G_{y y}\right]  \tag{B.19}\\
& \quad=-\frac{1}{2} \beta n_{d} \operatorname{tr}\left[\left(I_{p}+\beta \sigma_{d}^{2} K\right)^{-1} G_{y y}\right] \tag{B.20}
\end{align*}
$$

where we have defined the normalized Gram matrix of the outputs

$$
\begin{equation*}
G_{y y} \equiv \frac{1}{n_{d}} Y Y^{\top} \tag{B.21}
\end{equation*}
$$

and noticed that

$$
\begin{equation*}
I_{p}-\beta \sigma_{d}^{2}\left(I_{p}+\beta \sigma_{d}^{2} K\right)^{-1} K=\left(I_{p}+\beta \sigma_{d}^{2} K\right)^{-1} \tag{B.22}
\end{equation*}
$$

Therefore, we conclude that

$$
\begin{equation*}
Z=\mathbb{E}_{\mathcal{W}} \exp \left[\operatorname{tr}\left(J^{\top} O\right)-\frac{n_{d}}{2}\left(\beta \operatorname{tr}\left[\left(I_{p}+\beta \sigma_{d}^{2} K\right)^{-1} G_{y y}\right]+\log \operatorname{det}\left(I_{p}+\beta \sigma_{d}^{2} K\right)\right)\right] \tag{B.23}
\end{equation*}
$$

at any width.

## B. 2 Perturbative expansion

We now consider how this expression behaves in the large-width limit. We assume that this limit is well-defined in the sense that the readout kernel $K$ tends in probability to the constant GP kernel $K_{\infty}$ [4-7], and that the observable $O$ similarly tends to a deterministic limit $O_{\infty}$. Then, we formally write $K$ and $O$ as their infinite-width limits plus corrections which are small at large hidden layer widths:

$$
\begin{align*}
K & =K_{\infty}+\lambda \delta K,  \tag{B.24}\\
O & =O_{\infty}+\lambda \delta O, \tag{B.25}
\end{align*}
$$

where the parameter $\lambda$ is used to track powers of the small deviations.
We first expand the term resulting from integrating out the readout layer into its infinite-width limit and a finite-width correction. We define the constant matrix

$$
\begin{equation*}
\Gamma \equiv K_{\infty}+\frac{1}{\beta \sigma_{d}^{2}} I_{p} \tag{B.26}
\end{equation*}
$$

which is invertible at any finite temperature. Then, by the Woodbury identity [3], we have,

$$
\begin{equation*}
\beta \sigma_{d}^{2}\left(I_{p}+\beta \sigma_{d}^{2} K\right)^{-1}=(\Gamma+\lambda \delta K)^{-1}=\Gamma^{-1}-\lambda \Gamma^{-1} \delta K(\Gamma+\lambda \delta K)^{-1} \tag{B.27}
\end{equation*}
$$

and, similarly,

$$
\begin{equation*}
\log \operatorname{det}\left(I_{p}+\beta \sigma_{d}^{2} K\right)=\log \operatorname{det}\left(\beta \sigma_{d}^{2} \Gamma\right)+\log \operatorname{det}\left(I_{p}+\lambda \Gamma^{-1} \delta K\right) \tag{B.28}
\end{equation*}
$$

Noting that that both $\lambda \Gamma^{-1} \delta K(\Gamma+\lambda \delta K)^{-1}$ and $\log \operatorname{det}\left(I_{p}+\lambda \Gamma^{-1} \delta K\right)$ are $\mathcal{O}(\lambda)$, we expand the logarithm of the partition function as

$$
\begin{equation*}
\log Z=\log Z_{\infty}+\operatorname{tr}\left(J^{\top} O_{\infty}\right)+\log \mathbb{E}_{\mathcal{W}} \exp \left[\lambda \operatorname{tr}\left(J^{\top} \delta O\right)+\lambda \Omega\right] \tag{B.29}
\end{equation*}
$$

where

$$
\begin{equation*}
Z_{\infty} \equiv \operatorname{det}\left(\beta \sigma_{d}^{2} \Gamma\right)^{-n_{d} / 2} \exp \left(-\frac{1}{2} n_{d} \sigma_{d}^{-2} \operatorname{tr}\left(\Gamma^{-1} G_{y y}\right)\right) \tag{B.30}
\end{equation*}
$$

is the GP partition function and

$$
\begin{equation*}
\Omega \equiv \frac{1}{2} n_{d} \operatorname{tr}\left[\sigma_{d}^{-2} \Gamma^{-1} \delta K(\Gamma+\lambda \delta K)^{-1} G_{y y}\right]-\frac{1}{2} n_{d} \lambda^{-1} \log \operatorname{det}\left(I_{p}+\lambda \Gamma^{-1} \delta K\right) \tag{B.31}
\end{equation*}
$$

is the remainder. $\log \mathbb{E}_{\mathcal{W}} \exp \left[\lambda \operatorname{tr}\left(J^{\top} \delta O\right)+\lambda \Omega\right]$ has the form of a cumulant generating function, hence it has a formal series expansion in $\lambda$ given by

$$
\begin{align*}
\log \mathbb{E}_{\mathcal{W}} \exp \left[\lambda \operatorname{tr}\left(J^{\top} \delta O\right)+\lambda \Omega\right]= & \lambda \mathbb{E}_{\mathcal{W}}\left[\operatorname{tr}\left(J^{\top} \delta O\right)+\Omega\right] \\
& \left.+\frac{1}{2} \lambda^{2} \mathbb{E}_{\mathcal{W}}\left\{\operatorname{tr}\left[J^{\top}\left(\delta O-\mathbb{E}_{\mathcal{W}} \delta O\right)\right]+\Omega-\mathbb{E}_{\mathcal{W}} \Omega\right]\right\}^{2} \\
& \left.+\frac{1}{6} \lambda^{3} \mathbb{E}_{\mathcal{W}}\left\{\operatorname{tr}\left[J^{\top}\left(\delta O-\mathbb{E}_{\mathcal{W}} \delta O\right)\right]+\Omega-\mathbb{E}_{\mathcal{W}} \Omega\right]\right\}^{3} \\
& +\mathcal{O}\left(\lambda^{4}\right) \tag{B.32}
\end{align*}
$$

We can then see that the $k$-th cumulant is $\mathcal{O}\left(J^{k}\right)$, hence the $k$-th posterior cumulant of $O$ will be $\mathcal{O}\left(\lambda^{k}\right)$. Specifically, we can read off the posterior mean

$$
\begin{equation*}
\langle O\rangle=O_{\infty}+\lambda \mathbb{E}_{\mathcal{W}} \delta O+\lambda^{2} \mathbb{K}_{\mathcal{W}}(\delta O, \Omega)+\frac{1}{2} \lambda^{3} \mathbb{K}_{\mathcal{W}}(\delta O, \Omega, \Omega)+\mathcal{O}\left(\lambda^{4}\right) \tag{B.33}
\end{equation*}
$$

and covariance

$$
\begin{equation*}
\operatorname{cov}\left(O_{\rho \gamma}, O_{\omega \chi}\right)=\lambda^{2} \mathbb{K}_{\mathcal{W}}\left(\delta O_{\rho \gamma}, \delta O_{\omega \chi}\right)+\lambda^{3} \mathbb{K}_{\mathcal{W}}\left(\delta O_{\rho \gamma}, \delta O_{\omega \chi}, \Omega\right)+\mathcal{O}\left(\lambda^{4}\right) \tag{B.34}
\end{equation*}
$$

To make further progress, we expand $\Omega$ in powers of $\lambda$. Using the Neumann series for the matrix inverse (see Appendix A), we have

$$
\begin{equation*}
(\Gamma+\lambda \delta K)^{-1}=\Gamma^{-1}-\lambda \Gamma^{-1} \delta K \Gamma^{-1}+\mathcal{O}\left(\lambda^{2}\right) \tag{B.35}
\end{equation*}
$$

and, using the series expansion of the log-determinant near the identity (see Appendix A), we have

$$
\begin{equation*}
\lambda^{-1} \log \operatorname{det}\left(I_{p}+\lambda \Gamma^{-1} \delta K\right)=\operatorname{tr}\left(\Gamma^{-1} \delta K\right)-\frac{1}{2} \lambda \operatorname{tr}\left(\Gamma^{-1} \delta K \Gamma^{-1} \delta K\right)+\mathcal{O}\left(\lambda^{2}\right) \tag{B.36}
\end{equation*}
$$

This yields

$$
\begin{align*}
\Omega= & \frac{n_{d}}{2} \operatorname{tr}\left[\left(\sigma_{d}^{-2} \Gamma^{-1} G_{y y} \Gamma^{-1}-\Gamma^{-1}\right) \delta K\right] \\
& -\frac{n_{d}}{2} \lambda \operatorname{tr}\left[\left(\sigma_{d}^{-2} \Gamma^{-1} G_{y y} \Gamma^{-1}-\frac{1}{2} \Gamma^{-1}\right) \delta K \Gamma^{-1} \delta K\right] \\
& +\mathcal{O}\left(\lambda^{2}\right) . \tag{B.37}
\end{align*}
$$

The leading term is simple because it is linear in $\delta K$. Then, keeping only the leading non-trivial corrections and recognizing that

$$
\begin{align*}
O_{\infty}+\lambda \mathbb{E}_{\mathcal{W}} \delta O & =\mathbb{E}_{\mathcal{W}} O  \tag{B.38}\\
\lambda^{2} \mathbb{K}_{\mathcal{W}}\left(\delta O, \delta K_{\mu \nu}\right) & =\mathbb{K}_{\mathcal{W}}\left(O, K_{\mu \nu}\right),  \tag{B.39}\\
\lambda^{2} \mathbb{K}_{\mathcal{W}}\left(\delta O_{\rho \gamma}, \delta O_{\omega \chi}\right) & =\mathbb{K}_{\mathcal{W}}\left(O_{\rho \gamma}, O_{\omega \chi}\right), \quad \text { and }  \tag{B.40}\\
\lambda^{3} \mathbb{K}_{\mathcal{W}}\left(\delta O_{\rho \gamma}, \delta O_{\omega \chi}, \delta K_{\mu \nu}\right) & =\mathbb{K}_{\mathcal{W}}\left(O_{\rho \gamma}, O_{\omega \chi}, K_{\mu \nu}\right), \tag{B.41}
\end{align*}
$$

we have

$$
\begin{equation*}
\langle O\rangle=\mathbb{E}_{\mathcal{W}} O+\frac{1}{2} n_{d} \sum_{\mu, \nu=1}^{p}\left(\sigma_{d}^{-2} \Gamma^{-1} G_{y y} \Gamma^{-1}-\Gamma^{-1}\right)_{\mu \nu} \mathbb{K}_{\mathcal{W}}\left(O, K_{\mu \nu}\right)+\mathcal{O}\left(\lambda^{3}\right) \tag{B.42}
\end{equation*}
$$

and

$$
\begin{align*}
\operatorname{cov}\left(O_{\rho \gamma}, O_{\omega \chi}\right)= & \mathbb{K}_{\mathcal{W}}\left(O_{\rho \gamma}, O_{\omega \chi}\right) \\
& +\frac{1}{2} n_{d} \sum_{\mu, \nu=1}^{p}\left(\sigma_{d}^{-2} \Gamma^{-1} G_{y y} \Gamma^{-1}-\Gamma^{-1}\right)_{\mu \nu} \mathbb{K}_{\mathcal{W}}\left(O_{\rho \gamma}, O_{\omega \chi}, K_{\mu \nu}\right) \\
& +\mathcal{O}\left(\lambda^{4}\right) \tag{B.43}
\end{align*}
$$

Restoring the layer indices to $K=K^{(d-1)}$, the above result for $\langle O\rangle$ yields the expression given in the main text. From the structure of these expressions, we can see that higher-order terms (in $\lambda$ ) will involve higher joint cumulants of the kernel deviations $\delta K^{(\ell)}$, which can in turn be converted into joint cumulants of the kernels $K^{(\ell)}$. Therefore, to show that the perturbative expansion yields a valid asymptotic series, one would need to show that these joint cumulants themselves have asymptotic series expansions at large width, with leading terms that are successively suppressed by powers of $n^{-1}$.

## C Explicit covariance computations in deep linear networks

In this appendix, we detail how to compute the prior covariances appearing in (B.42) for the hidden layer kernels of deep linear fully-connected and convolutional networks.

## C. 1 Fully-connected linear networks

In this brief subsection, we provide a self-contained derivation of the behavior of the prior cumulants of the kernels of a deep fully-connected linear network with no bias terms. This is a special case of Yaida [8]'s results, and provides some intuition for his results on general MLPs. As in the main text, we consider a network with activations $\mathbf{h}^{(\ell)} \in \mathbb{R}^{n_{\ell}}$ recursively defined as

$$
\begin{equation*}
\mathbf{h}^{(\ell)}=n_{\ell-1}^{-1 / 2} W^{(\ell)} \mathbf{h}^{(\ell-1)} \quad(\ell=1, \ldots, d) \tag{C.1}
\end{equation*}
$$

with base case $\mathbf{h}^{(0)}=\mathbf{x}$. We take the prior distribution over weights to be $\left[W^{(\ell)}\right]_{i j} \sim_{\text {i.i.d. }} \mathcal{N}\left(0, \sigma_{\ell}^{(2)}\right)$, and define the hidden layer kernels $\left[K^{(\ell)}\right]_{\mu \nu} \equiv n_{\ell}^{-1} \mathbf{h}_{\mu}^{(\ell)} \cdot \mathbf{h}_{\nu}^{(\ell)}$ for $\ell=1, \ldots, d-1$. Then, we have

$$
\begin{align*}
\mathbb{E}_{\mathcal{W}} K_{\mu \nu}^{(\ell)} & =\frac{1}{n_{\ell} \cdots n_{0}} \mathbb{E}_{\mathcal{W}} \mathbf{x}_{\mu}^{\top}\left(W^{(1)}\right)^{\top} \cdots\left(W^{(\ell)}\right)^{\top} W^{(\ell)} \cdots W^{(1)} \mathbf{x}_{\nu}  \tag{C.2}\\
& =\sigma_{1}^{2} \cdots \sigma_{\ell}^{2} \frac{\mathbf{x}_{\mu} \cdot \mathbf{x}_{\nu}}{n_{0}}  \tag{C.3}\\
& =\left[K_{\infty}^{(\ell)}\right]_{\mu \nu} \tag{C.4}
\end{align*}
$$

at any width, as $\mathbb{E}_{W^{(\ell)}}\left(W^{(\ell)}\right)^{\top} W^{(\ell)} / n_{\ell}=\sigma_{\ell}^{2} I_{n_{\ell-1}}$. We now consider the second moments of the kernels. We first note that

$$
\begin{equation*}
\mathbb{E}_{\mathcal{W}} K_{\mu \nu}^{(\ell)} K_{\rho \lambda}^{(\ell+\tau)}=\sigma_{\ell+\tau}^{2} \cdots \sigma_{\ell+1}^{2} \mathbb{E}_{\mathcal{W}} K_{\mu \nu}^{(\ell)} K_{\rho \lambda}^{(\ell)} \tag{C.5}
\end{equation*}
$$

for any $\tau \geq 1$. By Isserlis' theorem (see Appendix A), we have

$$
\begin{equation*}
\mathbb{E}_{W^{(\ell)}} W_{i k}^{(\ell)} W_{i l}^{(\ell)} W_{j m}^{(\ell)} W_{j r}^{(\ell)}=\sigma_{\ell}^{4} \delta_{i j}\left(\delta_{k m} \delta_{l r}+\delta_{k r} \delta_{l m}\right)+\sigma_{\ell}^{4} \delta_{k l} \delta_{m r}, \tag{C.6}
\end{equation*}
$$

hence we have the exact recursion

$$
\begin{align*}
\mathbb{E}_{\mathcal{W}} K_{\mu \nu}^{(\ell)} K_{\rho \lambda}^{(\ell)}=\frac{1}{\left(n_{\ell} \cdots n_{0}\right)^{2}} \mathbb{E}_{\mathcal{W}} \sum_{i, j=1}^{n_{\ell}} & \sum_{k, l, m, r=1}^{n_{\ell-1}} W_{i k}^{(\ell)} W_{i l}^{(\ell)} W_{j m}^{(\ell)} W_{j r}^{(\ell)} \\
& \times\left[W^{(\ell-1)} \cdots W^{(1)} \mathbf{x}_{\mu}\right]_{k}\left[W^{(\ell-1)} \cdots W^{(1)} \mathbf{x}_{\nu}\right]_{l} \\
& \times\left[W^{(\ell-1)} \cdots W^{(1)} \mathbf{x}_{\rho}\right]_{m}\left[W^{(\ell-1)} \cdots W^{(1)} \mathbf{x}_{\lambda}\right]_{r}  \tag{C.7}\\
=\sigma_{\ell}^{4} \mathbb{E}_{\mathcal{W}} K_{\mu \nu}^{(\ell-1)} K_{\rho \lambda}^{(\ell-1)} & +\frac{1}{n_{\ell}} \sigma_{\ell}^{4}\left(\mathbb{E}_{\mathcal{W}} K_{\mu \rho}^{(\ell-1)} K_{\nu \lambda}^{(\ell-1)}+\mathbb{E}_{\mathcal{W}} K_{\mu \lambda}^{(\ell-1)} K_{\nu \rho}^{(\ell-1)}\right) \tag{C.8}
\end{align*}
$$

with base case

$$
\begin{align*}
\mathbb{E}_{\mathcal{W}} K_{\mu \nu}^{(1)} K_{\rho \lambda}^{(1)} & =\frac{1}{\left(n_{1} n_{0}\right)^{2}} \sum_{i, j=1}^{n_{1}} \sum_{k, l, m, r=1}^{n_{0}} \mathbb{E}_{\mathcal{W}} W_{i k}^{(1)} W_{i l}^{(1)} W_{j m}^{(1)} W_{j r}^{(1)} x_{\mu, k} x_{\nu, l} x_{\rho, m} x_{\lambda, r}  \tag{C.9}\\
& =\sigma_{1}^{4} \frac{\mathbf{x}_{\mu} \cdot \mathbf{x}_{\nu}}{n_{0}} \frac{\mathbf{x}_{\rho} \cdot \mathbf{x}_{\lambda}}{n_{0}}+\frac{1}{n_{1}} \sigma_{1}^{4}\left(\frac{\mathbf{x}_{\mu} \cdot \mathbf{x}_{\rho}}{n_{0}} \frac{\mathbf{x}_{\nu} \cdot \mathbf{x}_{\lambda}}{n_{0}}+\frac{\mathbf{x}_{\mu} \cdot \mathbf{x}_{\lambda}}{n_{0}} \frac{\mathbf{x}_{\nu} \cdot \mathbf{x}_{\rho}}{n_{0}}\right)  \tag{C.10}\\
& =\left[K_{\infty}^{(1)}\right]_{\mu \nu}\left[K_{\infty}^{(1)}\right]_{\rho \lambda}+\frac{1}{n_{1}}\left(\left[K_{\infty}^{(1)}\right]_{\mu \rho}\left[K_{\infty}^{(1)}\right]_{\nu \lambda}+\left[K_{\infty}^{(1)}\right]_{\mu \lambda}\left[K_{\infty}^{(1)}\right]_{\nu \rho}\right) \tag{C.11}
\end{align*}
$$

for the second moments of the kernels at each layer. This recurrence relation is in principle exactly solvable for any finite width, but we are interested only in its leading-order behavior at large widths. In particular, we can read off that

$$
\begin{align*}
\operatorname{cov}_{\mathcal{W}}\left(K_{\mu \nu}^{(\ell)}, K_{\rho \lambda}^{(\ell+\tau)}\right)= & \sigma_{\ell+\tau}^{2} \cdots \sigma_{\ell+1}^{2}\left(\sum_{\ell^{\prime}=1}^{\ell} \frac{1}{n_{\ell^{\prime}}}\right)\left(\left[K_{\infty}^{(\ell)}\right]_{\mu \rho}\left[K_{\infty}^{(\ell)}\right]_{\nu \lambda}+\left[K_{\infty}^{(\ell)}\right]_{\mu \lambda}\left[K_{\infty}^{(\ell)}\right]_{\nu \rho}\right) \\
& +\mathcal{O}\left(n^{-2}\right) \tag{C.12}
\end{align*}
$$

Moreover, one can see by Isserlis' theorem that the third and higher cumulants will be $\mathcal{O}\left(n^{-2}\right)$. Substituting this result into (B.42) with the hidden layer kernel as the observable of interest, we obtain the expression (8) given in the main text.

