BEYOND SPECTRA: EIGENVECTOR OVERLAPS IN LOSS GEOMETRY

000

001

002 003 004

010 011

012

013

014

015

016

017

018

019

021

025

026

027

028

029

031

033

035

037

040

041

042

043

044

045

046

047

048

051

052

Paper under double-blind review

ABSTRACT

Local loss geometry in machine learning is fundamentally a two-operator concept. When only a single loss is considered, geometry is fully summarized by the Hessian spectrum; in practice, however, both training and test losses are relevant, and the resulting geometry depends on their spectra together with the alignment of their eigenspaces. We first establish general foundations for two-loss geometry by formulating a universal local fluctuation law, showing that the expected testloss increment under small training perturbations is a trace that combines train and test spectral data with a critical additional factor quantifying eigenspace overlap, and by proving a novel transfer law that describes how overlaps transform in response to noise. As a solvable analytical model, we next apply these laws to ridge regression with arbitrary covariate shift, where operator-valued free probability yields asymptotically exact overlap decompositions that reveal overlaps as the natural quantities specifying shift and that resolve the puzzle of multiple descent: peaks are controlled by eigenspace (mis-)alignment rather than by Hessian ill-conditioning alone. Finally, for empirical validation and scalability, we confirm the fluctuation law in multilayer perceptrons, develop novel algorithms based on subspace iteration and kernel polynomial methods to estimate overlap functionals, and apply them to a ResNet-20 trained on CIFAR10, showing that class imbalance reshapes train-test loss geometry via induced misalignment. Together, these results establish overlaps as the critical missing ingredient for understanding local loss geometry, providing both theoretical foundations and scalable estimators for analyzing generalization in modern neural networks.

1 Introduction

Modern learning algorithms are inherently local, and sources of randomness (stochastic gradients, finite-sample variability, and distributional drift) are often small relative to the underlying signal. A local quadratic approximation to the loss thus provides a natural setting for analyzing learning. This observation underpins the considerable literature studying loss geometry via Hessians, and in particular Hessian spectra (see references below). When the focus is a single loss, local geometry is indeed fully captured by the Hessian spectrum. Crucially, however, in machine learning contexts there are (at least) *two* losses of interest - train and test - and so local loss geometry involves two quadratic approximations. The joint geometry of these approximations is not captured by Hessian spectra alone; it requires a critical additional ingredient: *eigenvector alignment*, or overlaps.

Despite the fundamental importance of eigenvector overlaps, most studies to date have centered on Hessian eigenvalue distributions - often explicitly equating spectra with loss geometry. The literature is extensive and examines Hessians from several complementary angles, including: (i) empirical measurement of eigenvalue distributions and their training-time evolution, with links to optimization stability (Sagun et al., 2017; Ghorbani et al., 2019; Yao et al., 2019); (ii) random-matrix-theoretic and mean field models (Pennington & Bahri, 2017; Pennington & Worah, 2018; 2019; Liao & Mahoney, 2021; Karakida et al., 2019); (iii) class- and layer-structured spectral phenomena, such as identifiable outliers tied to data and architecture (Papyan, 2020; Sankar et al., 2021); and (iv) Hessian/Fisher analyses relating sharpness (as measured by eigenvalue magnitude) to stability and generalization (Keskar et al., 2017; Foret et al., 2021; Cohen et al., 2021; Yao et al., 2019). These studies give fundamental insight into aspects of loss curvature, but ignore directional information that becomes relevant as soon as one compares two operators.

The need to go beyond spectra is well understood in random matrix theory, where eigenvector consistency and overlaps with population directions are central objects in spiked models and correlated ensembles (Johnstone, 2001; Paul, 2007; Nadler, 2008; Benaych-Georges & Nadakuditi, 2011; Bun et al., 2017; Landau et al., 2023). There, eigenvalues alone do not determine statistical performance; rather, risk depends on how sample and population eigenvectors align. Our perspective adapts this principle to learning-theoretic questions by focusing on train–test alignment. We show that overlap measures between the training fluctuations covariance (intimately related to the training Hessian; see Results) and the test Hessian yield a decomposition of generalization error into components associated with eigenvector alignment.

Applying this perspective to ridge regression resolves the puzzle of anisotropy-induced multiple descent (see Chen et al. (2021); Li & Wei (2021); Mel & Ganguli (2021); Meng et al. (2023) for several distinct forms of multiple descent). High-dimensional analyses of double/multiple descent often focus on the connection between interpolation peaks and the eigenvalue distribution of the design matrix (Singh et al., 2022; Chen & Mei, 2022). However, in anisotropic settings where error exhibits multiple peaks despite monotonically decreasing minimum training eigenvalue, spectra alone do not explain generalization. By making overlaps explicit, we show that the presence and location of multiple descent peaks are governed by the alignment between train and test Hessian eigenspaces, rather than by the monotone behavior of the smallest eigenvalue itself. This corrects interpretations that implicitly attribute sample-wise multiple descent purely to spectrum ill-conditioning (Chen & Mei, 2022; Mel & Ganguli, 2021; Mel & Pennington, 2022), and suggests a simple geometric picture for anisotropy-induced generalization error that may prove useful for understanding more complex models.

A second arena where eigenvector orientation matters is generalization under covariate shift. Recent work derives high-dimensional risk formulas under shift for the random feature model (Tripuraneni et al., 2021) and proposes regularizers that encourage cross-domain invariance in gradient, Fisher, or Hessian statistics (Rame et al., 2022; Hemati et al., 2023). Our analysis complements these views by decomposing the effect of covariate shift into eigenvalue and eigenvector components. Holding spectra fixed, changes in overlap structure alone can increase or decrease test risk; the formalism therefore predicts when shift will help or hurt, and by how much, in a way that a spectrum-only analysis cannot.

We next perform controlled experiments on multilayer perceptrons to probe the quadratic/noise regime: near a minimizer, stochastic gradient fluctuations induce parameter perturbations that are filtered by the (inverse) training Hessian, and the ensuing test-loss increment is determined by the alignment between these fluctuations and the eigenspaces of the test Hessian. We directly observe this "train-Hessian filtering" effect, validating the fluctuation-overlap formulas derived from the local two-loss setting.

Translating our overlap-centric theory into practice at modern scale requires algorithms that go beyond spectral density estimation. A substantial literature has developed linear algebraic tools for implicit matrices, including polynomial/quadrature approaches and stochastic trace methods such as Hutchinson and Lanczos-based quadrature (Golub & Meurant, 2009; Lin et al., 2016; Ubaru et al., 2017). These and related techniques have been adapted to deep learning to estimate Hessian spectral densities and extremal eigenpairs efficiently (Adams et al., 2018; Papyan, 2019; Ghorbani et al., 2019; Yao et al., 2019). Building on these foundations, we develop novel estimators for overlap functionals between pairs of Hessians (train-test, population-sample), and apply these to demonstrate how class imbalance impacts train-test loss geometry.

The resulting picture is that local geometry in machine learning is fundamentally bivariate - spectra *and* overlaps: spectra characterize the curvatures of train and test losses, while eigenvector overlaps determine how these curvatures combine to produce test error.

2 Contributions

1. **Two-loss theory of local geometry.** We introduce a novel two-loss framework for local loss geometry that incorporates both spectra and overlaps 3.1, rectifying a widespread oversimplification that equates spectra with geometry.

- 2. **General foundations** We derive and test a universal local fluctuation law showing how overlaps impact generalization 3.1.1, and a general transfer law dictating how eigenvector overlaps are transformed by noise 3.1.2.
- 3. **Explicit formulas for high-dimensional ridge regression.** Combining tools from random matrix theory with our overlap transfer law, we provide closed-form expressions for the overlap function between train and test Hessians in anisotropic ridge regression 3.2.
- 4. **Unified explanation of covariate shift and multiple descent.** We show that covariate shift is naturally quantified by eigenvector overlaps 3.2.1, and that overlaps analytically resolve the puzzle of multiple descent 3.2.2.
- 5. **Empirical validation in neural networks.** We confirm our theoretical predictions in multilayer perceptrons, and use overlap machinery to show that the training Hessian acts as a filter shaping optimization 3.3.
- 6. **Scalable algorithms for Hessian overlaps.** We develop novel, scalable numerical methods for estimating Hessian eigenvector overlaps in large-scale models, enabling practical use of our theory in modern deep learning 3.4.
- 7. **Train–test misalignment under class imbalance.** We show that class imbalance in CI-FAR induces misalignment between train and test Hessians, explaining the effects of class imbalance in terms of train-test loss geometry. 3.4.