## C. 2 Convolutional linear networks

In this subsection, we derive the prior cumulants required to compute corrections to the average feature kernels of deep convolutional linear networks. As described in the main text, following the setup of Novak et al. [6] and Xiao et al. [9], we consider a network consisting of $d-1$ linear convolutional layers followed by a fully-connected linear readout layer. For simplicity, we assume circular padding and no internal pooling. As discussed in Novak et al. [6], this setup could be easily extended to other padding strategies, strided convolutions, and average pooling in intermediate layers.
We write the activations at the $\ell$-th hidden layer as $h_{i, \mathfrak{a}}^{(\ell)}$, where $i$ indexes the $n_{\ell}$ channels of the layer and $\mathfrak{a}$ is a $q$-dimensional spatial multi-index. We take the filters to be of size $(2 k+1) \times \cdots \times(2 k+1)$ in all convolutional layers; the extension to differently-sized filters would be straightforward but notationally cumbersome. The ranges of all spatial summations will be implied.

The hidden layer activations are then defined through the recurrence

$$
\begin{equation*}
h_{i, \mathfrak{a}}^{(\ell)}(x)=\frac{1}{\sqrt{n_{\ell-1}}} \sum_{j=1}^{n_{\ell-1}} \sum_{\mathfrak{b}} w_{i j, \mathfrak{b}}^{(\ell)} h_{j, \mathfrak{a}+\mathfrak{b}}^{(\ell-1)}(x) \tag{C.13}
\end{equation*}
$$

with base case $h_{i, \mathfrak{a}}^{(0)}(x)=x_{i, \mathfrak{a}}$. We fix the prior distribution of the filter elements to be

$$
\begin{equation*}
w_{i j, \mathfrak{a}}^{(\ell)} \underset{\text { i.i.d. }}{\sim} \mathcal{N}\left(0, \sigma_{\ell}^{2} v_{\mathfrak{a}}\right), \tag{C.14}
\end{equation*}
$$

where $v_{\mathfrak{a}}>0$ is a weighting factor that sets the fraction of receptive field variance at location $\mathfrak{a}$ (and is thus subject to the constraint $\sum_{\mathfrak{a}} v_{\mathfrak{a}}=1$ ). For inputs $\left[x_{\mu}\right]_{i, \mathfrak{a}}$ and $\left[x_{\nu}\right]_{i, \mathfrak{a}}$, we introduce the hidden layer kernels

$$
\begin{equation*}
K_{\mu \nu, \mathfrak{a b}}^{(\ell)} \equiv \frac{1}{n_{\ell}} \sum_{i=1}^{n_{\ell}} h_{i, \mathfrak{a}}^{(\ell)}\left(x_{\mu}\right) h_{i, \mathfrak{b}}^{(\ell)}\left(x_{\nu}\right) \tag{C.15}
\end{equation*}
$$

We will first compute the prior mean and covariance of these four-indexed kernels, and then address how to handle readout across space.
As shown by Xiao et al. [9], the prior mean obeys the recurrence

$$
\begin{align*}
& \mathbb{E}_{\mathcal{W}} K_{\mu \nu, \mathfrak{a b}}^{(\ell)} \\
& \quad=\mathbb{E}_{W^{(1)} \ldots W^{(\ell-1)}} \frac{1}{n_{\ell} n_{\ell-1}} \sum_{i=1}^{n_{\ell}} \sum_{j, j^{\prime}=1}^{n_{\ell-1}} \sum_{\mathfrak{c}, \mathfrak{d}} h_{j, \mathfrak{a}+\mathfrak{c}}^{(\ell-1)}\left(x_{\mu}\right) h_{j^{\prime}, \mathfrak{b}+\mathfrak{d}}^{(\ell-1)}\left(x_{\nu}\right) \mathbb{E}_{W^{(\ell)}} w_{i j, \mathfrak{c}}^{(\ell)} w_{i j^{\prime}, \mathfrak{d}}^{(\ell)}  \tag{C.16}\\
& \quad=\sigma_{\ell}^{2} \mathbb{E}_{W^{(1)} \ldots W^{(\ell-1)}} \sum_{\mathfrak{c}} v_{\mathfrak{c}} \frac{1}{n_{\ell-1}} \sum_{j=1}^{n_{\ell-1}} h_{j, \mathfrak{a}+\mathfrak{c}}^{(\ell-1)}\left(x_{\mu}\right) h_{j, \mathfrak{b}+\mathfrak{c}}^{(\ell-1)}\left(x_{\nu}\right)  \tag{C.17}\\
& \quad=\sigma_{\ell}^{2} \sum_{\mathfrak{c}} v_{\mathfrak{c}} \mathbb{E}_{\mathcal{W}} K_{\mu \nu,(\mathfrak{a}+\mathfrak{c})(\mathfrak{b}+\mathfrak{c})}^{(\ell-1)} \tag{C.18}
\end{align*}
$$

with base case

$$
\begin{equation*}
\mathbb{E}_{\mathcal{W}} K_{\mu \nu, \mathfrak{a b}}^{(1)}=\sigma_{1}^{2} \sum_{\mathfrak{c}} v_{\mathfrak{c}}\left[G_{x x}\right]_{\mu \nu,(\mathfrak{a}+\mathfrak{c})(\mathfrak{b}+\mathfrak{c})} \tag{C.19}
\end{equation*}
$$

for

$$
\begin{equation*}
\left[G_{x x}\right]_{\mu \nu, \mathfrak{a} \mathfrak{b}} \equiv \frac{1}{n_{0}} \sum_{i=1}^{n_{0}}\left[x_{\mu}\right]_{i, \mathfrak{a}}\left[x_{\nu}\right]_{i, \mathfrak{b}} \tag{C.20}
\end{equation*}
$$

This recurrence yields

$$
\begin{equation*}
\mathbb{E}_{\mathcal{W}} K_{\mu \nu, \mathfrak{a b}}^{(\ell)}=\sigma_{1}^{2} \cdots \sigma_{\ell}^{2} \sum_{\mathfrak{c}_{1}, \ldots, \mathfrak{c}_{\ell}} v_{\mathfrak{c}_{1}} \cdots v_{\mathfrak{c}_{\ell}}\left[G_{x x}\right]_{\mu \nu,\left(\mathfrak{a}+\mathfrak{c}_{1}+\cdots+\mathfrak{c}_{\ell}\right)\left(\mathfrak{b}+\mathfrak{c}_{1}+\cdots+\mathfrak{c}_{\ell}\right)} . \tag{C.21}
\end{equation*}
$$

Moreover, as in the fully-connected case considered in the preceding section, we have

$$
\begin{equation*}
\left[K_{\infty}^{(\ell)}\right]_{\mu \nu, \mathfrak{a b}}=\mathbb{E}_{\mathcal{W}} K_{\mu \nu, \mathfrak{a b}}^{(\ell)} \tag{C.22}
\end{equation*}
$$

at any width.
We now consider the prior covariance of the kernels of two different hidden layers $\ell$ and $\ell+\tau$. As the weight prior factors across layers, we have

$$
\begin{align*}
& \mathbb{E}_{\mathcal{W}} K_{\mu \nu, \mathfrak{a b}}^{(\ell)} K_{\rho \lambda, \mathfrak{c} \mathfrak{d}}^{(\ell+\tau)}=\sigma_{\ell+1}^{2} \cdots \sigma_{\ell+\tau}^{2} \sum_{\mathfrak{e}_{1}, \ldots, \mathfrak{e}_{\tau}} \\
& v_{\mathfrak{e}_{1}} \cdots v_{\mathfrak{e}_{\tau}}  \tag{C.23}\\
& \times \mathbb{E}_{\mathcal{W}} K_{\mu \nu, \mathfrak{a b}}^{(\ell)} K_{\rho \lambda,\left(\mathfrak{c}+\mathfrak{e}_{1}+\cdots+\mathfrak{e}_{\tau}\right)\left(\mathfrak{d}+\mathfrak{e}_{1}+\cdots+\mathfrak{e}_{\tau}\right)}^{(\ell)}
\end{align*}
$$

By Isserlis' theorem (see Appendix A),

$$
\begin{align*}
\mathbb{E}_{W^{(\ell)}} w_{i j, \mathfrak{e}}^{(\ell)} w_{i j^{\prime}, \mathfrak{f}}^{(\ell)} w_{i^{\prime} j^{\prime \prime}, \mathfrak{g}}^{(\ell)} w_{i^{\prime} j^{\prime \prime \prime}, \mathfrak{h}}^{(\ell)}= & \sigma_{\ell}^{4} v_{\mathfrak{e}} v_{\mathfrak{g}} \delta_{j j^{\prime}} \delta_{j^{\prime \prime} j^{\prime \prime \prime}} \delta_{\mathfrak{e} \mathfrak{f}} \delta_{\mathfrak{g h}} \\
& +\sigma_{\ell}^{4} v_{\mathfrak{e}} v_{\mathfrak{f}} \delta_{i i^{\prime}} \delta_{j j^{\prime \prime}} \delta_{j^{\prime} j^{\prime \prime \prime}} \delta_{\mathfrak{e g}} \delta_{\mathfrak{f h}} \\
& +\sigma_{\ell}^{4} v_{\mathfrak{e}} v_{\mathfrak{f}} \delta_{i i^{\prime}} \delta_{j j^{\prime \prime \prime}} \delta_{j^{\prime} j^{\prime \prime}} \delta_{\mathfrak{e h}} \delta_{\mathfrak{f g}} \tag{C.24}
\end{align*}
$$

hence we have the recurrence

$$
\begin{align*}
& \mathbb{E}_{\mathcal{W}} K_{\mu \nu, \mathfrak{a b}}^{(\ell)} K_{\rho \lambda, \mathfrak{c o d}}^{(\ell)} \\
&=\mathbb{E}_{W^{(1)} \ldots W^{(\ell-1)}} \frac{1}{n_{\ell}^{2} n_{\ell-1}^{2}} \sum_{i, i^{\prime}=1}^{n_{\ell}} \sum_{j, j^{\prime}, j^{\prime \prime}, j^{\prime \prime \prime}=1}^{n_{\ell-1}} \sum_{\mathfrak{e}, \mathfrak{f}, \mathfrak{g}, \mathfrak{h}} \\
& \times h_{j, \mathfrak{a}+\mathfrak{e}}^{(\ell-1)}\left(x_{\mu}\right) h_{j^{\prime}, \mathfrak{b}+\mathfrak{f}}^{(\ell-1)}\left(x_{\nu}\right) h_{j^{\prime \prime}, \mathfrak{c}+\mathfrak{g}}^{(\ell-1)}\left(x_{\rho}\right) h_{j^{\prime \prime \prime}, \mathfrak{d}+\mathfrak{h}}^{(\ell-1)}\left(x_{\lambda}\right) \\
& \times \mathbb{E}_{W^{(\ell)}} w_{i j, \mathfrak{e}}^{(\ell)} w_{i^{\prime}, \mathfrak{f}}^{(\ell)} w_{i^{\prime} j^{\prime \prime}, \mathfrak{g}}^{(\ell)} w_{i^{\prime} j^{\prime \prime \prime}, \mathfrak{h}}^{(\ell)}  \tag{C.25}\\
&=\sigma_{\ell}^{4} \sum_{\mathfrak{e}, \mathfrak{f}} v_{\mathfrak{e}} v_{\mathfrak{f}}\left[\mathbb{E}_{\mathcal{W}} K_{\mu \nu,(\mathfrak{a}+\mathfrak{e})(\mathfrak{b}+\mathfrak{e})}^{(\ell-1)} K_{\rho \lambda,(\mathfrak{c}+\mathfrak{f})(\mathfrak{d}+\mathfrak{f})}^{(\ell-1)}\right. \\
&+\frac{1}{n_{\ell}} \mathbb{E}_{\mathcal{W}} K_{\mu \rho,(\mathfrak{a}+\mathfrak{e})(\mathfrak{c}+\mathfrak{e})}^{(\ell-1)} K_{\nu \lambda,(\mathfrak{b}+\mathfrak{f})(\mathfrak{d}+\mathfrak{f})}^{(\ell-1)} \\
&\left.+\frac{1}{n_{\ell}} \mathbb{E}_{\mathcal{W}} K_{\mu \lambda,(\mathfrak{a}+\mathfrak{e})(\mathfrak{d}+\mathfrak{e})}^{(\ell-1)} K_{\nu \rho,(\mathfrak{a}+\mathfrak{f})(\mathfrak{c}+\mathfrak{f})}^{(\ell-1)}\right] \tag{C.26}
\end{align*}
$$

with base case

$$
\begin{align*}
\mathbb{E}_{\mathcal{W}} K_{\mu \nu, \mathfrak{a b}}^{(1)} K_{\rho \lambda, \mathfrak{e d}}^{(1)}= & \sigma_{1}^{4} \sum_{\mathfrak{e}, \mathfrak{f}} v_{\mathfrak{e}} v_{\mathfrak{f}}\left[\left[G_{x x}\right]_{\mu \nu,(\mathfrak{a}+\mathfrak{e})(\mathfrak{b}+\mathfrak{e})}\left[G_{x x}\right]_{\rho \lambda,(\mathfrak{c}+\mathfrak{f})(\mathfrak{d}+\mathfrak{f})}\right. \\
& +\frac{1}{n_{\ell}}\left[G_{x x}\right]_{\mu \rho,(\mathfrak{a}+\mathfrak{e})(\mathfrak{c}+\mathfrak{e})}\left[G_{x x}\right]_{\nu \lambda,(\mathfrak{b}+\mathfrak{f})(\mathfrak{d}+\mathfrak{f})} \\
& \left.+\frac{1}{n_{\ell}}\left[G_{x x}\right]_{\mu \lambda,(\mathfrak{a}+\mathfrak{e})(\mathfrak{d}+\mathfrak{e})}\left[G_{x x}\right]_{\nu \rho,(\mathfrak{a}+\mathfrak{f})(\mathfrak{c}+\mathfrak{f})}\right]  \tag{C.27}\\
= & {\left[K_{\infty}^{(1)}\right]_{\mu \nu, \mathfrak{a b}}\left[K_{\infty}^{(1)}\right]_{\rho \lambda, \mathfrak{c d}} } \\
& +\frac{1}{n_{\ell}}\left[\left[K_{\infty}^{(1)}\right]_{\mu \rho, \mathfrak{a c}}\left[K_{\infty}^{(1)}\right]_{\nu \lambda, \mathfrak{b d}}+\left[K_{\infty}^{(1)}\right]_{\mu \lambda, \mathfrak{a d}}\left[K_{\infty}^{(1)}\right]_{\nu \rho, \mathfrak{b c}}\right] \tag{C.28}
\end{align*}
$$

for the second prior moments of the kernels. As in the fully-connected case, these recurrence relations could in principle be solved exactly, but we are only interested in their large-width behavior. Using the forward recurrence for the GP kernels, we can easily read off that

$$
\begin{align*}
\operatorname{cov}_{\mathcal{W}}\left(K_{\mu \nu, \mathfrak{a b}}^{(\ell)}, K_{\rho \lambda, \mathfrak{c o}}^{(\ell)}\right)= & \left(\sum_{\ell^{\prime}=1}^{\ell} \frac{1}{n_{\ell^{\prime}}}\right)\left(\left[K_{\infty}^{(\ell)}\right]_{\mu \rho, \mathfrak{a c}}\left[K_{\infty}^{(\ell)}\right]_{\nu \lambda, \mathfrak{b d}}+\left[K_{\infty}^{(\ell)}\right]_{\mu \lambda, \mathfrak{a d}}\left[K_{\infty}^{(\ell)}\right]_{\nu \rho, \mathfrak{b c}}\right) \\
& +\mathcal{O}\left(n^{-2}\right) \tag{C.29}
\end{align*}
$$

which can then be substituted into the desired cross-layer covariance:

$$
\begin{align*}
\operatorname{cov}_{\mathcal{W}}\left(K_{\mu \nu, \mathfrak{a b}}^{(\ell)}, K_{\rho \lambda, \mathfrak{c}}^{(\ell+\tau)}\right)=\sigma_{\ell+1}^{2} \cdots \sigma_{\ell+\tau}^{2} & \sum_{\mathfrak{e}_{1}, \ldots, \mathfrak{e}_{\tau}} \\
& v_{\mathfrak{e}_{1}} \cdots v_{\mathfrak{e}_{\tau}}  \tag{C.30}\\
& \times \operatorname{cov}_{\mathcal{W}}\left(K_{\mu \nu, \mathfrak{a b}}^{(\ell)}, K_{\rho \lambda,\left(\mathfrak{c}+\mathfrak{e}_{1}+\cdots+\mathfrak{e}_{\tau}\right)\left(\mathfrak{d}+\mathfrak{e}_{1}+\cdots+\mathfrak{e}_{\tau}\right)}^{(\ell)}\right) .
\end{align*}
$$

We now address the question of how to read out the convolutional layer activities across space. Following Novak et al. [6], we consider two strategies: vectorization and projection. With vectorization, the output of the final convolutional layer is flattened into a $n_{d-1} s$-dimensional vector before readout, i.e., $\psi_{i+s(\mathfrak{a}-1)}(x)=h_{i, \mathfrak{a}}^{(d-1)}(x)$ or $\psi_{n_{d}(i-1)+\mathfrak{a}}(x)=h_{i, \mathfrak{a}}^{(d-1)}(x)$. The two-index feature map kernel appearing in Conjecture 1 is then related to the four-index convolutional hidden layer kernel analyzed above via

$$
\begin{equation*}
K_{\mu \nu}^{(d-1)}=\frac{1}{s} \sum_{\mathfrak{a}} K_{\mu \nu, \mathfrak{a a}}^{(d-1)} . \tag{C.31}
\end{equation*}
$$

With projection, the feature map is formed by contracting the final convolutional layer with a fixed vector u, i.e.,

$$
\begin{equation*}
\psi_{i}(x)=\sum_{\mathfrak{a}} u_{\mathfrak{a}} h_{i, \mathfrak{a}}^{(d-1)}(x) \tag{C.32}
\end{equation*}
$$

The feature map kernel is then given as

$$
\begin{equation*}
K_{\mu \nu}^{(d-1)}=\sum_{\mathfrak{a}, \mathfrak{b}} u_{\mathfrak{a}} u_{\mathfrak{b}} K_{\mu \nu, \mathfrak{a b}}^{(d-1)} \tag{C.33}
\end{equation*}
$$