3 RESULTS

3.1 Theoretical foundations

3.1.1 A UNIVERSAL LOCAL FLUCTUATION LAW

Let $w \in \mathbb{R}^d$ denote the d-dimensional parameter vector of a model f_w . The parameters are selected by minimizing a training objective $J_{\text{train}}(w)$, while performance is measured by a test objective $J_{\text{test}}(w)$. The local geometry is thus characterized by two local quadratic approximations, one for J_{train} , which determines optimization dynamics, and one for J_{test} , which measures how changes to w impact generalization.

How do these approximations combine to determine model performance? Consider a pretrained model w_0, f_{w_0} at the minimum of the unperturbed loss $J_{\text{train}}(w,0)$. Generically, the effect of a perturbation to the loss, $J_{\text{train}}(w,\epsilon)$, will be to create a small gradient displacing the minimum. Write $z := d \nabla J_{\text{train}}(w_0,\epsilon)$ for the gradient and $H_{\text{train}} := d \nabla^2 J_{\text{train}}(w_0,0)$ for the Hessian. (These scalings ensure relevant quantities are O(1).) Quadratic approximation of J_{train} gives

$$J_{\text{train}}(w,\epsilon) \approx J_{\text{train}}(w_0,\epsilon) + \frac{1}{d}z \cdot (w - w_0) + \frac{1}{2d}(w - w_0)^{\top} H_{\text{train}}(w - w_0). \tag{1}$$

Minimizing the perturbed training objective yields a noise-induced fluctuation $\Delta w = -H_{\rm train}^{-1}z$, so that retraining can be seen as filtering the perturbation by $H_{\rm train}^{-1}$. We will refer to the perturbation gradient z as the *injected noise* and the post-learning fluctuation Δw as the *filtered noise*.

After retraining, loss is measured through the test loss J_{test} . Writing $H_{\text{test}} := d \nabla^2 J_{\text{test}}(w_0)$ for the test Hessian, quadratic approximation of J_{test} yields

$$\Delta J_{\text{test}} \approx -\nabla J_{\text{test}} \left(w_0 \right)^{\top} H_{\text{train}}^{-1} z + \frac{1}{2d} z^{\top} H_{\text{train}}^{-1} H_{\text{test}} H_{\text{train}}^{-1} z \tag{2}$$

Equation (2) represents the simplest model capturing the interaction of nontrivial train and test geometry in the context of noisy learning. Of particular interest is the case of centered noise $\mathbb{E}\,z=0$, so that the mean ΔJ_{test} reduces to the average of the quadratic term only. Writing $C_{\mathrm{train}}:=\mathbb{E}\,\Delta w \Delta w^{\top}=\mathbb{E}\,(H_{\mathrm{train}}^{-1}z)(H_{\mathrm{train}}^{-1}z)^{\top}$ for the covariance of the post-learning fluctuations, we have $\mathbb{E}\,\Delta J_{\mathrm{test}}=\frac{1}{2d}\mathrm{tr}[H_{\mathrm{test}}C_{\mathrm{train}}]$. This simple expression already suggests the importance of alignment between directions of large training fluctuation and directions of large test Hessian eigenvalue. Writing $\Delta J:=\mathbb{E}\Delta J_{\mathrm{test}}$, this last trace can be rewritten exactly as the integral

$$\Delta J = \frac{1}{2} \iint \lambda_1 \, \lambda_2 \, O(\lambda_1, \lambda_2) \, \mu_{\text{test}}(d\lambda_1) \, \mu_{\text{train}}(d\lambda_2), \tag{3}$$

where μ_{train} , μ_{test} are the spectral measures of H_{test} , C_{train} , and $O(\lambda_1, \lambda_2)$ is the average squared overlap between eigenspaces at eigenvalues λ_1, λ_2 (see Appendix B for explicit formula). Equation 3, which puts training fluctuations λ_2 , test sensitivity λ_1 , and overlaps of the corresponding directions $O(\lambda_1, \lambda_2)$, on equal footing, attests to the fundamental role of overlaps in loss geometry.

3.1.2 OVERLAP TRANSFER LAW

Often, in settings involving two *population* operators A, B, model behavior depends on the eigenspace overlaps of A and a noisy *sample* operator \tilde{B} . Informally, one needs a way of combining the given population overlaps $O_{A,B}$ with the noise, specified by $O_{B,\tilde{B}}$. We prove the following appealing transfer law in Appendix B.1:

Proposition 1 (Free transfer law for overlap functions). Let $\tilde{B} = F(B, X)$ be a matrix rational expression. If X is free from A, B, then

$$O_{A,\tilde{B}}(a,\tilde{b}) = \int O_{A,B}(a,b)O_{B,\tilde{B}}(b,\tilde{b})\,\mu_B(db)\,. \tag{4}$$

Proposition 1 entails a simple overlap calculus that can be used to compute overlap functions of complex matrix models from simpler ones. In Appendix C, we use (4) to quickly derive expressions for train-test Hessian overlap functions in anisotropic ridge regression.

3.2 HESSIAN OVERLAPS GOVERN GENERALIZATION IN LINEAR REGRESSION

Let training inputs $x \in \mathbb{R}^d$ have covariance $\Sigma_{\text{train}} := \mathbb{E}[xx^{\top}]$, and let the ground-truth output be linear with Gaussian label noise:

$$y(x) = \frac{1}{\sqrt{d}} w_*^\top x + \xi, \qquad \xi \sim \mathcal{N}(0, \sigma^2).$$

We will also assume for convenience that $w_* \sim \mathcal{N}(0,I)$, so that the signal to noise ratio is $\operatorname{tr}\Sigma_{\operatorname{train}}/(\operatorname{tr}\Sigma_{\operatorname{train}}+\sigma^2)$ (tr denotes the dimension normalized trace). Given a training set consisting of $X \in \mathbb{R}^{m \times d}$ (rows x^\top) and labels $y \in \mathbb{R}^m$, ridge regression chooses $w \in \mathbb{R}^d$ to minimize

$$J_{\text{train}}(w) = \frac{1}{2m} \|y - \frac{1}{\sqrt{d}} Xw\|^2 + \frac{\lambda}{2d} \|w\|^2, \qquad \lambda \ge 0.$$
 (5)

The (excess) test loss is measured with test inputs with covariance Σ_{test} :

$$J_{\text{test}}(w) := \frac{1}{2} \mathbb{E}_{x,\xi} \left[\left(\frac{1}{\sqrt{d}} w^{\top} x - y(x) \right)^{2} \right] - \frac{1}{2} \sigma^{2} = \frac{1}{2d} (w - w_{*})^{\top} \Sigma_{\text{test}} (w - w_{*}).$$

With the scalings of 3.1.1, the train and test Hessians are $H_{\text{train}} = A + \lambda I$ and $H_{\text{test}} = \Sigma_{\text{test}}$, where $A := X^{\top} X/m$ is the training set sample covariance. (Note this approaches Σ_{train} for large m/d.)

We now apply the local fluctuation formula (3), which in the setting of ridge regression is exact. Letting the label noise supply the perturbation, the injected noise $z := \nabla_w J_{\text{train}}(w_0, \xi)$ and parameter perturbation covariance $C_{\text{train}} := \mathbb{E}\left(H_{\text{train}}^{-1}z\right)(H_{\text{train}}^{-1}z)^{\top}$ are

$$z = -X^{\top} \xi / m \sqrt{d}, \qquad C_{\text{train}} = \sigma^2 q A (A + \lambda I)^{-2},$$
 (6)

The test loss increment ΔJ is obtained by substituting into the overlap formula (3). The training-side operators A, $H_{\rm train} = A + \lambda I$ and $C_{\rm train}$ commute and share eigenvectors, so for simplicity all formulas are written in terms of A's eigenvalues and vectors:

$$\Delta J = \frac{\sigma^2 q}{2} \iint \lambda_1 \frac{\lambda_2}{(\lambda_2 + \lambda)^2} O(\lambda_1, \lambda_2) \,\mu_{\Sigma_{\text{test}}}(d\lambda_1) \,\mu_A(d\lambda_2), \tag{7}$$

where $\mu_{\Sigma_{\mathrm{test}}}$ and μ_A are the empirical spectral measures of Σ_{test} and A, and $O(\lambda_1,\lambda_2)$ is the eigenvector-overlap function between their eigenbases (see (13)). Since we will be interested primarily in the ridgeless limit $\lambda \to 0$, we will loosely refer to $A = \frac{1}{m} X^\top X$ as the train Hessian.

The fundamental conclusion from (7) that we will apply toward analyzing covariate shift and multiple descent is that error is amplified when there is significantly alignment $O(\lambda_1, \lambda_2)$ between a

high-variance direction of the learned parameters (small training eigenvalue λ_2) and a high-curvature direction of the test loss (large test eigenvalue λ_1).