Examples of common projection readout strategies include global average pooling ( $u_{\mathfrak{a}}=1 / \mathrm{s}$ ) and single-pixel subsampling ( $u_{\mathfrak{a}}=\delta_{\mathfrak{a c}}$ for some desired location $\mathfrak{c}$ ). These readout approaches endow the network with differing properties under spatial transformations; global average pooling has the particular property of making the output translation-invariant.
We now seek to simplify the resulting expression for the leading-order correction to the posterior mean of some four-index feature kernel $K_{\mu \nu, \mathfrak{a b}}^{(\ell)}$. Per Conjecture 1, the general form of this correction is

$$
\begin{equation*}
\frac{1}{2} n_{d} \sum_{\rho, \lambda=1}^{p} \Phi_{\rho \lambda} \operatorname{cov}_{\mathcal{W}}\left(K_{\mu \nu, \mathfrak{a b}}^{(\ell)}, K_{\rho \lambda}^{(d-1)}\right) \tag{C.34}
\end{equation*}
$$

where we have defined $\Phi_{\rho \lambda}=\left[\sigma_{d}^{-2} \Gamma^{-1} G_{y y} \Gamma^{-1}-\Gamma^{-1}\right]_{\rho \lambda}$ for notational convenience. As elsewhere, $\Gamma \equiv K_{\infty}^{(d-1)}+\beta^{-1} \sigma_{d}^{-2} I_{p}$ for $K_{\infty}^{(d-1)}$ the two-index kernel determined by the chosen readout strategy. Depending on the chosen readout strategy, this general expression can be simplified dramatically. In particular, for vectorization or global average pooling, the correction does not depend on the particular form of $v_{\mathfrak{a}}$.
To show this for vectorization (the strategy used in our experiments), we substitute the definition of $K_{\infty}^{(d-1)}$ from (C.31) and the expression for the cross-layer kernel covariance from (C.30) into the general expression for the correction to obtain

$$
\begin{align*}
& \frac{n_{d}}{2 s} \sigma_{\ell+1}^{2} \cdots \sigma_{d-1}^{2} \\
& \times \sum_{\rho, \lambda} \Phi_{\rho \lambda} \sum_{\mathfrak{e}_{1}, \cdots, \mathfrak{e}_{d-\ell-1}} v_{\mathfrak{e}_{1}} \cdots v_{\mathfrak{e}_{d-\ell-1}} \sum_{\mathfrak{c}} \operatorname{cov}_{\mathcal{W}}\left(K_{\mu \nu, \mathfrak{a} \mathfrak{b}}^{(\ell)}, K_{\rho \lambda,\left(\mathfrak{c}+\mathfrak{e}_{1}+\cdots+\mathfrak{e}_{d-\ell-1}\right)\left(\mathfrak{c}+\mathfrak{e}_{1}+\cdots+\mathfrak{e}_{d-\ell-1}\right)}^{(\ell)}\right) . \tag{C.35}
\end{align*}
$$

Thanks to the periodic boundary conditions, the summation over $\mathfrak{c}$ is independent of the index shift $\mathfrak{e}_{1}+\cdots+\mathfrak{e}_{d-\ell-1}$. Then, the sums over $\mathfrak{e}_{1}, \cdots, \mathfrak{e}_{d-\ell-1}$ factor, yielding

$$
\begin{equation*}
\frac{n_{d}}{2 s} \sigma_{\ell+1}^{2} \cdots \sigma_{d-1}^{2} \sum_{\rho, \lambda=1}^{p} \Phi_{\rho \lambda} \sum_{\mathfrak{c}} \operatorname{cov}_{\mathcal{W}}\left(K_{\mu \nu, \mathfrak{a b}}^{(\ell)}, K_{\rho \lambda, \mathfrak{c c}}^{(\ell)}\right) \tag{C.36}
\end{equation*}
$$

thanks to the normalization constraint $\sum_{\mathfrak{e}} v_{\mathfrak{e}}=1$. We now notice that $\Phi_{\rho \lambda}$ is a symmetric matrix, and that the kernel remains invariant under the simultaneous exchange of indices $\rho \leftrightarrow \lambda$ and $\mathfrak{c} \leftrightarrow \mathfrak{d}$. Then, substituting in the expression for the same-layer kernel covariance (C.29), it is easy to show that the correction reduces to

$$
\begin{equation*}
\sigma_{\ell+1}^{2} \cdots \sigma_{d-1}^{2}\left(\sum_{\ell^{\prime}=1}^{\ell} \frac{n_{d}}{n_{\ell^{\prime}}}\right) \frac{1}{s} \sum_{\mathfrak{c}} \sum_{\rho, \lambda=1}^{p}\left[K_{\infty}^{(\ell)}\right]_{\mu \rho, \mathfrak{a c}} \Phi_{\rho \lambda}\left[K_{\infty}^{(\ell)}\right]_{\lambda \nu, \mathfrak{c b}} . \tag{C.37}
\end{equation*}
$$

This yields the expression given in the main text.
For projection, an analogous simplification is possible in the case of global average pooling ( $u_{\mathfrak{a}}=$ $1 / s$ ). Substituting the definition of $K_{\infty}^{(d-1)}$ from (C.33) and expression for the cross-layer kernel covariance (C.30) into the correction, we have

$$
\begin{equation*}
\frac{n_{d}}{2 s^{2}} \sigma_{\ell+1}^{2} \cdots \sigma_{d-1}^{2} \sum_{\rho, \lambda=1}^{p} \Phi_{\rho \lambda} \sum_{\mathfrak{c}, \mathfrak{d}} \operatorname{cov}_{\mathcal{W}}\left(K_{\mu \nu, \mathfrak{a b}}^{(\ell)}, K_{\rho \lambda, \mathfrak{c o}}^{(\ell)}\right) . \tag{C.38}
\end{equation*}
$$

Substituting in the expression for the same-layer kernel covariance (C.29), it is again easy to show that the correction reduces to

$$
\begin{equation*}
\sigma_{\ell+1}^{2} \cdots \sigma_{d-1}^{2}\left(\sum_{\ell^{\prime}=1}^{\ell} \frac{n_{d}}{n_{\ell^{\prime}}}\right) \frac{1}{s^{2}} \sum_{\mathfrak{c}, \mathfrak{d}} \sum_{\rho, \lambda=1}^{p}\left[K_{\infty}^{(\ell)}\right]_{\mu \rho, \mathfrak{a c}} \Phi_{\rho \lambda}\left[K_{\infty}^{(\ell)}\right]_{\lambda \nu, \mathfrak{d b}} . \tag{C.39}
\end{equation*}
$$

For projection strategies other than global average pooling (more precisely, for strategies for which $u_{\mathfrak{a}}$ is not constant), the sum over indices in the cross-layer covariance is not independent of the shift, hence we cannot simplify the correction in a similar fashion. This can be seen explicitly when treating the case of single-pixel subsampling ( $u_{\mathfrak{a}}=\delta_{\mathfrak{a} \mathfrak{c}}$ for some desired location $\mathfrak{c}$ ). In this case, the correction reduces to

$$
\begin{align*}
& \frac{n_{d}}{2} \sigma_{\ell+1}^{2} \cdots \sigma_{d-1}^{2} \\
& \times \sum_{\rho, \lambda} \Phi_{\rho \lambda} \sum_{\mathfrak{e}_{1}, \cdots, \mathfrak{e}_{d-\ell-1}} v_{\mathfrak{e}_{1}} \cdots v_{\mathfrak{e}_{d-\ell-1}} \operatorname{cov}_{\mathcal{W}}\left(K_{\mu \nu, \mathfrak{a b}}^{(\ell)}, K_{\rho \lambda,\left(\mathfrak{c}+\mathfrak{e}_{1}+\cdots+\mathfrak{e}_{d-\ell-1}\right)\left(\mathfrak{c}+\mathfrak{e}_{1}+\cdots+\mathfrak{e}_{d-\ell-1}\right)}^{(\ell)}\right) . \tag{C.40}
\end{align*}
$$

Unlike for vectorization or for projection using global average pooling, this expression is manifestly dependent on the form of $v_{\mathfrak{a}}$.

Naïvely, the computation of the corrections to the linear convolutional kernels requires the computa-
 tical amounts of compute time and storage. We only found it practical to compute the theoretical kernels in the special cases presented above.

## D Direct computation of the average hidden layer kernels of a deep linear MLP

In this appendix, we provide a self-contained derivation of the average hidden layer kernels of a deep linear fully-connected network (MLP). This derivation relies upon neither the results of Appendices $B$ and $C$ nor those of Yaida [8].

## D. 1 The cumulant generating function of learned features for a MLP

In this section, we briefly describe the full partition function of the Bayes posterior for a general fully connected network, or multi-layer perceptron (MLP), with no bias terms. An MLP $\mathbf{f}: \mathbb{R}^{n_{0}} \rightarrow \mathbb{R}^{n_{d}}$ with $d$ layers, no biases, and parameters $\Theta=\left\{W^{(\ell)}\right\}_{\ell=1}^{d}$ can be defined recursively in terms of its layer-wise preactivations $\mathbf{h}^{(\ell)} \in \mathbb{R}^{n_{\ell}}$ as

$$
\begin{align*}
\mathbf{h}^{(0)} & =\mathbf{x}  \tag{D.1}\\
\mathbf{h}^{(\ell)} & =\frac{1}{\sqrt{n_{\ell-1}}} W^{(\ell)} \phi_{\ell-1}\left(\mathbf{h}^{(\ell-1)}\right) \quad(\ell=1, \ldots, d),  \tag{D.2}\\
\mathbf{f} & =\phi_{d}\left(\mathbf{h}^{(d)}\right), \tag{D.3}
\end{align*}
$$

where the activation functions $\phi_{\ell}$ act elementwise. As always, we focus on networks with linear readout, i.e., $\phi_{d}(x)=x$, and assume Gaussian priors over the weights:

$$
\begin{equation*}
W_{i j}^{(\ell)} \underset{\text { i.i.d. }}{\sim} \mathcal{N}\left(0, \sigma_{\ell}^{2}\right) . \tag{D.4}
\end{equation*}
$$

We enforce the definition of the network architecture via Fourier representations of the Dirac distribution, with $\mathbf{q}_{\mu}^{(\ell)}$ being the Lagrange multiplier that enforces the definition of the preactivation $\mathbf{h}_{\mu}^{(\ell)}$. Then, after integrating out the weights using the fact that the relevant integrals are Gaussian, this allows us to write the partition function as

$$
\begin{equation*}
Z=\int \prod_{\mu=1}^{p} \prod_{\ell=1}^{d} \frac{d \mathbf{h}_{\mu}^{(\ell)} d \mathbf{q}_{\mu}^{(\ell)}}{(2 \pi)^{n_{\ell}}} \exp \left[S\left(\left\{\mathbf{h}_{\mu}^{(\ell)}\right\},\left\{\mathbf{q}_{\mu}^{(\ell)}\right\}\right)\right] \tag{D.5}
\end{equation*}
$$

where the "effective action" for the preactivations and Lagrange multipliers is

$$
\begin{align*}
S= & -\frac{1}{2} \beta \sum_{\mu=1}^{p}\left\|\mathbf{h}_{\mu}^{(d)}-\mathbf{y}_{\mu}\right\|^{2}+\sum_{\ell=1}^{d} \sum_{\mu=1}^{p} i \mathbf{q}_{\mu}^{(\ell)} \cdot \mathbf{h}_{\mu}^{(\ell)} \\
& -\frac{1}{2} \sum_{\ell=1}^{d} \frac{\sigma_{\ell}^{2}}{n_{\ell-1}} \sum_{\mu, \nu=1}^{p} \mathbf{q}_{\mu}^{(\ell)} \cdot \mathbf{q}_{\nu}^{(\ell)} \phi_{\ell-1}\left(\mathbf{h}_{\mu}^{(\ell-1)}\right) \cdot \phi_{\ell-1}\left(\mathbf{h}_{\nu}^{(\ell-1)}\right) . \tag{D.6}
\end{align*}
$$

As described in Appendix B, source terms can be added to the effective action to allow computation of various averages. For deep linear networks, it is convenient to scale the source terms by an overall factor of $-1 / 2$, for which we must correct when computing the averages:

$$
\begin{equation*}
S_{\mathrm{J}}=-\frac{1}{2} \sum_{\ell=1}^{d-1} \sum_{\mu, \nu=1}^{p} J_{\mu \nu}^{(\ell)} \phi_{\ell}\left(\mathbf{h}_{\mu}^{(\ell)}\right) \cdot \phi_{\ell}\left(\mathbf{h}_{\nu}^{(\ell)}\right) \tag{D.7}
\end{equation*}
$$

For an MLP, our task is therefore to integrate out the preactivations and corresponding Lagrange multipliers. We will do so sequentially from the first layer to the last, keeping terms up to the desired order at each step, akin to the approach of Yaida [8]. So long as $n_{d}$ and $d$ are fixed and small relative to the width of the hidden layers, this is a consistent perturbative approach, as noted by Yaida [8].

## D. 2 General form of the perturbative layer integrals for a deep linear network

In this section, we evaluate the general form of the integrals required to perturbatively marginalize out a given layer of a deep linear network to $\mathcal{O}\left(n^{-1}\right)$. These integrals are generically of the form

$$
\begin{align*}
I=\int \prod_{\mu=1}^{p} \frac{d \mathbf{h}_{\mu} d \mathbf{q}_{\mu}}{(2 \pi)^{n_{2}}} \exp \left(\sum_{\mu=1}^{p}\right. & i \mathbf{q}_{\mu} \cdot \mathbf{h}_{\mu}-\frac{1}{2} \sum_{\mu, \nu=1}^{p} G_{\mu \nu}\left(\mathbf{q}_{\mu} \cdot \mathbf{q}_{\nu}\right)+\sum_{\mu=1}^{p} \mathbf{j}_{\mu} \cdot \mathbf{h}_{\mu} \\
& -\frac{1}{2} \frac{1}{n_{2}} \sum_{\mu, \nu=1}^{p} A_{\mu \nu}\left(\mathbf{h}_{\mu} \cdot \mathbf{h}_{\nu}\right) \\
& +\frac{1}{4} \frac{g}{n_{1}} \sum_{\mu, \nu, \rho, \lambda=1}^{p} G_{\mu \nu}\left(\mathbf{q}_{\nu} \cdot \mathbf{q}_{\rho}\right) G_{\rho \lambda}\left(\mathbf{q}_{\lambda} \cdot \mathbf{q}_{\mu}\right) \\
& \left.+\frac{1}{2} \frac{1}{n_{1}} \sum_{\mu, \nu=1}^{p} B_{\mu \nu}\left(\mathbf{q}_{\mu} \cdot \mathbf{q}_{\nu}\right)\right) \tag{D.8}
\end{align*}
$$

where $\mathbf{h}_{\mu}, \mathbf{q}_{\mu} \in \mathbb{R}^{n_{2}}$. Here, $G$ is a positive semidefinite matrix, while $A$ and $B$ are symmetric matrices that need not be positive semidefinite. Furthermore, $\mathbf{j}_{\mu}$ is some source, while $g$ is a coupling constant. We will first evaluate this integral up to terms of $\mathcal{O}\left(n_{1}^{-1}\right)$ for $n_{1} \gg 1$, assuming that $G, A, B, \mathbf{j}_{\mu}$, and $g$ are $\mathcal{O}(1)$ functions of $n_{1}$, and then evaluate it up to terms of $\mathcal{O}\left(n_{1}^{-1}, n_{2}^{-1}\right)$ for $n_{1}, n_{2} \gg 1$, assuming that $G, A, B, \mathbf{j}_{\mu}$, and $g$ are also $\mathcal{O}(1)$ functions of $n_{2}$.
We will proceed by evaluating the integrals for $G$ invertible, and then infer the general case by a continuity argument. We treat the quartic term perturbatively, and all other terms directly. Writing

$$
\begin{equation*}
C \equiv G-\frac{1}{n_{1}} B \tag{D.9}
\end{equation*}
$$

the leading term in the integral over $\mathbf{q}_{\mu}$ is

$$
\begin{equation*}
\frac{1}{(2 \pi)^{n_{2} p / 2} \operatorname{det}(C)^{n_{2} / 2}} \exp \left(-\frac{1}{2} \sum_{\mu, \nu=1}^{p} C_{\mu \nu}^{-1}\left(\mathbf{h}_{\mu} \cdot \mathbf{h}_{\nu}\right)\right) \tag{D.10}
\end{equation*}
$$

Multiplying and dividing by this quantity, we can compute the perturbative correction from the quartic term using the fact that $\mathbf{q}_{\mu}$ then behaves as a Gaussian random vector of mean $\overline{\mathbf{q}}_{\mu}=i \sum_{\nu=1}^{p} C_{\mu \nu}^{-1} \mathbf{h}_{\nu}$ ${ }_{\tilde{\mathbf{q}}}^{\mu}$ and covariance $C_{\mu \nu}^{-1} I_{n_{2}}$. Denoting expectation with respect to this distribution as $\langle\langle\cdot\rangle\rangle_{q}$ and writing $\tilde{\mathbf{q}}_{\mu} \equiv \mathbf{q}_{\mu}-\overline{\mathbf{q}}_{\mu}$, Isserlis' theorem yields

$$
\begin{align*}
\left\langle\left\langle\left(\mathbf{q}_{\nu} \cdot \mathbf{q}_{\rho}\right)\left(\mathbf{q}_{\lambda} \cdot \mathbf{q}_{\mu}\right)\right\rangle\right\rangle_{q}= & \left\langle\left\langle\left(\left[\tilde{\mathbf{q}}_{\nu}+\overline{\mathbf{q}}_{\nu}\right] \cdot\left[\tilde{\mathbf{q}}_{\rho}+\overline{\mathbf{q}}_{\rho}\right]\right)\left(\left[\tilde{\mathbf{q}}_{\lambda}+\overline{\mathbf{q}}_{\lambda}\right] \cdot\left[\tilde{\mathbf{q}}_{\mu}+\overline{\mathbf{q}}_{\mu}\right]\right)\right\rangle\right\rangle_{q}  \tag{D.11}\\
= & \left\langle\left\langle\left(\tilde{\mathbf{q}}_{\nu} \cdot \tilde{\mathbf{q}}_{\rho}+\tilde{\mathbf{q}}_{\nu} \cdot \overline{\mathbf{q}}_{\rho}+\overline{\mathbf{q}}_{\nu} \cdot \tilde{\mathbf{q}}_{\rho}+\overline{\mathbf{q}}_{\nu} \cdot \overline{\mathbf{q}}_{\rho}\right)\right.\right. \\
& \left.\left.\times\left(\tilde{\mathbf{q}}_{\lambda} \cdot \tilde{\mathbf{q}}_{\mu}+\tilde{\mathbf{q}}_{\lambda} \cdot \overline{\mathbf{q}}_{\mu}+\overline{\mathbf{q}}_{\lambda} \cdot \tilde{\mathbf{q}}_{\mu}+\overline{\mathbf{q}}_{\lambda} \cdot \overline{\mathbf{q}}_{\mu}\right)\right\rangle\right\rangle_{q}  \tag{D.12}\\
= & \left\langle\left\langle\left(\tilde{\mathbf{q}}_{\nu} \cdot \tilde{\mathbf{q}}_{\rho}\right)\left(\tilde{\mathbf{q}}_{\lambda} \cdot \tilde{\mathbf{q}}_{\mu}\right)\right\rangle\right\rangle_{q}+\left\langle\left\langle\left(\tilde{\mathbf{q}}_{\nu} \cdot \tilde{\mathbf{q}}_{\rho}\right)\right\rangle\right\rangle_{q}\left(\overline{\mathbf{q}}_{\lambda} \cdot \overline{\mathbf{q}}_{\mu}\right) \\
& +\left\langle\left\langle\left(\tilde{\mathbf{q}}_{\nu} \cdot \overline{\mathbf{q}}_{\rho}\right)\left(\tilde{\mathbf{q}}_{\lambda} \cdot \overline{\mathbf{q}}_{\mu}\right)\right\rangle\right\rangle_{q}+\left\langle\left\langle\left(\tilde{\mathbf{q}}_{\nu} \cdot \overline{\mathbf{q}}_{\rho}\right)\left(\overline{\mathbf{q}}_{\lambda} \cdot \tilde{\mathbf{q}}_{\mu}\right)\right\rangle\right\rangle_{q} \\
& +\left\langle\left\langle\left(\overline{\mathbf{q}}_{\nu} \cdot \tilde{\mathbf{q}}_{\rho}\right)\left(\tilde{\mathbf{q}}_{\lambda} \cdot \overline{\mathbf{q}}_{\mu}\right)\right\rangle\right\rangle_{q}+\left\langle\left\langle\left(\overline{\mathbf{q}}_{\nu} \cdot \tilde{\mathbf{q}}_{\rho}\right)\left(\overline{\mathbf{q}}_{\lambda} \cdot \tilde{\mathbf{q}}_{\mu}\right)\right\rangle\right\rangle_{q} \\
& +\left(\overline{\mathbf{q}}_{\nu} \cdot \overline{\mathbf{q}}_{\rho}\right)\left\langle\left\langle\left(\tilde{\mathbf{q}}_{\lambda} \cdot \tilde{\mathbf{q}}_{\mu}\right)\right\rangle\right\rangle_{q}+\left(\overline{\mathbf{q}}_{\nu} \cdot \overline{\mathbf{q}}_{\rho}\right)\left(\overline{\mathbf{q}}_{\lambda} \cdot \overline{\mathbf{q}}_{\mu}\right)  \tag{D.13}\\
= & n_{2}^{2} C_{\nu \rho}^{-1} C_{\lambda \mu}^{-1}+n_{2} C_{\nu \lambda}^{-1} C_{\rho \mu}^{-1}+n_{2} C_{\nu \mu}^{-1} C_{\rho \lambda}^{-1}+n_{2} C_{\nu \rho}^{-1}\left(\overline{\mathbf{q}}_{\lambda} \cdot \overline{\mathbf{q}}_{\mu}\right) \\
& +C_{\nu \lambda}^{-1}\left(\overline{\mathbf{q}}_{\rho} \cdot \overline{\mathbf{q}}_{\mu}\right)+C_{\nu \mu}^{-1}\left(\overline{\mathbf{q}}_{\rho} \cdot \overline{\mathbf{q}}_{\lambda}\right) \\
& +C_{\rho \lambda}^{-1}\left(\overline{\mathbf{q}}_{\nu} \cdot \overline{\mathbf{q}}_{\mu}\right)+C_{\rho \mu}^{-1}\left(\overline{\mathbf{q}}_{\nu} \cdot \overline{\mathbf{q}}_{\lambda}\right) \\
& +n_{2}\left(\overline{\mathbf{q}}_{\nu} \cdot \overline{\mathbf{q}}_{\rho}\right) C_{\mu \lambda}^{-1}+\left(\overline{\mathbf{q}}_{\nu} \cdot \overline{\mathbf{q}}_{\rho}\right)\left(\overline{\mathbf{q}}_{\lambda} \cdot \overline{\mathbf{q}}_{\mu}\right) \tag{D.14}
\end{align*}
$$