In Appendix C, we use techniques from operator-valued free probability to derive asymptotically exact expressions for J_{test} , ΔJ_{test} and the overlap function $O_{H_{\mathrm{train}},H_{\mathrm{test}}}$ in proportional asymptotics where $m,d\to\infty$ with q:=d/m fixed. The main conceptual contribution of this work is that while the spectral densities of train/test operators set the relevant scales, it is their relative orientation - as quantified by the overlap function - that determines how fluctuations translate into test loss. We illustrate these points in two settings: first, a simple covariate shift experiment that provides geometric intuition and positions $O_{\Sigma_{\mathrm{test}},\Sigma_{\mathrm{train}}}$ as the natural object quantifying shift; second, the puzzle of multiple descent (cf. Mel & Ganguli (2021)), where the overlap function allows a full analytical account. For both settings, we use the simplest possible model of anisotropic data: the "two-scale" covariance with spectral measure

$$\mu_{\Sigma} := p_1 \delta_{s_1}^2 + p_2 \delta_{s_2}^2. \tag{8}$$

3.2.1 COVARIATE SHIFT INCREASES LOSS THROUGH TRAIN-TEST MISALIGNMENT

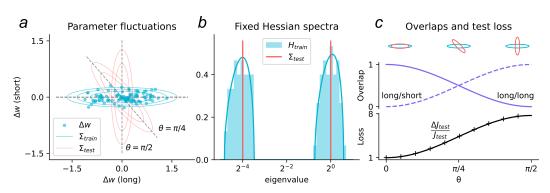


Figure 1: Covariate shift and test error. a) Points: individual parameter fluctuations due to label noise. Cyan lines: fluctuation covariance $C_{\rm train} \approx \Sigma_{\rm train}^{-1}$. Red lines: different test Hessians. θ measures the angle between the large test and training eigenspaces. b) Eigenspaces are rotated while Hessian spectra are fixed. Blue line shows $d,m\to\infty$ theory. c) For small θ , large fluctuations in learned parameters are aligned with the low-eigenvalue directions of the test Hessian (aligned ellipses at top left, and purple overlap traces), and error is small (black trace). For large θ , large fluctuation directions are aligned with sensitive directions of the test loss and test error is large (black trace). Traces: theory; crosses: average from simulations. $d, \alpha, \lambda, \sigma = 10^2, 10, 10^{-4}, 10^{-1/2}$.

Equation (7) expresses ΔJ in terms of the overlap function $O_{\Sigma_{\mathrm{test}},A}$. Relative to the population overlap $O_{\Sigma_{\mathrm{test}},\Sigma_{\mathrm{train}}}$, this overlap is deformed by the finite sampling ratio of the training set $\alpha=m/d$, in a way that is captured exactly by the free transfer law of Proposition 1. In Appendix C.3 we use the transfer law to state an explicit formula for $O_{\Sigma_{\mathrm{test}},A}$, and then show that ΔJ can be written

$$\Delta J = \frac{1}{2} q \sigma^2 \tilde{\lambda}' \int \lambda_1 \frac{\lambda_2}{(\tilde{\lambda} + \lambda_2)^2} O_{\Sigma_{\text{test}}, \Sigma_{\text{train}}} (\lambda_1, \lambda_2) d\mu_{\Sigma_{\text{test}}} (\lambda_1) d\mu_{\Sigma_{\text{train}}} (\lambda_2), \qquad (9)$$

which parallels (7) but averages out the random training inputs and label noise to express ΔJ purely in terms of the population operators $\Sigma_{\rm train}, \Sigma_{\rm test}$ (see Appendix C.3 for definition of $\tilde{\lambda}$). This expression illustrates how $O_{\Sigma_{\rm test},\Sigma_{\rm train}}$ - as the only quantity in (9) that can change under isospectral transformations to $\Sigma_{\rm train},\Sigma_{\rm test}$ - is the natural object for quantifying covariate shift.

To illustrate this point, we perform a simple experiment where both $\Sigma_{\rm train}$, $\Sigma_{\rm test}$ have fixed two-level spectra (equation (8)) with scales $s_1, s_2 = 2^0, 2^{-4}$ and equal multiplicities. $\lambda = 10^{-4}$ and $\alpha = m/d = 10$ so that $H_{\rm train} \approx A \approx \Sigma_{\rm train}$, while $H_{\rm test} = \Sigma_{\rm test}$. Fig. 1a shows the distribution of learned parameters for different label noise realizations. As predicted, fluctuations have larger variance along long directions of $C_{\rm train} \approx \sigma^2 q \, \Sigma_{\rm train}^{-1}$, corresponding to low-curvature directions of the train Hessian. At the same time, the test loss contours are determined by the test Hessian $\Sigma_{\rm test}$. We construct a controlled perturbation in which $\Sigma_{\rm test}$ is systematically rotated with respect to $\Sigma_{\rm train}$