Then, the quartic correction to the integral over $\mathbf{q}_{\mu}$ is proportional to

$$
\begin{align*}
\sum_{\mu, \nu, \rho, \lambda=1}^{p} G_{\mu \nu} G_{\rho \lambda}\left\langle\left\langle\left(\mathbf{q}_{\nu} \cdot \mathbf{q}_{\rho}\right)\left(\mathbf{q}_{\lambda} \cdot \mathbf{q}_{\mu}\right)\right\rangle\right\rangle_{q}= & n_{2}\left(n_{2}+1\right) \operatorname{tr}\left(G C^{-1} G C^{-1}\right)+n_{2} \operatorname{tr}\left(G C^{-1}\right)^{2} \\
& -2\left(n_{2}+1\right) \operatorname{tr}\left(G C^{-1} G C^{-1} H C^{-1}\right) \\
& -2 \operatorname{tr}\left(G C^{-1}\right) \operatorname{tr}\left(G C^{-1} H C^{-1}\right) \\
& +\operatorname{tr}\left(G C^{-1} H C^{-1} G C^{-1} H C^{-1}\right), \tag{D.15}
\end{align*}
$$

where we write $H_{\mu \nu} \equiv \mathbf{h}_{\mu} \cdot \mathbf{h}_{\nu}$.
We now must integrate over $\mathbf{h}_{\mu}$. The leading term is simply

$$
\begin{equation*}
\operatorname{det}(C D)^{-n_{2} / 2} \exp \left(\frac{1}{2} \sum_{\mu, \nu=1}^{p} D_{\mu \nu}^{-1} J_{\mu \nu}\right) \tag{D.16}
\end{equation*}
$$

where we have defined

$$
\begin{equation*}
D \equiv C^{-1}+\frac{1}{n_{2}} A \tag{D.17}
\end{equation*}
$$

and $J_{\mu \nu} \equiv \mathbf{j}_{\mu} \cdot \mathbf{j}_{\nu}$. Multiplying and dividing by this quantity, we can compute the perturbative correction from the quartic term using the fact that $\mathbf{h}_{\mu}$ then behaves as a Gaussian random vector of mean $\overline{\mathbf{h}}_{\mu}=\sum_{\nu=1}^{p} D_{\mu \nu}^{-1} \mathbf{j}_{\nu}$ and covariance $D_{\mu \nu}^{-1} I_{n_{2}}$. We denote expectations with respect to this distribution by $\langle\langle\cdot\rangle\rangle_{h}$, and define $\tilde{\mathbf{h}}_{\mu} \equiv \mathbf{h}_{\mu}-\overline{\mathbf{h}}_{\mu}$. Then, we have

$$
\begin{equation*}
\left\langle\left\langle H_{\mu \nu}\right\rangle\right\rangle_{h}=\left\langle\left\langle\mathbf{h}_{\mu} \cdot \mathbf{h}_{\nu}\right\rangle\right\rangle_{h}=\overline{\mathbf{h}}_{\mu} \cdot \overline{\mathbf{h}}_{\nu}+n_{2} D_{\mu \nu}^{-1} \tag{D.18}
\end{equation*}
$$

and, by analogy to the corresponding four-point average for $\mathbf{q}_{\mu}$,

$$
\begin{align*}
\left\langle\left\langle\left(\mathbf{h}_{\nu} \cdot \mathbf{h}_{\rho}\right)\left(\mathbf{h}_{\lambda} \cdot \mathbf{h}_{\mu}\right)\right\rangle\right\rangle_{h}= & n_{2}^{2} D_{\nu \rho}^{-1} D_{\lambda \mu}^{-1}+n_{2} D_{\nu \lambda}^{-1} D_{\rho \mu}^{-1}+n_{2} D_{\nu \mu}^{-1} D_{\rho \lambda}^{-1}+n_{2} D_{\nu \rho}^{-1}\left(\overline{\mathbf{h}}_{\lambda} \cdot \overline{\mathbf{h}}_{\mu}\right) \\
& +D_{\nu \lambda}^{-1}\left(\overline{\mathbf{h}}_{\rho} \cdot \overline{\mathbf{h}}_{\mu}\right)+D_{\nu \mu}^{-1}\left(\overline{\mathbf{h}}_{\rho} \cdot \overline{\mathbf{h}}_{\lambda}\right) \\
& +D_{\rho \lambda}^{-1}\left(\overline{\mathbf{h}}_{\nu} \cdot \overline{\mathbf{h}}_{\mu}\right)+D_{\rho \mu}^{-1}\left(\overline{\mathbf{h}}_{\nu} \cdot \overline{\mathbf{h}}_{\lambda}\right) \\
& +n_{2}\left(\overline{\mathbf{h}}_{\nu} \cdot \overline{\mathbf{h}}_{\rho}\right) D_{\mu \lambda}^{-1}+\left(\overline{\mathbf{h}}_{\nu} \cdot \overline{\mathbf{h}}_{\rho}\right)\left(\overline{\mathbf{h}}_{\lambda} \cdot \overline{\mathbf{h}}_{\mu}\right) . \tag{D.19}
\end{align*}
$$

Then, the correction to the integral over $\mathbf{h}_{\mu}$ is proportional to

$$
\begin{align*}
\sum_{\mu, \nu, \rho, \lambda=1}^{p} G_{\mu \nu} G_{\rho \lambda}\left\langle\left\langle\left(\mathbf{q}_{\nu} \cdot \mathbf{q}_{\rho}\right)\left(\mathbf{q}_{\lambda} \cdot \mathbf{q}_{\mu}\right)\right\rangle\right\rangle= & n_{2}\left(n_{2}+1\right) \operatorname{tr}\left(G C^{-1} G C^{-1}\right)+n_{2} \operatorname{tr}\left(G C^{-1}\right)^{2} \\
& -2\left(n_{2}+1\right) \operatorname{tr}\left(G C^{-1} G C^{-1} D^{-1} J D^{-1} C^{-1}\right) \\
& -2 n_{2}\left(n_{2}+1\right) \operatorname{tr}\left(G C^{-1} G C^{-1} D^{-1} C^{-1}\right) \\
& -2 \operatorname{tr}\left(G C^{-1}\right) \operatorname{tr}\left(G C^{-1} D^{-1} J D^{-1} C^{-1}\right) \\
& -2 n_{2} \operatorname{tr}\left(G C^{-1}\right) \operatorname{tr}\left(G C^{-1} D^{-1} C^{-1}\right) \\
& +n_{2}\left(n_{2}+1\right) \operatorname{tr}\left(C^{-1} G C^{-1} D^{-1} C^{-1} G C^{-1} D^{-1}\right) \\
& +n_{2} \operatorname{tr}\left(C^{-1} G C^{-1} D^{-1}\right)^{2} \\
& +2\left(n_{2}+1\right) \operatorname{tr}\left(C^{-1} G C^{-1} D^{-1} C^{-1} G C^{-1} D^{-1} J D^{-1}\right) \\
& +2 \operatorname{tr}\left(C^{-1} G C^{-1} D^{-1}\right) \operatorname{tr}\left(C^{-1} G C^{-1} D^{-1} J D^{-1}\right) \\
& +\operatorname{tr}\left(C^{-1} G C^{-1} D^{-1} J D^{-1} C^{-1} G C^{-1} D^{-1} J D^{-1}\right), \tag{D.20}
\end{align*}
$$

where we have noted that

$$
\begin{align*}
&\left\langle\left\langle\operatorname{tr}\left(G C^{-1} H C^{-1} G C^{-1} H C^{-1}\right)\right\rangle\right\rangle_{h} \\
&= \sum_{\mu, \nu, \rho, \lambda=1}^{p}\left(C^{-1} G C^{-1}\right)_{\mu \nu}\left(C^{-1} G C^{-1}\right)_{\rho \lambda}\left\langle\left\langle\left(\mathbf{h}_{\nu} \cdot \mathbf{h}_{\rho}\right)\left(\mathbf{h}_{\lambda} \cdot \mathbf{h}_{\mu}\right)\right\rangle\right\rangle_{h}  \tag{D.21}\\
&= n_{2}\left(n_{2}+1\right) \operatorname{tr}\left(C^{-1} G C^{-1} D^{-1} C^{-1} G C^{-1} D^{-1}\right)+n_{2} \operatorname{tr}\left(C^{-1} G C^{-1} D^{-1}\right)^{2} \\
&+2\left(n_{2}+1\right) \operatorname{tr}\left(C^{-1} G C^{-1} D^{-1} C^{-1} G C^{-1} D^{-1} J D^{-1}\right) \\
&+2 \operatorname{tr}\left(C^{-1} G C^{-1} D^{-1}\right) \operatorname{tr}\left(C^{-1} G C^{-1} D^{-1} J D^{-1}\right) \\
&+\operatorname{tr}\left(C^{-1} G C^{-1} D^{-1} J D^{-1} C^{-1} G C^{-1} D^{-1} J D^{-1}\right) \tag{D.22}
\end{align*}
$$

by analogy with the corresponding quartic expectation for $\mathbf{q}_{\mu}$.
We must now expand our results in $n_{1}^{-1}$. The inverses of the matrices $C$ and $D$ have Neumann series

$$
\begin{equation*}
C^{-1}=G^{-1}+\frac{1}{n_{1}} G^{-1} B G^{-1}+\mathcal{O}\left(n_{1}^{-2}\right) \tag{D.23}
\end{equation*}
$$

and

$$
\begin{align*}
D^{-1} & =\left(C^{-1}+\frac{1}{n_{2}} A\right)^{-1}  \tag{D.24}\\
& =\left(G^{-1}+\frac{1}{n_{1}} G^{-1} B G^{-1}+\frac{1}{n_{2}} A+\mathcal{O}\left(n_{1}^{-2}\right)\right)^{-1}  \tag{D.25}\\
& =F^{-1} G-\frac{1}{n_{1}} F^{-1} B F^{-\top}+\mathcal{O}\left(n_{1}^{-2}\right) \tag{D.26}
\end{align*}
$$

where we have defined

$$
\begin{equation*}
F \equiv I_{p}+\frac{1}{n_{2}} G A \tag{D.27}
\end{equation*}
$$

and we write $F^{-\top}=\left(F^{-1}\right)^{\top}=\left(F^{\top}\right)^{-1}$. Then, using the series expansion of the log-determinant, we find that the logarithm of the leading term expands as

$$
\begin{align*}
\frac{1}{2} \operatorname{tr}\left(D^{-1} J\right)-\frac{1}{2} n_{2} \log \operatorname{det}(C D)= & \frac{1}{2} \operatorname{tr}\left(F^{-1} G J\right)-\frac{1}{2} n_{2} \log \operatorname{det}(F) \\
& -\frac{1}{2} \frac{1}{n_{1}} \operatorname{tr}\left(F^{-1} B F^{-\top} J\right)+\frac{1}{2} \frac{1}{n_{1}} \operatorname{tr}\left(F^{-1} B A\right) \\
& +\mathcal{O}\left(n_{1}^{-2}\right) \tag{D.28}
\end{align*}
$$

while the quartic correction simplifies to

$$
\begin{align*}
\frac{1}{4} \frac{g}{n_{1}} & \sum_{\mu, \nu, \rho, \lambda=1}^{p} G_{\mu \nu} G_{\rho \lambda}\left\langle\left\langle\left(\mathbf{q}_{\nu} \cdot \mathbf{q}_{\rho}\right)\left(\mathbf{q}_{\lambda} \cdot \mathbf{q}_{\mu}\right)\right\rangle\right\rangle \\
= & \frac{1}{4} \frac{g}{n_{1}} n_{2}\left(n_{2}+p+1\right) p \\
& +\frac{1}{4} \frac{n_{2} g}{n_{1}}\left(\left(n_{2}+1\right) \operatorname{tr}\left(F^{-2}\right)+\operatorname{tr}\left(F^{-1}\right)^{2}-2\left(n_{2}+p+1\right) \operatorname{tr}\left(F^{-1}\right)\right) \\
& +\frac{1}{2} \frac{g}{n_{1}}\left(\left(n_{2}+1\right) \operatorname{tr}\left(F^{-3} G J\right)+\operatorname{tr}\left(F^{-1}\right) \operatorname{tr}\left(F^{-2} G J\right)-\left(n_{2}+p+1\right) \operatorname{tr}\left(F^{-2} G J\right)\right) \\
& +\frac{1}{4} \frac{g}{n_{1}} \operatorname{tr}\left(F^{-2} G J F^{-2} G J\right) \\
& +\mathcal{O}\left(n_{1}^{-2}\right) . \tag{D.29}
\end{align*}
$$

Combining these results, we find that the result of integrating out the layer to $\mathcal{O}\left(n_{1}^{-1}\right)$ is

$$
\begin{align*}
\log I= & \frac{1}{2} \operatorname{tr}\left(F^{-1} G J\right)-\frac{1}{2} n_{2} \log \operatorname{det}(F) \\
& -\frac{1}{2} \frac{1}{n_{1}} \operatorname{tr}\left(F^{-1} B F^{-\top} J\right)+\frac{1}{2} \frac{1}{n_{1}} \operatorname{tr}\left(F^{-1} B A\right) \\
& +\frac{1}{4} \frac{g}{n_{1}} n_{2}\left(n_{2}+p+1\right) p \\
& +\frac{1}{4} \frac{n_{2} g}{n_{1}}\left(\left(n_{2}+1\right) \operatorname{tr}\left(F^{-2}\right)+\operatorname{tr}\left(F^{-1}\right)^{2}-2\left(n_{2}+p+1\right) \operatorname{tr}\left(F^{-1}\right)\right) \\
& +\frac{1}{2} \frac{g}{n_{1}}\left(\left(n_{2}+1\right) \operatorname{tr}\left(F^{-3} G J\right)+\operatorname{tr}\left(F^{-1}\right) \operatorname{tr}\left(F^{-2} G J\right)-\left(n_{2}+p+1\right) \operatorname{tr}\left(F^{-2} G J\right)\right) \\
& +\frac{1}{4} \frac{g}{n_{1}} \operatorname{tr}\left(F^{-2} G J F^{-2} G J\right) \\
& +\mathcal{O}\left(n_{1}^{-2}\right) . \tag{D.30}
\end{align*}
$$

As this result is a continuous function of $G$, as the set of full-rank positive definite matrices is dense in the space of positive semidefinite matrices, this result holds for all positive-semidefinite $G$.
We now further expand this result in $n_{2}^{-1}$. This yields

$$
\begin{equation*}
F^{-1}=I_{p}-\frac{1}{n_{2}} G A+\frac{1}{n_{2}^{2}} G A G A+\mathcal{O}\left(n_{2}^{-3}\right) \tag{D.31}
\end{equation*}
$$

and

$$
\begin{equation*}
\log \operatorname{det}(F)=\frac{1}{n_{2}} \operatorname{tr}(G A)-\frac{1}{2} \frac{1}{n_{2}^{2}} \operatorname{tr}(G A G A)+\mathcal{O}\left(n_{2}^{-3}\right) \tag{D.32}
\end{equation*}
$$

hence we find that the logarithm of the leading term yields

$$
\begin{align*}
\frac{1}{2} \operatorname{tr}\left(D^{-1} J\right)-\frac{1}{2} n_{2} \log \operatorname{det}(C D)= & \frac{1}{2} \operatorname{tr}(G J)-\frac{1}{2} \operatorname{tr}(G A)+\frac{1}{4} \frac{1}{n_{2}} \operatorname{tr}(G A G A) \\
& -\frac{1}{2} \frac{1}{n_{2}} \operatorname{tr}(G A G J)+\frac{1}{2} \frac{1}{n_{1}} \operatorname{tr}(B(A-J)) \\
& +\mathcal{O}\left(n_{1}^{-2}, n_{2}^{-2}, n_{1}^{-1} n_{2}^{-1}\right) \tag{D.33}
\end{align*}
$$

After some straightforward but tedious algebra, the quartic term reduces to

$$
\begin{align*}
\frac{1}{4} \frac{g}{n_{1}} \sum_{\mu, \nu, \rho, \lambda=1}^{p} G_{\mu \nu} G_{\rho \lambda}\left\langle\left\langle\left(\mathbf{q}_{\nu} \cdot \mathbf{q}_{\rho}\right)\left(\mathbf{q}_{\lambda} \cdot \mathbf{q}_{\mu}\right)\right\rangle\right\rangle= & \frac{1}{4} \frac{g}{n_{1}} \operatorname{tr}(G(A-J) G(A-J)) \\
& +\mathcal{O}\left(n_{1}^{-2}, n_{2}^{-2}, n_{1}^{-1} n_{2}^{-1}\right) \tag{D.34}
\end{align*}
$$

Combining these results, we find that the result of integrating out the layer is

$$
\begin{align*}
\log I= & \frac{1}{2} \operatorname{tr}(G J)-\frac{1}{2} \operatorname{tr}(G A)+\frac{1}{4} \frac{1}{n_{2}}\left(1+\frac{n_{2}}{n_{1}} g\right) \operatorname{tr}(G A G A) \\
& -\frac{1}{2}\left(1+\frac{n_{2}}{n_{1}} g\right) \operatorname{tr}(G A G J)+\frac{1}{2} \frac{1}{n_{1}} \operatorname{tr}(B(A-J))+\frac{1}{4} \frac{1}{n_{1}} g \operatorname{tr}(G J G J) \\
& +\mathcal{O}\left(n_{1}^{-2}, n_{2}^{-2}, n_{1}^{-1} n_{2}^{-1}\right) . \tag{D.35}
\end{align*}
$$

Again, this result is continuous in $G$, hence it holds even if $G$ is rank-deficient.

## D. 3 Perturbative computation of the partition function of a deep linear network

We now apply the results of Appendix D. 2 to compute the partition function for a deep linear network to the desired order. Our starting point is the effective action before any of the layers have been integrated out, including a source term:

$$
\begin{align*}
S= & -\frac{1}{2} \beta \sum_{\mu=1}^{p}\left\|\mathbf{h}_{\mu}^{(d)}-\mathbf{y}_{\mu}\right\|^{2}+\sum_{\ell=1}^{d} \sum_{\mu=1}^{p} i \mathbf{q}_{\mu}^{(\ell)} \cdot \mathbf{h}_{\mu}^{(\ell)}-\frac{1}{2} \sum_{\mu, \nu=1}^{p}\left(\sigma_{1}^{2} G_{x x}\right)_{\mu \nu}\left(\mathbf{q}_{\mu}^{(1)} \cdot \mathbf{q}_{\nu}^{(1)}\right) \\
& -\frac{1}{2} \sum_{\ell=1}^{d-1} \frac{1}{n_{\ell}} \sum_{\mu, \nu=1}^{p}\left(J_{\mu \nu}^{(\ell)}+\sigma_{\ell+1}^{2} \mathbf{q}_{\mu}^{(\ell+1)} \cdot \mathbf{q}_{\nu}^{(\ell+1)}\right)\left(\mathbf{h}_{\mu}^{(\ell)} \cdot \mathbf{h}_{\nu}^{(\ell)}\right) . \tag{D.36}
\end{align*}
$$

Applying the results of Appendix D. 2 with

$$
\begin{align*}
G & =\sigma_{1}^{2} G_{x x}, \\
\mathbf{j}_{\mu} & =0 \\
A & =J^{(1)}+\sigma_{2}^{2} Q^{(2)},  \tag{D.37}\\
B & =0, \quad \text { and } \\
g & =0,
\end{align*}
$$

we find that the effective action after integrating out the first layer is

$$
\begin{align*}
S^{(1)}= & -\frac{1}{2} \beta \sum_{\mu=1}^{p}\left\|\mathbf{h}_{\mu}^{(d)}-\mathbf{y}_{\mu}\right\|^{2}+\sum_{\ell=2}^{d} \sum_{\mu=1}^{p} i \mathbf{q}_{\mu}^{(\ell)} \cdot \mathbf{h}_{\mu}^{(\ell)}-\frac{1}{2} \sum_{\mu, \nu=1}^{p}\left(m_{2}^{2} G_{x x}\right)_{\mu \nu}\left(\mathbf{q}_{\mu}^{(2)} \cdot \mathbf{q}_{\nu}^{(2)}\right) \\
& -\frac{1}{2} \sum_{\ell=2}^{d-1} \frac{1}{n_{\ell}} \sum_{\mu, \nu=1}^{p}\left(J_{\mu \nu}^{(\ell)}+\sigma_{\ell+1}^{2} \mathbf{q}_{\mu}^{(\ell+1)} \cdot \mathbf{q}_{\nu}^{(\ell+1)}\right)\left(\mathbf{h}_{\mu}^{(\ell)} \cdot \mathbf{h}_{\nu}^{(\ell)}\right) \\
& +\frac{1}{4} \frac{g_{1}}{n_{1}} m_{2}^{4} \operatorname{tr}\left(G_{x x} Q^{(2)} G_{x x} Q^{(2)}\right)+\frac{1}{2} \frac{g_{1}}{n_{1}} m_{2}^{2} \operatorname{tr}\left(G_{x x} \tilde{J}_{1} G_{x x} Q^{(2)}\right) \\
& -\frac{1}{2} \operatorname{tr}\left(m_{1}^{2} G_{x x} J^{(1)}\right)+\frac{1}{4} \frac{g_{1}}{n_{1}} m_{1}^{4} \operatorname{tr}\left(G_{x x} J^{(1)} G_{x x} J^{(1)}\right) \\
& +\mathcal{O}\left(n^{-2}\right) \tag{D.38}
\end{align*}
$$

where we have defined

$$
\begin{align*}
m_{1} & \equiv \sigma_{1}, \\
m_{2} & \equiv \sigma_{2} m_{1}, \\
g_{1} & \equiv 1, \quad \text { and }  \tag{D.39}\\
\tilde{J}_{1} & \equiv m_{1}^{2} J^{(1)} .
\end{align*}
$$

Assuming that the network has more than one hidden layer, if we now again apply the results of Appendix D. 2 with

$$
\begin{align*}
G & =m_{2}^{2} G_{x x} \\
\mathbf{j}_{\mu} & =0 \\
A & =J^{(2)}+\sigma_{3}^{2} Q^{(3)}  \tag{D.40}\\
B & =g_{1} m_{2}^{2} G_{x x} \tilde{J}_{1} G_{x x}, \quad \text { and } \\
g & =g_{1}
\end{align*}
$$

we find that the effective action after integrating out the first two layers is

$$
\begin{align*}
S^{(2)}= & -\frac{1}{2} \beta \sum_{\mu=1}^{p}\left\|\mathbf{h}_{\mu}^{(d)}-\mathbf{y}_{\mu}\right\|^{2}+\sum_{\ell=3}^{d} \sum_{\mu=1}^{p} i \mathbf{q}_{\mu}^{(\ell)} \cdot \mathbf{h}_{\mu}^{(\ell)}-\frac{1}{2} \sum_{\mu, \nu=1}^{p}\left(m_{3}^{2} G_{x x}\right)_{\mu \nu}\left(\mathbf{q}_{\mu}^{(3)} \cdot \mathbf{q}_{\nu}^{(3)}\right) \\
& -\frac{1}{2} \sum_{\ell=3}^{d-1} \frac{1}{n_{\ell}} \sum_{\mu, \nu=1}^{p}\left(J_{\mu \nu}^{(\ell)}+\sigma_{\ell+1}^{2} \mathbf{q}_{\mu}^{(\ell+1)} \cdot \mathbf{q}_{\nu}^{(\ell+1)}\right)\left(\mathbf{h}_{\mu}^{(\ell)} \cdot \mathbf{h}_{\nu}^{(\ell)}\right) \\
& +\frac{1}{4} \frac{g_{2}}{n_{2}} m_{3}^{4} \operatorname{tr}\left(G_{x x} Q^{(3)} G_{x x} Q^{(3)}\right) \\
& +\frac{1}{2} \frac{g_{2}}{n_{2}} m_{3}^{2} \operatorname{tr}\left(G_{x x} \tilde{J}_{2} G_{x x} Q^{(3)}\right) \\
& -\frac{1}{2} \operatorname{tr}\left(m_{1}^{2} G_{x x} J^{(1)}\right)-\frac{1}{2} \operatorname{tr}\left(m_{2}^{2} G_{x x} J^{(2)}\right) \\
& +\frac{1}{4} \frac{g_{1}}{n_{1}} m_{1}^{4} \operatorname{tr}\left(G_{x x} J^{(1)} G_{x x} J^{(1)}\right)+\frac{1}{4} \frac{g_{2}}{n_{2}} m_{2}^{4} \operatorname{tr}\left(G_{x x} J^{(2)} G_{x x} J^{(2)}\right) \\
& +\frac{1}{2} \frac{g_{1}}{n_{1}} m_{2}^{2} \operatorname{tr}\left(G_{x x} \tilde{J}_{1} G_{x x} J^{(2)}\right) \\
& +\mathcal{O}\left(n^{-2}\right) \tag{D.41}
\end{align*}
$$

where we have defined

$$
\begin{align*}
m_{3} & \equiv \sigma_{3} m_{2} \\
g_{2} & \equiv 1+\frac{n_{2}}{n_{1}} g_{1}, \quad \text { and }  \tag{D.42}\\
\tilde{J}_{2} & \equiv m_{2}^{2} J^{(2)}+\frac{n_{2}}{n_{1}} \frac{g_{1}}{g_{2}} \tilde{J}_{1}
\end{align*}
$$

Then, by induction, we can see that we can iterate this procedure to integrate out all of the hidden layers, yielding

$$
\begin{align*}
S^{(d-1)}= & -\frac{1}{2} \beta \sum_{\mu=1}^{p}\left\|\mathbf{h}_{\mu}^{(d)}-\mathbf{y}_{\mu}\right\|^{2}+\sum_{\mu=1}^{p} i \mathbf{q}_{\mu}^{(d)} \cdot \mathbf{h}_{\mu}^{(d)}-\frac{1}{2} \sum_{\mu, \nu=1}^{p}\left(m_{d}^{2} G_{x x}\right)_{\mu \nu}\left(\mathbf{q}_{\mu}^{(d)} \cdot \mathbf{q}_{\nu}^{(d)}\right) \\
& +\frac{1}{4} \frac{g_{d-1}}{n_{d-1}} m_{d}^{4} \operatorname{tr}\left(G_{x x} Q^{(d)} G_{x x} Q^{(d)}\right) \\
& +\frac{1}{2} \frac{g_{d-1}}{n_{d-1}} m_{d}^{2} \operatorname{tr}\left(G_{x x} \tilde{J}_{d-1} G_{x x} Q^{(d)}\right) \\
& -\frac{1}{2} \sum_{\ell=1}^{d-1} \operatorname{tr}\left(m_{\ell}^{2} G_{x x} J^{(\ell)}\right) \\
& +\frac{1}{4} \sum_{\ell=1}^{d-1} \frac{g_{\ell}}{n_{\ell}} m_{\ell}^{4} \operatorname{tr}\left(G_{x x} J^{(\ell)} G_{x x} J^{(\ell)}\right) \\
& +\frac{1}{2} \sum_{\ell=1}^{d-2} \frac{g_{\ell}}{n_{\ell}} m_{\ell+1}^{2} \operatorname{tr}\left(G_{x x} \tilde{J}_{\ell} G_{x x} J^{(\ell+1)}\right) \\
& +\mathcal{O}\left(n^{-2}\right) \tag{D.43}
\end{align*}
$$

where $m_{d}, g_{d-1}$, and $\tilde{J}_{d-1}$ are defined by the closed recurrences

$$
\begin{align*}
m_{\ell} & \equiv \sigma_{\ell} m_{\ell-1}  \tag{D.44}\\
g_{\ell} & \equiv 1+\frac{n_{\ell}}{n_{\ell-1}} g_{\ell-1}, \quad \text { and }  \tag{D.45}\\
\tilde{J}_{\ell} & \equiv m_{\ell}^{2} J^{(\ell)}+\frac{n_{\ell}}{n_{\ell-1}} \frac{g_{\ell-1}}{g_{\ell}} \tilde{J}_{\ell-1} \tag{D.46}
\end{align*}
$$

Applying the results of Appendix D. 2 one final time with

$$
\begin{align*}
G & =m_{d}^{2} G_{x x} \\
\mathbf{j}_{\mu} & =\beta \mathbf{y} \\
A & =\beta n_{d} I_{p}  \tag{D.47}\\
B & =g_{d-1} m_{d}^{2} G_{x x} \tilde{J}_{d-1} G_{x x}, \quad \text { and } \\
g & =g_{d-1}
\end{align*}
$$

we conclude that

$$
\begin{align*}
\log Z= & -\frac{1}{2} \beta n_{d} \operatorname{tr}\left(\tilde{\Gamma}^{-1} G_{y y}\right)-\frac{1}{2} n_{d} \log \operatorname{det}(\tilde{\Gamma}) \\
& +\frac{1}{4} \frac{n_{d} g_{d-1}}{n_{d-1}}\left(\left(n_{d}+p+1\right) p+\left(n_{d}+1\right) \operatorname{tr}\left(\tilde{\Gamma}^{-2}\right)+\operatorname{tr}\left(\tilde{\Gamma}^{-1}\right)^{2}-2\left(n_{d}+p+1\right) \operatorname{tr}\left(\tilde{\Gamma}^{-1}\right)\right) \\
& +\frac{1}{2} \frac{g_{d-1}}{n_{d-1}} \beta^{2} n_{d} m_{d}^{2}\left(\left(n_{d}+1\right) \operatorname{tr}\left(\tilde{\Gamma}^{-3} G_{x x} G_{y y}\right)+\operatorname{tr}\left(\tilde{\Gamma}^{-1}\right) \operatorname{tr}\left(\tilde{\Gamma}^{-2} G_{x x} G_{y y}\right)\right. \\
& \left.-\left(n_{d}+p+1\right) \operatorname{tr}\left(\tilde{\Gamma}^{-2} G_{x x} G_{y y}\right)\right) \\
& +\frac{1}{4} \frac{g_{d-1}}{n_{d-1}} \beta^{4} n_{d}^{2} m_{d}^{4} \operatorname{tr}\left(\tilde{\Gamma}^{-2} G_{x x} G_{y y} \tilde{\Gamma}^{-2} G_{x x} G_{y y}\right) \\
& -\frac{1}{2} \frac{g_{d-1}}{n_{d-1}} n_{d} m_{d}^{2} \operatorname{tr}\left[\left(\beta^{2} G_{x x} \tilde{\Gamma}^{-1} G_{y y} \tilde{\Gamma}^{-1} G_{x x}-\beta G_{x x} \tilde{\Gamma}^{-1} G_{x x}\right) \tilde{J}_{d-1}\right] \\
& -\frac{1}{2} \sum_{\ell=1}^{d-1} \operatorname{tr}\left(m_{\ell}^{2} G_{x x} J^{(\ell)}\right) \\
& +\frac{1}{4} \sum_{\ell=1}^{d-1} \frac{g_{\ell}}{n_{\ell}} m_{\ell}^{4} \operatorname{tr}\left(G_{x x} J^{(\ell)} G_{x x} J^{(\ell)}\right) \\
& +\frac{1}{2} \sum_{\ell=1}^{d-2} \frac{g_{\ell}}{n_{\ell}} m_{\ell+1}^{2} \operatorname{tr}\left(G_{x x} \tilde{J}_{\ell} G_{x x} J^{(\ell+1)}\right) \\
& +\mathcal{O}\left(n^{-2}\right), \tag{D.48}
\end{align*}
$$

where we have defined the matrix

$$
\begin{equation*}
\tilde{\Gamma} \equiv I_{p}+\beta m_{d}^{2} G_{x x} \tag{D.49}
\end{equation*}
$$

and absorbed the normalizing constant using the fact that $I_{p}-\beta m_{d}^{2} \tilde{\Gamma}^{-1} G_{x x}=\tilde{\Gamma}^{-1}$. As was the case for the individual layer integrals, a continuity argument implies that this expression can be applied even if $G_{x x}$ is rank-deficient.

## D. 4 Computing the average hidden layer kernels of a deep linear network

With the relevant partition function in hand, we can finally compute the average hidden layer kernels. In particular, we can immediately read off that

$$
\begin{align*}
\left\langle K^{(\ell)}\right\rangle= & m_{\ell}^{2} G_{x x} \\
& +\frac{g_{d-1}}{n_{d-1}} n_{d} m_{d}^{2} \operatorname{tr}\left[\left.\left(\beta^{2} G_{x x} \tilde{\Gamma}^{-1} G_{y y} \tilde{\Gamma}^{-1} G_{x x}-\beta G_{x x} \tilde{\Gamma}^{-1} G_{x x}\right) \frac{\delta \tilde{J}_{d-1}}{\delta J^{(\ell)}}\right|_{J^{(\ell)}=0}\right] \\
& +\mathcal{O}\left(n^{-2}\right) \tag{D.50}
\end{align*}
$$

hence our only task is to determine how the effective source $\tilde{J}_{d-1}$ depends on the source for a given layer. Fortunately, the recurrence relation for the effective source is extremely easy to solve, yielding

$$
\begin{equation*}
\tilde{J}_{d-1}=\sum_{\ell=1}^{d-1} m_{\ell}^{2} \frac{n_{d-1}}{n_{\ell}} \frac{g_{\ell}}{g_{d-1}} J^{(\ell)} \tag{D.51}
\end{equation*}
$$

Thus, defining the matrix

$$
\begin{equation*}
\Gamma \equiv \frac{1}{\beta m_{d}^{2}} \tilde{\Gamma}=G_{x x}+\frac{1}{\beta m_{d}^{2}} I_{p} \tag{D.52}
\end{equation*}
$$

we find that

$$
\begin{equation*}
\left\langle K^{(\ell)}\right\rangle=m_{\ell}^{2} G_{x x}+\frac{g_{\ell}}{n_{\ell}} n_{d} m_{\ell}^{2}\left(m_{d}^{-2} G_{x x} \Gamma^{-1} G_{y y} \Gamma^{-1} G_{x x}-G_{x x} \Gamma^{-1} G_{x x}\right)+\mathcal{O}\left(n^{-2}\right) \tag{D.53}
\end{equation*}
$$

To obtain the expression listed in the main text, we note that

$$
\begin{equation*}
\frac{g_{\ell}}{n_{\ell}}=\frac{1}{n_{\ell}}+\frac{g_{\ell-1}}{n_{\ell-1}} \tag{D.54}
\end{equation*}
$$

hence we have

$$
\begin{equation*}
\frac{g_{\ell}}{n_{\ell}}=\sum_{\ell^{\prime}=1}^{\ell} \frac{1}{n_{\ell^{\prime}}} \tag{D.55}
\end{equation*}
$$

mirroring the width dependence found by Yaida [8] in his study of the prior of deep linear networks.

## E Average kernels in a deep feedforward linear network with skip connections

In this appendix, we show that Conjecture 1 holds perturbatively for a linear feedforward network with arbitrary skip connections, following the method of Appendix D. Concretely, we consider a network defined as

$$
\begin{align*}
\mathbf{h}^{(0)} & =\mathbf{x}  \tag{E.1}\\
\mathbf{h}^{(\ell)} & =\sum_{\ell^{\prime}=0}^{\ell-1} \frac{\sigma_{\ell, \ell^{\prime}}}{\sqrt{n_{\ell^{\prime}}}} W^{\left(\ell, \ell^{\prime}\right)} \mathbf{h}^{\left(\ell^{\prime}\right)} \quad \ell=1, \ldots, d  \tag{E.2}\\
\mathbf{f} & =\mathbf{h}^{(d)} \tag{E.3}
\end{align*}
$$

where $\sigma_{\ell, \ell^{\prime}}$ is positive if layer $\ell$ receives input from an earlier layer $\ell^{\prime}<\ell$, and zero otherwise.