while all spectra are kept fixed (panel b), isolating the effect of *overlaps*. Fig. 1c demonstrates the consequence of varying overlap. When the long directions of $\Sigma_{\rm train}$ align with the long directions of $\Sigma_{\rm test}$ ($\theta=0$), fluctuations occur in directions where the test error is relatively flat, yielding low excess test loss (Fig. 1c, left column). In contrast, when the same train-long directions align with test-short directions ($\theta=\pi/2$), the same magnitude of parameter fluctuation is heavily penalized, and the test loss rises sharply (Fig. 1c, right column). This simple experiment illustrates the central role of eigenvector overlaps in the context of covariate shift.

3.2.2 Multiple descent is explained by train-test overlaps

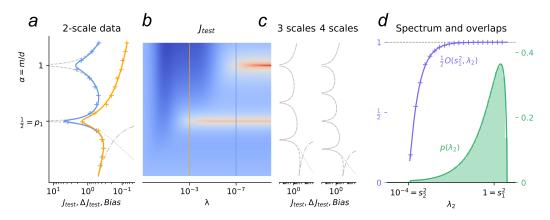


Figure 2: Multiple descent in ridge regression. a) $J_{\text{train}}, \Delta J_{\text{test}}$ and bias as a function of the measurement density $\alpha = m/d$ for two-scale data. Note the peaks at critical values of $\alpha = 1/2, 1$. Solid traces: theory; crosses: simulations with d = 5000. Dashed, dotted, and dash-dotted lines show theory error increment, and bias in the limit that the lower scale $s_2 \to 0$, where bumps become true singularities. b) Theory $J_{\text{test}}(\alpha,\lambda)$. Traces in panel a) correspond to gold and blue lines. c) 3 and 4 scale data which exhibit 3 and 4 peaks; legend same as a). d) Green histogram: empirical spectral density of the train Hessian $X^\top X/m$ at $\alpha = 0.496$; solid green line: theory. Purple line: overlap function, $O(s_1^2,\lambda_2)/2$, giving overlap between a train eigenspace at eigenvalue λ_2 with a the entire large-eigenvalue test space (ie. s_1^2). Note strong overlap for high train/test eigenspaces.

Double descent is a well-established phenomenon in machine learning in which test error exhibits a non-monotonic dependence on model size. More recently, several authors have described an extension of this effect, termed multiple descent, which arises in settings where input data are highly anisotropic and the covariance spectrum contains multiple separated scales (see introduction). Fig. 2a,b illustrate multiple descent for two-scale data with $s_1, s_2 = 1, 10^{-2}$, while panel c shows how a larger number of separated scales can create additional peaks in test error (see caption for details).

For a two-level covariance, the overlap function is determined by the solution to a cubic polynomial that is easily solved numerically (Appendix C). Fig. 2d shows the spectrum of the train Hessian (green histogram and theory line), and the overlap function (13), indicating overlap of a training eigenspace at eigenvalue λ_2 with the large-eigenvalue (s_1^2) eigenspace of the test Hessian. Theoretical and empirical overlaps are in excellent agreement (purple line and crosses).

The peaks of multiple descent are easily understood in terms of eigenvector overlaps. Fig. 3 reports the error, training spectrum, and overlap map for the two-scale covariance model of Fig. 2. The test-loss curve shows two singularities at critical sampling densities $\alpha=m/d$ (panel a). At the same densities the training spectrum undergoes phase transitions: at $\alpha=1/2$ an initially unimodal density splits into two bands centered near s_1^2 and s_2^2 , and at $\alpha=1$ the lower s_2^2 band develops a near-zero component (Fig. 3b). The corresponding overlap map $O(\lambda_1,\lambda_2)$ is approximately block-diagonal: modes near s_1^2 align predominantly with the s_1^2 test subspace, and modes near s_2^2 with the s_2^2 subspace (Fig. 3c). Thus, the first error spike occurs when near-null training directions overlap the sharp test subspace, whereas the second arises when an even smaller training component overlaps the flat subspace but with variance large enough to dominate its small curvature. Fig. 3d provides a geometric cartoon of the alignment of top and bottom eigenspaces of $H_{\rm train}$, $H_{\rm test}$ throughout this sequence. Until line 5, the minimum eigenvalue of $H_{\rm train}$ always decreases as a function of