## E. 1 Perturbative computation of the partition function

Upon integrating out the weights, we obtain an effective action for the preactivations and the corresponding Lagrange multipliers of

$$
\begin{align*}
S= & -\beta \sum_{\mu=1}^{p} \varepsilon\left(\mathbf{h}_{\mu}^{(d)}, \mathbf{y}_{\mu}\right)+\sum_{\mu=1}^{p} \sum_{\ell=1}^{d} i \mathbf{q}_{\mu}^{(\ell)} \cdot \mathbf{h}_{\mu}^{(\ell)} \\
& -\frac{1}{2} \sum_{\ell=1}^{d-1} \frac{1}{n_{\ell}} \sum_{\mu, \nu=1}^{p}\left[J^{(\ell)}+\sum_{\ell^{\prime}=\ell+1}^{d} \sigma_{\ell^{\prime}, \ell}^{2}\left(\mathbf{q}_{\mu}^{\left(\ell^{\prime}\right)} \cdot \mathbf{q}_{\nu}^{\left(\ell^{\prime}\right)}\right)\right]\left(\mathbf{h}_{\mu}^{(\ell)} \cdot \mathbf{h}_{\nu}^{(\ell)}\right) \\
& -\frac{1}{2} \sum_{\ell=1}^{d} \sigma_{\ell, 0}^{2} \sum_{\mu, \nu=1}^{p}\left(G_{x x}\right)_{\mu \nu}\left(\mathbf{q}_{\mu}^{(\ell)} \cdot \mathbf{q}_{\nu}^{(\ell)}\right) . \tag{E.4}
\end{align*}
$$

Applying the result of Appendix D. 2 with

$$
\begin{align*}
G & =\sigma_{1,0}^{2} G_{x x} \\
\mathbf{j}_{\mu} & =\mathbf{0} \\
A & =J^{(1)}+\sum_{\ell^{\prime}=2}^{d} \sigma_{\ell^{\prime}, 1}^{2} Q^{\left(\ell^{\prime}\right)},  \tag{E.5}\\
B & =0, \quad \text { and } \\
g & =0
\end{align*}
$$

we find that the effective action after integrating out the first layer is

$$
\begin{align*}
S^{(1)}= & -\beta \sum_{\mu=1}^{p} \varepsilon\left(\mathbf{h}_{\mu}^{(d)}, \mathbf{y}_{\mu}\right)+\sum_{\mu=1}^{p} \sum_{\ell=2}^{d} i \mathbf{q}_{\mu}^{(\ell)} \cdot \mathbf{h}_{\mu}^{(\ell)} \\
& -\frac{1}{2} \sum_{\ell=2}^{d-1} \frac{1}{n_{\ell}} \sum_{\mu, \nu=1}^{p}\left[J^{(\ell)}+\sum_{\ell^{\prime}=\ell+1}^{d} \sigma_{\ell^{\prime}, \ell}^{2}\left(\mathbf{q}_{\mu}^{\left(\ell^{\prime}\right)} \cdot \mathbf{q}_{\nu}^{\left(\ell^{\prime}\right)}\right)\right]\left(\mathbf{h}_{\mu}^{(\ell)} \cdot \mathbf{h}_{\nu}^{(\ell)}\right) \\
& -\frac{1}{2} \sum_{\ell=2}^{d} m_{\ell, 1}^{2} \operatorname{tr}\left(G_{x x} Q^{(\ell)}\right)+\frac{1}{4} \frac{1}{n_{1}} \sum_{\ell, \ell^{\prime}=2}^{d} g_{\ell, \ell^{\prime}, 1} \operatorname{tr}\left(G_{x x} Q^{(\ell)} G_{x x} Q^{\left(\ell^{\prime}\right)}\right) \\
& +\frac{1}{2} \frac{1}{n_{1}} \sum_{\ell=2}^{d} \operatorname{tr}\left(G_{x x} \tilde{J}_{\ell, 1} G_{x x} Q^{(\ell)}\right) \\
& -\frac{1}{2} m_{1,0}^{2} \operatorname{tr}\left(G_{x x} J^{(1)}\right) \\
& +\frac{1}{4} \frac{1}{n_{1}} \sigma_{1,0}^{4} \operatorname{tr}\left(G_{x x} J^{(1)} G_{x x} J^{(1)}\right) \\
& +\mathcal{O}\left(n^{-2}\right) \tag{E.6}
\end{align*}
$$

where we have defined

$$
\begin{align*}
m_{\ell, 0}^{2} & \equiv \sigma_{\ell, 0}^{2}  \tag{E.7}\\
m_{\ell, 1}^{2} & \equiv m_{\ell, 0}^{2}+\sigma_{\ell, 1}^{2} m_{1,0}^{2}  \tag{E.8}\\
g_{\ell, \ell^{\prime}, 1} & \equiv \sigma_{\ell, 1}^{2} \sigma_{\ell^{\prime}, 1}^{2} \sigma_{1,0}^{4}, \quad \text { and }  \tag{E.9}\\
\tilde{J}_{\ell, 1} & \equiv \sigma_{\ell, 1}^{2} \sigma_{1,0}^{4} J^{(1)} \tag{E.10}
\end{align*}
$$

where $\ell, \ell^{\prime}>1$ for all cases but $m_{1,0}^{2}$. Assuming the network has more than one hidden layer, if we now again apply the results of Appendix D. 2 with

$$
\begin{align*}
G & =m_{2,1}^{2} G_{x x} \\
\mathbf{j}_{\mu} & =\mathbf{0} \\
A & =J^{(2)}+\sum_{\ell^{\prime}=3}^{d} \sigma_{\ell^{\prime}, 2}^{2} Q^{\left(\ell^{\prime}\right)}  \tag{E.11}\\
B & =G_{x x} \tilde{J}_{2,1} G_{x x}+\sum_{\ell^{\prime}=3}^{d} g_{\ell^{\prime}, 2,1} G_{x x} Q^{\left(\ell^{\prime}\right)} G_{x x}, \quad \text { and } \\
g & =g_{2,2,1} / m_{2,1}^{4}
\end{align*}
$$

we find that the effective action after integrating out the first two layers of the network is

$$
\begin{align*}
S^{(2)}= & -\beta \sum_{\mu=1}^{p} \varepsilon\left(\mathbf{h}_{\mu}^{(d)}, \mathbf{y}_{\mu}\right)+\sum_{\mu=1}^{p} \sum_{\ell=1}^{d} i \mathbf{q}_{\mu}^{(\ell)} \cdot \mathbf{h}_{\mu}^{(\ell)} \\
& -\frac{1}{2} \sum_{\ell=3}^{d-1} \frac{1}{n_{\ell}} \sum_{\mu, \nu=1}^{p}\left[J_{\mu \nu}^{(\ell)}+\sum_{\ell^{\prime}=\ell+1}^{d} \sigma_{\ell^{\prime}, \ell}^{2}\left(\mathbf{q}_{\mu}^{\left(\ell^{\prime}\right)} \cdot \mathbf{q}_{\nu}^{\left(\ell^{\prime}\right)}\right)\right]\left(\mathbf{h}_{\mu}^{(\ell)} \cdot \mathbf{h}_{\nu}^{(\ell)}\right) \\
& -\frac{1}{2} \sum_{\ell=3}^{d} m_{\ell, 2}^{2} \operatorname{tr}\left(G_{x x} Q^{(\ell)}\right)+\frac{1}{4} \frac{1}{n_{2}} \sum_{\ell, \ell^{\prime}=3}^{d} g_{\ell, \ell^{\prime}, 2} \operatorname{tr}\left(G_{x x} Q^{(\ell)} G_{x x} Q^{\left(\ell^{\prime}\right)}\right) \\
& +\frac{1}{2} \frac{1}{n_{2}} \sum_{\ell=3}^{d} \operatorname{tr}\left(G_{x x} \tilde{J}_{\ell, 2} G_{x x} Q^{(\ell)}\right) \\
& -\frac{1}{2} \sum_{\ell=1}^{2} m_{\ell, \ell-1}^{2} \operatorname{tr}\left(G_{x x} J^{(\ell)}\right) \\
& +\frac{1}{4} \frac{1}{n_{1}} m_{1,0}^{4} \operatorname{tr}\left(G_{x x} J^{(1)} G_{x x} J^{(1)}\right)+\frac{1}{4} \frac{1}{n_{2}}\left(m_{2,1}^{4}+\frac{n_{2}}{n_{1}} g_{2,2,1}\right) \operatorname{tr}\left(G_{x x} J^{(2)} G_{x x} J^{(2)}\right) \\
& +\frac{1}{2} \frac{1}{n_{1}} \operatorname{tr}\left(G_{x x} \tilde{J}_{2,1} G_{x x} J^{(2)}\right) \\
& +\mathcal{O}\left(n^{-2}\right) \tag{E.12}
\end{align*}
$$

where we now define

$$
\begin{align*}
m_{\ell, 2}^{2} & \equiv m_{\ell, 1}^{2}+m_{2,1}^{2} \sigma_{\ell, 2}^{2}  \tag{E.13}\\
g_{\ell, \ell^{\prime}, 2} & \equiv m_{2,1}^{4}+\frac{n_{2}}{n_{1}}\left(g_{\ell, \ell^{\prime}, 1}+g_{2,2,1} \sigma_{\ell, 2}^{2} \sigma_{\ell^{\prime}, 2}^{2}+g_{\ell, 2,1} \sigma_{\ell^{\prime}, 2}^{2}+\sigma_{\ell, 2}^{2} g_{2, \ell^{\prime}, 1}\right), \quad \text { and }  \tag{E.14}\\
\tilde{J}_{\ell, 2} & \equiv \frac{n_{2}}{n_{1}} \tilde{J}_{\ell, 1}+\left(m_{2,1}^{4}+\frac{n_{2}}{n_{1}} g_{2,2,1}\right) \sigma_{\ell, 2}^{2} J^{(2)}+\frac{n_{2}}{n_{1}} \sigma_{\ell, 2}^{2} \tilde{J}_{2,1}+\frac{n_{2}}{n_{1}} g_{\ell, 2,1} J^{(2)} \tag{E.15}
\end{align*}
$$

for $\ell, \ell^{\prime}>2$. We can now see that we can repeat this procedure to integrate out all of the hidden layers of the network, yielding an effective action of

$$
\begin{align*}
S^{(d-1)}= & -\beta \sum_{\mu=1}^{p} \varepsilon\left(\mathbf{h}_{\mu}^{(d)}, \mathbf{y}_{\mu}\right)+\sum_{\mu=1}^{p} i \mathbf{q}_{\mu}^{(d)} \cdot \mathbf{h}_{\mu}^{(d)} \\
& -\frac{1}{2} m_{d, d-1}^{2} \operatorname{tr}\left(G_{x x} Q^{(d)}\right)+\frac{1}{4} \frac{1}{n_{d-1}} g_{d, d, d-1} \operatorname{tr}\left(G_{x x} Q^{(d)} G_{x x} Q^{(d)}\right) \\
& +\frac{1}{2} \frac{1}{n_{d-1}} \operatorname{tr}\left(G_{x x} \tilde{J}_{d, d-1} G_{x x} Q^{(d)}\right) \\
& -\frac{1}{2} \sum_{\tau=1}^{d-1} m_{\tau, \tau-1}^{2} \operatorname{tr}\left(G_{x x} J^{(\tau)}\right) \\
& +\frac{1}{4} \sum_{\tau=2}^{d-1}\left(\frac{1}{n_{\tau}} m_{\tau, \tau-1}^{4}+\frac{1}{n_{\tau-1}} g_{\tau, \tau, \tau-1}\right) \operatorname{tr}\left(G_{x x} J^{(\tau)} G_{x x} J^{(\tau)}\right) \\
& +\frac{1}{2} \sum_{\tau=2}^{d-1} \frac{1}{n_{\tau-1}} \operatorname{tr}\left(G_{x x} \tilde{J}_{\tau, \tau-1} G_{x x} J^{(\tau)}\right) \\
& +\mathcal{O}\left(n^{-2}\right) \tag{E.16}
\end{align*}
$$

where the coupling constants and effective source obey the recurrences

$$
\begin{align*}
m_{\ell, \tau}^{2} \equiv & m_{\ell, \tau-1}^{2}+m_{\tau, \tau-1}^{2} \sigma_{\ell, \tau}^{2}  \tag{E.17}\\
g_{\ell, \ell^{\prime}, \tau} \equiv & m_{\tau, \tau-1}^{4} \sigma_{\ell, \tau}^{2} \sigma_{\ell^{\prime}, \tau}^{2} \\
& +\frac{n_{\tau}}{n_{\tau-1}}\left(g_{\ell, \ell^{\prime}, \tau-1}+g_{\tau, \tau, \tau-1} \sigma_{\ell, \tau}^{2} \sigma_{\ell^{\prime}, \tau}^{2}+g_{\ell, \tau, \tau-1} \sigma_{\ell^{\prime}, \tau}^{2}+\sigma_{\ell, \tau}^{2} g_{\tau, \ell^{\prime}, \tau-1}\right), \quad \text { and }  \tag{E.18}\\
\tilde{J}_{\ell, \tau} \equiv & \frac{n_{\tau}}{n_{\tau-1}} \tilde{J}_{\ell, \tau-1}+\frac{n_{\tau}}{n_{\tau-1}} \sigma_{\ell, \tau}^{2} \tilde{J}_{\tau, \tau-1} \\
& +\left(m_{\tau, \tau-1}^{4} \sigma_{\ell, \tau}^{2}+\frac{n_{\tau}}{n_{\tau-1}} g_{\tau, \tau, \tau-1} \sigma_{\ell, \tau}^{2}+\frac{n_{\tau}}{n_{\tau-1}} g_{\ell, \tau, \tau-1}\right) J^{(\tau)} \tag{E.19}
\end{align*}
$$

for $\ell, \ell^{\prime}>\tau$. Applying the results of Appendix D. 2 once more with

$$
\begin{align*}
G & =m_{d, d-1}^{2} G_{x x} \\
\mathbf{j}_{\mu} & =\beta \mathbf{y}_{\mu} \\
A & =\beta n_{d} I_{p}  \tag{E.20}\\
B & =G_{x x} \tilde{J}_{d, d-1} G_{x x}, \quad \text { and } \\
g & =g_{d, d, d-1} / m_{d, d-1}^{4}
\end{align*}
$$

we find the source-dependent terms in the logarithm of the partition function are

$$
\begin{align*}
\log Z \supset & -\frac{1}{2} \frac{1}{n_{1}} \beta^{2} n_{d} \operatorname{tr}\left(\Gamma^{-1} G_{x x} \tilde{J}_{d, d-1} G_{x x} \Gamma^{-1} G_{y y}\right)+\frac{1}{2} \frac{1}{n_{d-1}} \beta n_{d} \operatorname{tr}\left(\Gamma^{-1} G_{x x} \tilde{J}_{d, d-1} G_{x x}\right) \\
& -\frac{1}{2} \sum_{\tau=1}^{d-1} m_{\tau, \tau-1}^{2} \operatorname{tr}\left(G_{x x} J^{(\tau)}\right) \\
& +\frac{1}{4} \sum_{\tau=2}^{d-1}\left(\frac{1}{n_{\tau}} m_{\tau, \tau-1}^{4}+\frac{1}{n_{\tau-1}} g_{\tau, \tau, \tau-1}\right) \operatorname{tr}\left(G_{x x} J^{(\tau)} G_{x x} J^{(\tau)}\right) \\
& +\frac{1}{2} \sum_{\tau=2}^{d-1} \frac{1}{n_{\tau-1}} \operatorname{tr}\left(G_{x x} \tilde{J}_{\tau, \tau-1} G_{x x} J^{(\tau)}\right) \\
& +\mathcal{O}\left(n^{-2}\right) \tag{E.21}
\end{align*}
$$

where

$$
\begin{equation*}
\Gamma \equiv I_{p}+\beta m_{d, d-1}^{2} G_{x x} \tag{E.22}
\end{equation*}
$$

## E. 2 Computing the average hidden layer kernels

With the source-dependent terms of the relevant partition function in hand, we can compute the average hidden layer kernels for a feedforward linear network with arbitrary skip connections. We can immediately read off that

$$
\begin{align*}
\left\langle K^{(\ell)}\right\rangle= & m_{\ell, \ell-1}^{2} G_{x x} \\
& +\frac{n_{d}}{n_{d-1}} \operatorname{tr}\left[\left.\left(\beta^{2} G_{x x} \Gamma^{-1} G_{y y} \Gamma^{-1} G_{x x}-\beta G_{x x} \Gamma^{-1} G_{x x}\right) \frac{\delta \tilde{J}_{d, d-1}}{\delta J^{(\ell)}}\right|_{J^{(\ell)}=0}\right] \\
& +\mathcal{O}\left(n^{-2}\right) \tag{E.23}
\end{align*}
$$

hence our only task is to compute the derivative of the effective source $\tilde{J}_{d, d-1}$ with respect to the source for the $\ell$-th hidden layer. Singling out the $\ell$-th layer, we can set all sources except $J^{(\ell)}$ to zero. Then, the 'earliest' effective source to be non-zero is

$$
\begin{equation*}
\tilde{J}_{\ell^{\prime}, \ell}=\left(m_{\ell, \ell-1}^{4} \sigma_{\ell^{\prime}, \ell}^{2}+\frac{n_{\ell}}{n_{\ell-1}} g_{\ell, \ell, \ell-1} \sigma_{\ell^{\prime}, \ell}^{2}+\frac{n_{\ell}}{n_{\ell-1}} g_{\ell^{\prime}, \ell, \ell-1}\right) J^{(\ell)} \tag{E.24}
\end{equation*}
$$

for $\ell^{\prime}>\ell$, and the recurrence relation for $\tau>\ell$ is

$$
\begin{equation*}
\tilde{J}_{\ell^{\prime}, \tau}=\frac{n_{\tau}}{n_{\tau-1}}\left(\tilde{J}_{\ell^{\prime}, \tau-1}+\sigma_{\ell^{\prime}, \tau}^{2} \tilde{J}_{\tau, \tau-1}\right) \tag{E.25}
\end{equation*}
$$

From the form of these recurrences, we can see that

$$
\begin{align*}
\left\langle K^{(\ell)}\right\rangle= & m_{\ell, \ell-1}^{2} G_{x x} \\
& +\frac{n_{d}}{n_{d-1}} \tilde{g}_{\ell} G_{x x}\left(\beta^{2} \Gamma^{-1} G_{y y} \Gamma^{-1}-\beta \Gamma^{-1}\right) G_{x x} \\
& +\mathcal{O}\left(n^{-2}\right) \tag{E.26}
\end{align*}
$$

where $\tilde{g}_{\ell}$ is a layer-dependent scalar. Even without explicitly solving the recurrences to obtain $\tilde{g}_{\ell}$, this shows that Conjecture 1 holds perturbatively for linear networks with arbitrary skip connections. We leave detailed study of these recurrences-and therefore of the precise dependence of the corrections on width, depth, and skip connection structure-as an interesting objective for future work.

## F Comparison to the results of Aitchison [10] and Li and Sompolinsky [11]

In this appendix, we compare our results for the average kernels of deep linear networks to those of Aitchison [10] and Li and Sompolinsky [11].

## F. 1 Comparison to the results of Aitchison [10]

We first show that our result (9) for the low-temperature limit of the average kernels of a deep linear network can be recovered from the results of Aitchison [10]. Working in what corresponds to the zero-temperature limit of our setup, Aitchison derives the following implicit recurrence

$$
\begin{equation*}
0=-\left(n_{\ell+1}-n_{\ell}\right)\left(K^{(\ell)}\right)^{-1}+n_{\ell+1}\left(K^{(\ell)}\right)^{-1}\left(K^{(\ell+1)}\right)\left(K^{(\ell)}\right)^{-1}-n_{\ell}\left(K^{(\ell-1)}\right)^{-1} \tag{F.1}
\end{equation*}
$$

for $\ell=1, \ldots, d-1$, where the boundary conditions of the recurrence are $K^{(0)}=G_{x x}$ and $K^{(d)}=G_{y y}$. We will self-consistently solve this recurrence relation in the limit $n_{1}, \ldots, n_{d-1} \rightarrow \infty$, $n_{0}, n_{d}, p=\mathcal{O}(1)$. Concretely, we make the ansatz that the zero-temperature kernels are of the form

$$
\begin{equation*}
K^{(\ell)}=K_{\infty}^{(\ell)}+\frac{1}{n_{\ell}} K_{1}^{(\ell)}+\mathcal{O}\left(n_{\ell}^{-2}\right) \tag{F.2}
\end{equation*}
$$

and solve the recurrence relations order-by-order using the resulting Neumann series

$$
\begin{equation*}
\left(K^{(\ell)}\right)^{-1}=\left(K_{\infty}^{(\ell)}\right)^{-1}-\frac{1}{n_{\ell}}\left(K_{\infty}^{(\ell)}\right)^{-1} K_{1}^{(\ell)}\left(K_{\infty}^{(\ell)}\right)^{-1}+\mathcal{O}\left(n_{\ell}^{-2}\right) \tag{F.3}
\end{equation*}
$$

The leading-order recurrence is simply

$$
\begin{equation*}
0=\left(1-\frac{n_{\ell+1}}{n_{\ell}}\right)\left(K_{\infty}^{(\ell)}\right)^{-1}+\frac{n_{\ell+1}}{n_{\ell}}\left(K_{\infty}^{(\ell)}\right)^{-1}\left(K_{\infty}^{(\ell+1)}\right)\left(K_{\infty}^{(\ell)}\right)^{-1}-\left(K_{\infty}^{(\ell-1)}\right)^{-1} \tag{F.4}
\end{equation*}
$$

with boundary conditions $K_{\infty}^{(0)}=G_{x x}$ and $K_{\infty}^{(d)}=G_{y y}$. For the last hidden layer, we have $n_{\ell+1} / n_{\ell}=n_{d} / n_{d-1} \rightarrow 0$, hence the recurrence reduces to

$$
\begin{equation*}
K_{\infty}^{(d-1)}=K_{\infty}^{(d-2)} \tag{F.5}
\end{equation*}
$$

If we iterate this procedure backwards through the network, it is easy to see that the $n_{\ell+1} / n_{\ell^{-}}$ dependent terms at each layer will cancel, leaving

$$
\begin{equation*}
K_{\infty}^{(d-1)}=K_{\infty}^{(d-2)}=\cdots=K_{\infty}^{(1)}=G_{x x} \tag{F.6}
\end{equation*}
$$

We now consider the leading finite-width correction. For the last hidden layer, we obtain

$$
\begin{equation*}
0=n_{d}\left(G_{y y}-G_{x x}\right)-K_{1}^{(d-1)}+\frac{n_{d-1}}{n_{d-2}} K_{1}^{(d-2)} \tag{F.7}
\end{equation*}
$$

after dropping all terms that are of $\mathcal{O}\left(n^{-2}\right)$ and multiplying on the left and right by $G_{x x}$. For the first hidden layer, we have

$$
\begin{equation*}
0=K_{1}^{(2)}-\left(1+\frac{n_{2}}{n_{1}}\right) K_{1}^{(1)} \tag{F.8}
\end{equation*}
$$

Finally, for intermediate hidden layers (i.e., $\ell=2,3, \ldots, d-2$ ), we have

$$
\begin{equation*}
0=K_{1}^{(\ell+1)}-\left(1+\frac{n_{\ell+1}}{n_{\ell}}\right) K_{1}^{(\ell)}+\frac{n_{\ell}}{n_{\ell-1}} K_{1}^{(\ell-1)} \tag{F.9}
\end{equation*}
$$

Based on the form of these recurrences, we make the ansatz that the solution is of the form

$$
\begin{equation*}
K_{1}^{(\ell)}=n_{d} a_{\ell}\left(G_{y y}-G_{x x}\right) \tag{F.10}
\end{equation*}
$$

for some sequence $a_{\ell}$, where we assume that $G_{y y} \neq G_{x x}$. Then, the recurrence for the last hidden layer is satisfied provided that

$$
\begin{equation*}
a_{d-1}=1+\frac{n_{d-1}}{n_{d-2}} a_{d-2} \tag{F.11}
\end{equation*}
$$

those for the intermediate layers if

$$
\begin{equation*}
0=a_{\ell+1}-\left(1+\frac{n_{\ell+1}}{n_{\ell}}\right) a_{\ell}+\frac{n_{\ell}}{n_{\ell-1}} a_{\ell-1} \tag{F.12}
\end{equation*}
$$

and that for the first hidden layer if

$$
\begin{equation*}
a_{2}=\left(1+\frac{n_{2}}{n_{1}}\right) a_{1} \tag{F.13}
\end{equation*}
$$

Substituting the expression for $a_{d-1}$ into the condition resulting from the recurrence relation centered on $a_{d-2}$, we find that we must have

$$
\begin{equation*}
a_{d-2}=1+\frac{n_{d-2}}{n_{d-3}} a_{d-3}, \tag{F.14}
\end{equation*}
$$

hence we can iterate this process backwards to the second hidden layer, yielding

$$
\begin{equation*}
a_{\ell}=1+\frac{n_{\ell}}{n_{\ell-1}} a_{\ell-1} \tag{F.15}
\end{equation*}
$$

for $\ell=2,3, \ldots, d-1$. Then, the condition relating $a_{2}$ and $a_{1}$ resulting from the recurrence relation for the first layer implies that we must have $a_{1}=1$. Thus, we recover our zero-temperature result from solving Aitchison's recurrence relations order-by-order.

## F. 2 Comparison to the results of Li and Sompolinsky [11]

We now show that our result (9) for the low-temperature limit of the average kernels of a deep linear network can be recovered as a limiting case of the result of Li and Sompolinsky [11]. Their result for the zero-temperature kernel in the limit $n_{0}, n, p \rightarrow \infty$ with $n_{1}=n_{2}=\cdots=n_{d-1}=n$, $n_{0} / n \in(0, \infty), \alpha \equiv p / n \in(0, \infty)$, and $\sigma_{1}=\cdots=\sigma_{d}=\sigma$ is, in our notation,

$$
\begin{equation*}
\sigma^{-2(\ell+1)}\left\langle K^{(\ell)}\right\rangle \sim\left(1-\frac{n_{d}}{n}\right)^{\ell} G_{x x}+\frac{1}{n} \sigma^{-2 d} Y V M_{\ell} V^{\top} Y^{\top}, \tag{F.16}
\end{equation*}
$$

where $Y \in \mathbb{R}^{p \times n_{d}}$ is the matrix of targets and $M_{\ell} \in \mathbb{R}^{n_{d} \times n_{d}}$ is a diagonal matrix with non-zero elements

$$
\begin{equation*}
\left[M_{\ell}\right]_{k k}=z_{k}^{-(d-1)} \frac{z_{k}^{\ell}-1}{z_{k}-1} \tag{F.17}
\end{equation*}
$$

Here, the orthogonal matrix $V$ is the matrix of eigenvectors of

$$
\begin{equation*}
R=\frac{1}{\sigma^{2} p} Y^{\top} G_{x x}^{+} Y=V \Omega V^{\top} \tag{F.18}
\end{equation*}
$$

for $G_{x x}^{+}$the pseudoinverse of $G_{x x}$, and the scalars $z_{k}$ are in turn defined in terms of the eigenvalues $\Omega_{k k}=\omega_{k}$ as

$$
\begin{equation*}
1-\alpha=z_{k}-\alpha \sigma^{-2(d-1)} z_{k}^{-(d-1)} \omega_{k} ; \tag{F.19}
\end{equation*}
$$

we note that Li and Sompolinsky [11] use variables $u_{k 0}=\sigma^{2} z_{k}$.
As we are interested in the limit $\alpha \downarrow 0$, it is useful to write the implicit equation for $z_{k}$ as

$$
\begin{equation*}
z_{k}=1+\alpha\left(\sigma^{-2 L} z_{k}^{-(d-1)} \omega_{k}-1\right) \tag{F.20}
\end{equation*}
$$

hence we expect $z_{k} \rightarrow 1$ as $\alpha \downarrow 0$. Thus, we have

$$
\begin{equation*}
\left[M_{\ell}\right]_{k k} \rightarrow \ell \tag{F.21}
\end{equation*}
$$

which gives

$$
\begin{equation*}
V M_{\ell} V^{\top} \rightarrow \ell I_{n_{d}} \tag{F.22}
\end{equation*}
$$

Using the expansion $\left(1-n_{d} / n\right)^{\ell}=1-n_{d} \ell / n+\mathcal{O}\left(n^{-2}\right)$, we therefore find that

$$
\begin{equation*}
\sigma^{-2(\ell+1)}\left\langle K^{(\ell)}\right\rangle \sim G_{x x}+\frac{n_{d} \ell}{n}\left(\sigma^{-2 d} G_{y y}-G_{x x}\right) \tag{F.23}
\end{equation*}
$$

in the limit in which $n_{d} / n \downarrow 0$ and $p / n \downarrow 0$. Therefore, combining this result with that of the previous subsection, our result (9) agrees with those of Aitchison [10] and of Li and Sompolinsky [11] in the appropriate limit. Whether the full result of Li and Sompolinsky [11] agrees with that of Aitchison [10] is an interesting question, but is well beyond the scope of the present work.

## G Predictor statistics and generalization in deep linear networks

Though the main focus of our work is on the asymptotics of representation learning, we have also computed the leading finite-width corrections to the predictor statistics. Though one can derive the analogy of Conjecture 1 for the predictor statistics of a general BNN with linear readout, the resulting formula is not particularly illuminating. We will therefore present results only for linear networks. As was true of the hidden layer kernels of deep linear networks, this calculation can be performed either using methods similar to those described in Appendix B or Appendix D. As the steps are largely identical to those calculations, we only briefly summarize the results.
In short, we fix a test dataset $\hat{\mathcal{D}}=\left\{\left(\hat{\mathbf{x}}_{\mu}, \hat{\mathbf{y}}_{\mu}\right)\right\}_{\mu=1}^{\hat{p}}$ of $\hat{p}$ examples, and define the Gram matrices

$$
\begin{align*}
\left(G_{\hat{x} \hat{x}}\right)_{\hat{\mu} \hat{\nu}} & \equiv n_{0}^{-1} \hat{\mathbf{x}}_{\hat{\mu}} \cdot \hat{\mathbf{x}}_{\hat{\nu}},  \tag{G.1}\\
\left(G_{\hat{y} \hat{y}}\right)_{\hat{\mu} \hat{\nu}} & \equiv n_{d}^{-1} \hat{\mathbf{y}}_{\hat{\mu}} \cdot \hat{\mathbf{y}}_{\hat{\nu}},  \tag{G.2}\\
\left(G_{x \hat{x}}\right)_{\mu \hat{\mu}} & \equiv n_{0}^{-1} \mathbf{x}_{\mu} \cdot \hat{\mathbf{x}}_{\hat{\mu}}, \quad \text { and }  \tag{G.3}\\
\left(G_{y \hat{y}}\right)_{\mu \hat{\nu}} & \equiv n_{d}^{-1} \mathbf{y}_{\mu} \cdot \hat{\mathbf{y}}_{\hat{\nu}} \tag{G.4}
\end{align*}
$$

Introducing appropriate source terms to allow us to compute predictor statistics, we then proceed perturbatively as before, assuming that the combined input Gram matrix

$$
\left[\begin{array}{ll}
G_{x x} & G_{x \hat{x}}  \tag{G.5}\\
G_{x \hat{x}}^{\top} & G_{\hat{x} \hat{x}}
\end{array}\right]
$$

is invertible. Again, the final result can be extended to the case in which this matrix is not invertible by a continuity argument.
Our notation in this appendix will follow that of Appendix B rather than Appendix D in that we will introduce matrices

$$
\begin{align*}
K_{\infty} & \equiv \sigma_{1}^{2} \cdots \sigma_{d-1}^{2} G_{x x}  \tag{G.6}\\
\hat{R}_{\infty} & \equiv \sigma_{1}^{2} \cdots \sigma_{d-1}^{2} G_{x \hat{x}}, \quad \text { and }  \tag{G.7}\\
\hat{K}_{\infty} & \equiv \sigma_{1}^{2} \cdots \sigma_{d-1}^{2} G_{\hat{x} \hat{x}} \tag{G.8}
\end{align*}
$$

to denote the blocks of the infinite-width kernel of the last hidden layer, rather than introducing scalar parameters to represent the products of variances. This will make our expressions somewhat more compact than they would be under the conventions of Appendix D.

## G. 1 Predictor statistics

Defining the matrix $\hat{F}_{\hat{\mu} j} \equiv f_{j}\left(\hat{\mathbf{x}}_{\hat{\mu}}\right)$, we find that the mean predictor can be written compactly as

$$
\begin{equation*}
\langle\hat{F}\rangle=\hat{R}_{\infty}^{\top}\left[\Gamma^{-1}-\frac{1}{\beta \sigma_{d}^{2}}\left(\sum_{\ell=1}^{d-1} \frac{1}{n_{\ell}}\right) \Gamma^{-1} M \Gamma^{-1}\right] Y+\mathcal{O}\left(n^{-2}\right) \tag{G.9}
\end{equation*}
$$

for

$$
\begin{equation*}
M \equiv \Gamma^{-1} K_{\infty}+\operatorname{tr}\left(\Gamma^{-1} K_{\infty}\right) I_{p}-n_{d}\left(\sigma_{d}^{-2} \Gamma^{-1} G_{y y} \Gamma^{-1}-\Gamma^{-1}\right) K_{\infty} \tag{G.10}
\end{equation*}
$$

The predictor covariance is given as

$$
\begin{align*}
& \sigma_{d}^{-2} \operatorname{cov}\left(\hat{F}_{\hat{\mu} j}, \hat{F}_{\hat{\nu} k}\right) \\
&=\left(\hat{K}_{\infty}-\hat{R}_{\infty}^{\top} \Gamma^{-1} \hat{R}_{\infty}\right)_{\hat{\mu} \hat{\nu}} \delta_{j k} \\
&+\left(\sum_{\ell=1}^{d-1} \frac{1}{n_{\ell}}\right) {\left[\hat{M}_{\hat{\mu} \hat{\nu}} \delta_{j k}\right.} \\
&+\sigma_{d}^{-2}\left(Y^{\top} \Gamma^{-1} K_{\infty} \Gamma^{-1} Y\right)_{j k}\left(\hat{K}_{\infty}-\hat{R}_{\infty}^{\top} \Gamma^{-1} \hat{R}_{\infty}\right)_{\hat{\mu} \hat{\nu}} \\
&-\frac{1}{\beta \sigma_{d}^{4}}\left(Y^{\top} \Gamma^{-1} K_{\infty} \Gamma^{-1} Y\right)_{j k}\left(\hat{R}_{\infty}^{\top} \Gamma^{-2} \hat{R}_{\infty}\right)_{\hat{\mu} \hat{\nu}} \\
&\left.+\frac{1}{\beta^{2} \sigma_{d}^{6}}\left(Y^{\top} \Gamma^{-2} \hat{R}_{\infty}\right)_{j \hat{\nu}}\left(Y^{\top} \Gamma^{-2} \hat{R}_{\infty}\right)_{k \hat{\mu}}\right]
\end{align*}
$$

for

$$
\begin{align*}
\hat{M} \equiv & -\operatorname{tr}\left(\Gamma^{-1} K_{\infty}\right)\left(\hat{K}_{\infty}-\hat{R}_{\infty}^{\top} \Gamma^{-1} \hat{R}_{\infty}\right)+\frac{1}{\beta \sigma_{d}^{2}} \operatorname{tr}\left(\Gamma^{-1} K_{\infty}\right) \hat{R}_{\infty}^{\top} \Gamma^{-2} \hat{R}_{\infty}-\frac{1}{\beta^{2} \sigma_{d}^{4}} \hat{R}_{\infty}^{\top} \Gamma^{-3} \hat{R}_{\infty} \\
& +n_{d} \frac{1}{\beta^{2} \sigma_{d}^{4}} \hat{R}_{\infty}^{\top} \Gamma^{-1}\left(\sigma_{d}^{-2} \Gamma^{-1} G_{y y} \Gamma^{-1}-\Gamma^{-1}\right) \Gamma^{-1} \hat{R}_{\infty} \tag{G.12}
\end{align*}
$$

The mean and covariance of the training set predictor $F_{\mu j} \equiv f_{j}\left(\mathbf{x}_{\mu}\right)$ can be obtained by setting $\hat{R}_{\infty}$ and $\hat{K}_{\infty}$ to $K_{\infty}$ in the above expressions.

## G. 2 Bias-variance decompositions and the low-temperature limit

These results allow us to define thermal bias-variance decompositions of the form

$$
\begin{equation*}
\langle E\rangle=\frac{1}{2} \sum_{\mu=1}^{p}\left\|\left\langle\mathbf{f}\left(\mathbf{x}_{\mu}\right)\right\rangle-\mathbf{y}_{\mu}\right\|_{2}^{2}+\frac{1}{2} \sum_{\mu=1}^{p} \sum_{k=1}^{n_{d}} \operatorname{cov}\left[f_{k}\left(\mathbf{x}_{\mu}\right), f_{k}\left(\mathbf{x}_{\mu}\right)\right] \equiv E_{b}+E_{v} \tag{G.13}
\end{equation*}
$$

for the mean training and test errors. However, the resulting expressions are not particularly illuminating except in the low-temperature limit $\beta \rightarrow \infty$. We will focus on the regime in which $G_{x x}$ (and thus $K_{\infty}$ ) is invertible, in which the underlying linear system $X W=Y$ is underdetermined and the training set can be interpolated. In this regime, $\Gamma^{-1}=K_{\infty}^{-1}+\mathcal{O}\left(\beta^{-1}\right)$, and the mean predictor reduces to the least-norm pseudoinverse solution to the linear system, with mean training and test predictions of

$$
\begin{equation*}
\langle F\rangle=Y+\mathcal{O}\left(\beta^{-1}\right) \tag{G.14}
\end{equation*}
$$

and

$$
\begin{equation*}
\langle\hat{F}\rangle=\hat{R}_{\infty}^{\top} K_{\infty}^{-1} Y+\mathcal{O}\left(\beta^{-1}\right)=G_{\hat{x} x}^{\top} G_{x x}^{-1} Y+\mathcal{O}\left(\beta^{-1}\right)=\hat{X} X^{\top}\left(X X^{\top}\right)^{-1} Y+\mathcal{O}\left(\beta^{-1}\right) \tag{G.15}
\end{equation*}
$$

respectively. The training and test set covariances have low-temperature limits of

$$
\begin{equation*}
\operatorname{cov}\left(F_{\mu j}, F_{\nu k}\right)=\mathcal{O}\left(\beta^{-1}\right) \tag{G.16}
\end{equation*}
$$

and

$$
\begin{align*}
& \operatorname{cov}\left(\hat{F}_{\hat{\mu} j}, \hat{F}_{\hat{\nu} k}\right)=\sigma_{d}^{2}\left(\hat{K}_{\infty}-\hat{R}_{\infty}^{\top} K_{\infty}^{-1} \hat{R}_{\infty}\right)_{\hat{\mu} \hat{\nu}}\left[\delta_{j k}+\left(\sum_{\ell=1}^{d-1} \frac{1}{n_{\ell}}\right)\left(\sigma_{d}^{-2} Y^{\top} K_{\infty}^{-1} Y-p I_{n_{2}}\right)_{j k}\right] \\
& +\mathcal{O}\left(\beta^{-1}, n^{-2}\right) \tag{G.17}
\end{align*}
$$

respectively. Then, it is easy to see that both $E_{b}$ and $E_{v}$ are $\mathcal{O}\left(\beta^{-1}\right)$, while

$$
\begin{equation*}
\hat{E}_{b}=\frac{1}{2}\left\|\hat{R}_{\infty}^{\top} K_{\infty}^{-1} Y-\hat{Y}\right\|_{F}^{2}+\mathcal{O}\left(\beta^{-1}\right) \tag{G.18}
\end{equation*}
$$

and

$$
\begin{align*}
& \hat{E}_{v}=n_{d} \sigma_{d}^{2} \operatorname{tr}\left(\hat{K}_{\infty}-\hat{R}_{\infty}^{\top} K_{\infty}^{-1} \hat{R}_{\infty}\right)\left[1+\left(\sum_{\ell=1}^{d-1} \frac{1}{n_{\ell}}\right)\left(\sigma_{d}^{-2} \operatorname{tr}\left(K_{\infty}^{-1} G_{y y}\right)-p\right)\right] \\
&+\mathcal{O}\left(\beta^{-1}, n^{-2}\right) \tag{G.19}
\end{align*}
$$

Thus, at least to leading order, width affects the low-temperature test error only through the variance term. Substituting in the definition of $K_{\infty}$, we find that to leading order the test error decreases with increasing width if

$$
\begin{equation*}
\frac{1}{p} \operatorname{tr}\left(G_{x x}^{-1} G_{y y}\right)>\sigma_{1}^{2} \cdots \sigma_{d}^{2} \tag{G.20}
\end{equation*}
$$

and increases with increasing width otherwise. This small-initialization condition is the generalization of that found by Li and Sompolinsky [11] to our asymptotic regime.