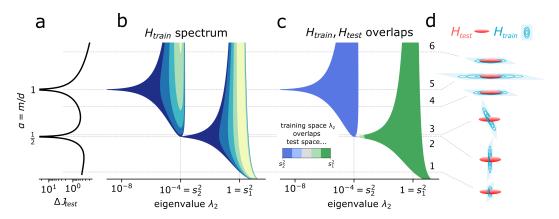


Figure 3: Multiple descent is explained by train-test Hessians overlaps. a) ΔJ due to label noise $(\lambda,s_2^2\to 0)$ limits taken to illustrate true singularities; see Appendix C.5.1 for details). b) Spectral density of $H_{\rm train}$ as a function of α (each density normalized to a maximum of 1). For high sampling density α , the training density acquires two components roughly centered on the true underlying data scales s_1^2, s_2^2 . c) The overlap function of $H_{\rm train}, H_{\rm test}$ as a function of α . For small α , all nonzero $H_{\rm train}$ spaces overlap strongly with the large eigenspace of $H_{\rm test}$. As α approaches the first critical value, a new spectral component appears, whose eigenspaces overlap almost entirely with the small eigenspace of $H_{\rm test}$. d) Minimal model of train-test Hessian geometry. Cyan plane and ellipses represent the top and bottom eigenspaces of $H_{\rm train}$. Red ellipsoid represents level sets of test error. Error is controlled by both train fluctuation magnitude and overlap onto test spaces.

 α - which, according to a spectrum-only analysis, should increase test error. Yet the error actually *decreases* between horizontal lines 3 and 4, precisely because the lowest train eigenspaces begin to overlap predominately with the low test eigenspace.

Summarizing, multiple descent arises from the interplay of (i) training components developing near-zero eigenvalues as α varies, and (ii) which test directions these overlap with - sharp or flat, illustrating the potentially extreme impact of (mis-)aligned train and test loss geometry.

3.3 LOCAL THEORY PREDICTS MLP GENERALIZATION AND LEARNING DYNAMICS

To test the applicability of the quadratic two-loss theory to nonlinear networks, we trained small, constant width multilayer perceptrons (MLPs) to reproduce the responses of an MLP teacher network. Student networks were batch trained for a large number of iterations to ensure near convergence to the noiseless training loss minimum. Noise was then added to the training set and the network was trained further - beginning from the initial trained state to determine the effect of the noise on the initial local minimum. After training, the training loss increment was computed and compared to prediction of the local quadratic theory. Fig. 4a,b show the measured test loss increment against the local quadratic prediction for several orders of magnitude of input (a) and label (b) noise strength. All later panels refer to the noise setting corresponding to the red point in (b).

Fig. 4c illustrates inverse Hessian filtering due to training dynamics. The gradient noise induced by the label noise has covariance $\mathbb{E}[zz^{\top}]$. Purple scatter represents the overlap function of $\mathbb{E}[zz^{\top}]$ and H_{train} . Dot x,y position is given by H_{train} , $\mathbb{E}[zz^{\top}]$ eigenvalue and size is proportional to overlap. Note strong alignment between gradient noise and test Hessian. After training, the parameter fluctuation covariance predicted by quadratic approximation is $C_{\text{train}} := \mathbb{E}\left(H_{\text{train}}^{-1}z\right)(H_{\text{train}}^{-1}z)^{\top}$. The overlap function of H_{train} and the actual post-training covariance is plotted in red. Note how in accordance with quadratic predictions, variance is strongly inflated/attenuated along low/high eigendirections of H_{train} - a phenomenon we refer to as inverse Hessian filtering. The large fluctuations do not translate into large test error since the train and test Hessians are well aligned (Fig. 6), meaning fluctuations occur primarily along low test Hessian (loss-insensitive) directions.

Loss landscape slices are shown in Fig. 4d for $J_{\text{train}}(w,0)$ (the unperturbed loss), $J_{\text{train}}(w,\epsilon)$ (the perturbed loss), and J_{test} . A single 2d slice was chosen to contain the unperturbed minimum w_0 (white crosses), perturbed minimum (white stars), and parameters predicted by the local quadratic