## G. 3 Effects of alternative regularization temperature-dependence

In this appendix, we comment on the possibility of alternative temperature-dependent posteriors. This possibility arises from the interpretation of the Bayes posterior (B.4) as the equilibrium distribution of the Langevin dynamics

$$
\begin{equation*}
d \Theta^{(\ell)}(t)=-\left(\lambda(\beta) \Sigma \Theta+\nabla_{\Theta} E\right) d t+\sqrt{2 \beta^{-1}} d B^{(\ell)}(t) \tag{G.21}
\end{equation*}
$$

at inverse temperature $\beta$, where $B^{(\ell)}(t)$ is a standard Wiener process, $\Sigma$ is the diagonal matrix of prior variances, and $\lambda(\beta)=1 / \beta$. As elsewhere, we focus on the regime in which the training dataset can be linearly interpolated, in which the thermal variance of the test set predictions need not vanish. Moreover, it suffices to consider only the GP contributions; the finite-width corrections computed above do not change the qualitative results. In these statistics, the case of general $\lambda(\beta)$ is related to $\lambda(\beta)=1 / \beta$ by the replacement

$$
\begin{equation*}
\sigma_{1}^{2} \cdots \sigma_{d}^{2} \leftarrow \frac{\sigma_{1}^{2} \cdots \sigma_{d}^{2}}{\beta^{d} \lambda(\beta)^{d}} \tag{G.22}
\end{equation*}
$$

Then, if we assume a low-temperature power-law dependence $\lambda(\beta) \sim \beta^{\omega}$ for simplicity, we find that the zero-temperature limits of the training set predictor mean and covariance are

$$
\lim _{\beta \rightarrow \infty}\langle F\rangle= \begin{cases}0 & \omega>1 / d-1  \tag{G.23}\\ K_{\infty}\left(\sigma_{d}^{-2} I_{p}+K_{\infty}\right)^{-1} Y & \omega=1 / d-1 \\ Y & \omega<1 / d-1\end{cases}
$$

and

$$
\begin{equation*}
\lim _{\beta \rightarrow \infty} \operatorname{cov}\left(F_{\mu j}, F_{\nu k}\right)=0 \tag{G.24}
\end{equation*}
$$

respectively, while those of the test set mean and covariance are

$$
\lim _{\beta \rightarrow \infty}\langle\hat{F}\rangle= \begin{cases}0 & \omega>1 / d-1  \tag{G.25}\\ \hat{R}_{\infty}^{\top}\left(\sigma_{d}^{-2} I_{p}+K_{\infty}\right)^{-1} Y & \omega=1 / d-1 \\ \hat{R}_{\infty}^{\top} K_{\infty}^{-1} Y & \omega<1 / d-1\end{cases}
$$

and

$$
\lim _{\beta \rightarrow \infty} \operatorname{cov}\left(\hat{F}_{\hat{\mu} j}, \hat{F}_{\hat{\nu} k}\right)= \begin{cases}0 & \omega>-1  \tag{G.26}\\ \sigma_{d}^{2}\left(\hat{K}_{\infty}-\hat{R}^{\top} K_{\infty}^{-1} \hat{R}_{\infty}\right)_{\hat{\mu} \hat{\nu}} \delta_{j k} & \omega=-1 \\ \infty & \omega<-1\end{cases}
$$

respectively. Therefore, taking $\lambda(\beta)=1 / \beta$ yields sensible zero-temperature infinite-width behavior for a linear network of any depth in the underdetermined regime.

## H Derivation of the average kernels for a depth-two network

In this appendix, we derive the average feature kernel for a network with a single (possibly nonlinear) hidden layer and a linear readout. This derivation is a simple extension of the perturbative derivation of Conjecture 1 in Appendix B, using the fact that the size of the terms in the expansion for two-layer networks can be directly controlled in terms of the inverse hidden layer width.
Concretely, we consider a network defined as

$$
\begin{align*}
\mathbf{h}^{(1)} & =\frac{\sigma_{1}}{\sqrt{n_{0}}} W^{(1)} \mathbf{x}  \tag{H.1}\\
\mathbf{h}^{(2)} & =\frac{\sigma_{2}}{\sqrt{n_{1}}} W^{(2)} \phi\left(\mathbf{h}^{(1)}\right)  \tag{H.2}\\
\mathbf{f} & =\mathbf{h}^{(2)} . \tag{H.3}
\end{align*}
$$

Our task is to control the prior cumulants of the hidden layer feature kernel

$$
\begin{equation*}
K_{\mu \nu} \equiv \frac{1}{n_{1}} \phi\left(\mathbf{h}_{\mu}^{(1)}\right) \cdot \phi\left(\mathbf{h}_{\nu}^{(1)}\right) . \tag{H.4}
\end{equation*}
$$

We can use the fact that the rows $\left[\mathbf{w}_{j}^{(1)}\right]^{\top}$ of $W^{(1)}$ are independent and identically distributed under the prior to obtain

$$
\begin{align*}
{\left[K_{\infty}\right]_{\mu \nu} } & =\mathbb{E}_{\mathcal{W}} K_{\mu \nu}  \tag{H.5}\\
& =\frac{1}{n_{1}} \sum_{j=1}^{n_{1}} \mathbb{E}_{\mathbf{w}_{j}^{(1)}}\left[\phi\left(\frac{\sigma_{1}}{\sqrt{n_{0}}} \mathbf{w}_{j}^{(1)} \cdot \mathbf{x}_{\mu}\right) \phi\left(\frac{\sigma_{1}}{\sqrt{n_{0}}} \mathbf{w}_{j}^{(1)} \cdot \mathbf{x}_{\nu}\right)\right]  \tag{H.6}\\
& =\mathbb{E}\left[\phi\left(h_{\mu}^{(1)}\right) \phi\left(h_{\nu}^{(1)}\right): \mathbf{h}^{(1)} \sim \mathcal{N}\left(\mathbf{0}, \sigma_{1}^{2} G_{x x}\right)\right] \tag{H.7}
\end{align*}
$$

at any hidden layer width $[12,13]$. Similarly, we can easily see that

$$
\begin{equation*}
\operatorname{cov}_{\mathcal{W}}\left(K_{\mu \nu}, K_{\rho \lambda}\right)=\frac{1}{n_{1}}\left(\mathbb{E}\left[\phi\left(h_{\mu}^{(1)}\right) \phi\left(h_{\nu}^{(1)}\right) \phi\left(h_{\rho}^{(1)}\right) \phi\left(h_{\lambda}^{(1)}\right)\right]-\left[K_{\infty}\right]_{\mu \nu}\left[K_{\infty}\right]_{\rho \lambda}\right), \tag{H.8}
\end{equation*}
$$

where $\mathbf{h}^{(1)} \sim \mathcal{N}\left(\mathbf{0}, \sigma_{1}^{2} G_{x x}\right)$, and that higher cumulants are $\mathcal{O}\left(n_{1}^{-2}\right)$. Then, we can directly apply the result of Appendix B to conclude that

$$
\begin{equation*}
\left\langle K_{\mu \nu}\right\rangle=\left[K_{\infty}\right]_{\mu \nu}+\frac{1}{2} n_{d} \sum_{\rho, \lambda=1}^{p}\left(\sigma_{d}^{-2} \Gamma^{-1} G_{y y} \Gamma^{-1}-\Gamma^{-1}\right)_{\rho \lambda} \operatorname{cov}_{\mathcal{W}}\left(K_{\mu \nu}, K_{\rho \lambda}\right)+\mathcal{O}\left(n_{1}^{-2}\right) \tag{H.9}
\end{equation*}
$$

for $\Gamma=\sigma_{1}^{2} K_{\infty}+I_{p} / \beta \sigma_{2}^{2}$. Depending on the nonlinearity, this result may be continuous in $G_{x x}$, and therefore extensible to the non-invertible case via a continuity argument. In particular, as noted in Appendix D, this holds for a linear network.

To gain some intuition for how different choices of nonlinear activation function affect the learned representations, we consider the case in which $G_{x x}$ is diagonal. In this special case, the four-point term simplifies dramatically. In particular, we have

$$
\begin{equation*}
\left(K_{\infty}\right)_{\mu \nu}=\operatorname{var}\left[\phi\left(h_{\mu}^{(1)}\right)\right] \delta_{\mu \nu}+\mathbb{E}\left[\phi\left(h_{\mu}^{(1)}\right)\right] \mathbb{E}\left[\phi\left(h_{\nu}^{(1)}\right)\right] \tag{H.10}
\end{equation*}
$$

and

$$
\begin{align*}
\operatorname{cov}_{\mathcal{W}}\left(K_{\mu \nu}, K_{\rho \lambda}\right)=\frac{1}{n_{1}}( & \operatorname{var}\left[\phi\left(h_{\mu}^{(1)}\right)^{2}\right] \delta_{\mu \nu} \delta_{\mu \rho} \delta_{\mu \lambda} \\
& \left.+\operatorname{var}\left[\phi\left(h_{\mu}^{(1)}\right)\right] \operatorname{var}\left[\phi\left(h_{\nu}^{(1)}\right)\right]\left(1-\delta_{\mu \nu}\right)\left(\delta_{\mu \rho} \delta_{\nu \lambda}+\delta_{\mu \lambda} \delta_{\nu \rho}\right)\right) \tag{H.11}
\end{align*}
$$

which yields

$$
\begin{align*}
&\left\langle K_{\mu \nu}\right\rangle=\left(K_{\infty}\right)_{\mu \nu}+ \frac{1}{2} \\
& \frac{n_{2}}{n_{1}}\left(\sigma_{2}^{-2} \Gamma^{-1} G_{y y} \Gamma^{-1}-\Gamma^{-1}\right)_{\mu \nu}  \tag{H.12}\\
& \times\left[\operatorname{var}\left[\phi\left(h_{\mu}^{(1)}\right)^{2}\right] \delta_{\mu \nu}+2 \operatorname{var}\left[\phi\left(h_{\mu}^{(1)}\right)\right] \operatorname{var}\left[\phi\left(h_{\nu}^{(1)}\right)\right]\left(1-\delta_{\mu \nu}\right)\right]+\mathcal{O}\left(n_{1}^{-2}\right)
\end{align*}
$$

Moreover, applying the Sherman-Morrison formula [3], we have

$$
\begin{equation*}
\frac{1}{\beta \sigma_{2}^{2}} \Gamma_{\mu \nu}^{-1}=\frac{\delta_{\mu \nu}}{\gamma_{\mu}}-\frac{1}{1+\sum_{\rho=1}^{p} \mathbb{E}\left[\phi\left(h_{\rho}^{(1)}\right)\right]^{2} / \gamma_{\rho}} \frac{\mathbb{E}\left[\phi\left(h_{\mu}^{(1)}\right)\right]}{\gamma_{\mu}} \frac{\mathbb{E}\left[\phi\left(h_{\nu}^{(1)}\right)\right]}{\gamma_{\nu}} \tag{H.13}
\end{equation*}
$$

where we have defined the vector $\gamma_{\mu} \equiv 1+\beta \sigma_{2}^{2} \operatorname{var}\left[\phi\left(h_{\mu}^{(1)}\right)\right]$ for brevity. Thus, in this simple setting, activation functions with $\mathbb{E} \phi(h) \neq 0$ yield qualitatively different behavior from those with $\mathbb{E} \phi(h)=0$ : non-vanishing $\mathbb{E} \phi(h)$ introduces a rank-1 component in the GP kernel, which in turn couples elements of $G_{y y}$ in the leading finite-width correction.

## I Numerical methods

In this appendix, we describe the numerical methods used in our experiments. We perform our simulations by sampling network parameters at each time step of the Langevin update G. 21 after some large burn-in period when the loss function stabilizes around a fixed number. We used EulerMaruyama method [14] to obtain the discretized Langevin equation:

$$
\begin{equation*}
\Theta(t+1)-\Theta(t)=-\beta^{-1} \Theta(t) d t-\nabla_{\Theta} E(t) d t+\xi \sqrt{2 \beta^{-1} d t} \tag{I.1}
\end{equation*}
$$

where $\xi \sim \mathcal{N}(0,1)$ is a standard Gaussian random variable sampled i.i.d. at each time step and $d t$ is the time step. The first, second and last terms represent the weight decay, the gradient descent update and the stochastic Wiener process, respectively. We also used stochastic gradient methods for posterior sampling such as Stochastic Gradient Langevin Dynamics [15, 16], although obtained the best results via full-batch gradient updates.

We used the Neural Tangents framework [17] and PyTorch deep learning library [18] to generate the neural networks and trained them according to the discretized full-batch Langevin update rule. A typical burn-in time was $\sim 2 \times 10^{6}$ iterations and after that the parameters were sampled over $\sim 2 \times 10^{6}$ iterations where we chose a learning rate of $d t \sim 10^{-4}$. Simulations have been performed on a cluster with NVIDIA Tesla V100 GPU's with 32 GB RAM and a typical simulation run took $\sim 2-6 \mathrm{hr}$ depending on the architecture and the network width. All code used throughout this work can be reached at https://github.com/Pehlevan-Group/finite-width-bayesian/.
All figures shown here are results of a single instance of a trained neural network on a fixed dataset. Since we performed all our experiments with $\beta=1$, we observed that the different initializations of a network did not influence the final posterior mean due to the weight decay and long burn-in periods.
Throughout all experiments, the MNIST digits were downsized from $28 \times 28$ pixels to $10 \times 10$ pixels without distorting the original digits. This was done to accelerate the training process since large input dimensions would take an order of magnitude more time to obtain well estimated posterior means. We considered 10 -dimensional outputs corresponding to one-hot encoded digits. Both inputs and labels were ordered according to their class. Figure 2 shows an example of MNIST digits and the input $G_{x x}$ and output $G_{y y}$ Gram matrices.

## References

[1] Leon Isserlis. On a formula for the product-moment coefficient of any order of a normal frequency distribution in any number of variables. Biometrika, 12(1/2):134-139, 1918.
[2] Gian-Carlo Wick. The evaluation of the collision matrix. Physical Review, 80(2):268, 1950.
[3] Roger A Horn and Charles R Johnson. Matrix Analysis. Cambridge University Press, 2012.
[4] 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.
[5] Alexander G. de 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-.
[6] Roman Novak, Lechao Xiao, Yasaman Bahri, Jaehoon Lee, Greg Yang, Daniel A. Abolafia, Jeffrey Pennington, and Jascha Sohl-Dickstein. Bayesian deep convolutional networks with many channels are Gaussian processes. In International Conference on Learning Representations, 2019. URL https://openreview.net/forum?id=B1g30j0qF7.
[7] Greg Yang. Scaling limits of wide neural networks with weight sharing: Gaussian process behavior, gradient independence, and neural tangent kernel derivation. arXiv preprint arXiv:1902.04760, 2019.
[8] Sho Yaida. Non-Gaussian processes and neural networks at finite widths. In Jianfeng Lu and Rachel Ward, editors, Proceedings of The First Mathematical and Scientific Machine Learning Conference, volume 107 of Proceedings of Machine Learning Research, pages 165192, Princeton University, Princeton, NJ, USA, July 2020. PMLR. URL http: //proceedings . mlr.press/v107/yaida20a.html.
[9] Lechao Xiao, Yasaman Bahri, Jascha Sohl-Dickstein, Samuel Schoenholz, and Jeffrey Pennington. Dynamical isometry and a mean field theory of CNNs: How to train 10,000-layer vanilla convolutional neural networks. In International Conference on Machine Learning, pages 5393-5402. PMLR, 2018.
[10] Laurence Aitchison. Why bigger is not always better: on finite and infinite neural networks. In Hal Daumé III and Aarti Singh, editors, Proceedings of the 37th International Conference on Machine Learning, volume 119 of Proceedings of Machine Learning Research, pages 156-164. PMLR, July 2020. URL http://proceedings.mlr.press/v119/aitchison20a.html.
[11] Qianyi Li and Haim Sompolinsky. Statistical mechanics of deep linear neural networks: The backpropagating kernel renormalization. Phys. Rev. X, 11:031059, Sep 2021. doi: 10. 1103/PhysRevX.11.031059. URL https://link.aps.org/doi/10.1103/PhysRevX.11. 031059.
[12] Radford M Neal. Priors for infinite networks. In Bayesian Learning for Neural Networks, pages 29-53. Springer, 1996.
[13] Christopher KI Williams. Computing with infinite networks. Advances in Neural Information Processing Systems, pages 295-301, 1997.
[14] Peter E Kloeden and Eckhard Platen. Stochastic differential equations. In Numerical Solution of Stochastic Differential Equations, pages 103-160. Springer, 1992.
[15] M. Welling and Y. Teh. Bayesian learning via stochastic gradient langevin dynamics. In ICML, 2011.
[16] Chunyuan Li, Changyou Chen, David Carlson, and Lawrence Carin. Preconditioned stochastic gradient langevin dynamics for deep neural networks. In Thirtieth AAAI Conference on Artificial Intelligence, 2016.
[17] Roman Novak, Lechao Xiao, Jiri Hron, Jaehoon Lee, Alexander A. Alemi, Jascha SohlDickstein, and Samuel S. Schoenholz. Neural tangents: Fast and easy infinite neural networks in python. In International Conference on Learning Representations, 2020. URL https: //github.com/google/neural-tangents.
[18] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, highperformance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'AlchéBuc, E. Fox, and R. Garnett, editors, Advances in Neural Information Processing Systems, volume 32. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper/ 2019/file/bdbca288fee7f92f2bfa9f7012727740-Paper.pdf.
[19] Yann LeCun, Corinna Cortes, and CJ Burges. MNIST handwritten digit database. ATT Labs [Online]. Available: http://yann.lecun.com/exdb/mnist, 2, 2010.

## Checklist

1. For all authors...
(a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
(b) Did you describe the limitations of your work? [Yes] As discussed in $\S 3$ and the Conclusion, the primary limitation of our work is that our main result is obtained through a non-rigorous perturbative argument. We therefore frame its most general form as a conjecture.
(c) Did you discuss any potential negative societal impacts of your work? [N/A] We do not anticipate that our work will have negative societal impacts as outlined in the ethics guidelines, as it is purely theoretical.
(d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
(a) Did you state the full set of assumptions of all theoretical results? [Yes] See §2 and the statement of Conjecture 1.
(b) Did you include complete proofs of all theoretical results? [Yes] See Appendices B, C, C , D , and H . As noted above, we frame our main result as a conjecture due to the fact that some of our methods are not fully rigorous.
3. If you ran experiments...
(a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] Our code is publicly available at https://github.com/Pehlevan-Group/ finite-width-bayesian/.
(b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] All experimental details have been detailed in supplemental material.
(c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]
(d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See details in Appendix I of the supplemental material.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
(a) If your work uses existing assets, did you cite the creators? [Yes] We use the MNIST dataset [19], which we cite appropriately. We also use open-source libraries such as PyTorch [18], which we also cite.
(b) Did you mention the license of the assets? [Yes] We mention licenses in Appendix I of the supplemental material.
(c) Did you include any new assets either in the supplemental material or as a URL? [N/A] Our work involves no new assets.
(d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
(a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]