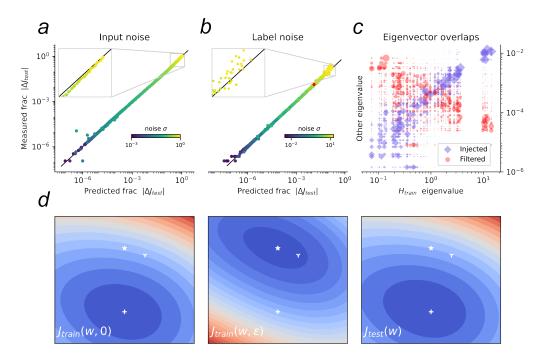


Figure 4: Validation of local fluctuation law in MLPs. Layer widths for both student and teacher were (5,5,5,1); nonlinearity: tanh; teacher network had gaussian weights with scale $4/\sqrt{d_{in}}$ for each layer; loss: MSE with ℓ_2 parameter $\lambda=1$. a,b) Predicted vs measured fluctuation-induced $\Delta J_{\rm test}/J_{\rm test}$ for increasing input (a) or label (b) noise amplitude σ . c) Eigenvector overlap function between $H_{\rm test}$ and the injected noise $E[zz^{\top}]$ (purple), and post-learning filtered noise $C_{\rm train}$ (red). Note how learning inflates/dampens variance along the low/high $H_{\rm test}$ eigenspaces. d) 2-dimensional loss landscapes for 1 example simulation: noiseless $J_{\rm train}$ (left), perturbed $J_{\rm train}$ (middle), and $J_{\rm test}$ (right). Cross: noiseless training minimum; star: minimum of perturbed $J_{\rm train}$ (ie. the new learned minimum); tri-star: parameters predicted by quadratic theory.

approximation (white "Y"s). Local geometry also predicts local gradient descent dynamics well (Appendix E.1; Fig. 7). Together, these results validate the predictions of the two-loss local theory.

3.4 CALCULATION OF OVERLAP FUNCTIONS FOR LARGE SCALE NETWORKS

Applying our theory to modern networks requires estimating the overlap function between the training and test operators. These operators have dimension equal to the number of parameters - often millions to billions - so any approach that forms them explicitly is infeasible.

Here we give a brief overview of our approach, deferring details to Appendix F. We apply two separate algorithms, one for computing overlaps among outlier eigenspaces and another for the remaining "bulk" spaces. Outlier overlaps are straightforward to obtain using subspace iteration, a generalization of power iteration wherein k vectors are repeatedly fed into the Hessian-vector product and orthonormalized (Appendix F.2; cf. Papyan (2019)). After sufficiently many iterations, the vectors give a good approximation to the top k eigenvectors of the matrix. Overlaps can then be computed directly. For the bulk eigenspaces, we generalize a well known approach to spectral density estimation known as the kernel polynomial method (KPM; Algorithm 1 in Appendix F.3).

Given self-adjoint matrices $A, B \in \mathbb{R}^{d \times d}$ and smoothing kernels $G(x; \sigma)$ of width σ , the smoothed total eigenvector overlap of A, B at eigenvalues λ_1, λ_2 can be written

$$\bar{\operatorname{tr}}\left[G_{A,\lambda_{1}}G_{B,\lambda_{2}}\right] = \frac{1}{d^{2}} \sum_{i,j=1}^{d} G\left(\lambda_{A,i} - \lambda_{1};\sigma\right) G\left(\lambda_{B,j} - \lambda_{2};\sigma\right) \left[d\left(v_{A,i} \cdot v_{B,j}\right)^{2}\right],$$
(10)

where $G_{A,\lambda_1} := G(A - \lambda_1 I; \sigma)$ and similarly for G_{B,λ_2} . To obtain the normalized overlap function treated above, one simply divides by the (smoothed) spectral densities of A, B at λ_1, λ_2 .

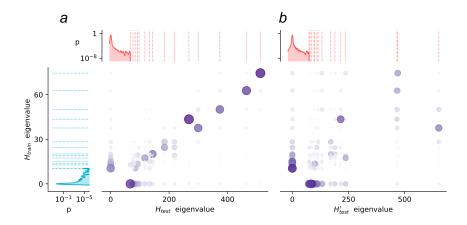


Figure 5: Overlap function for $H_{\rm train}$, $H_{\rm test}$, and class imbalanced Hessian $H'_{\rm test}$ for ResNet-20 trained on CIFAR10. Cyan/red data: spectra of $H_{\rm train}$, $H_{\rm test}$. Dashed lines indicate outlier eigenvalues. Purple scatters show overlap between each pair of eigenspaces/bulk spaces. Size and color reflect overlap magnitude. Left: $H_{\rm train}$, $H_{\rm test}$ overlaps. Note strong alignment indicated by large overlaps along the diagonal. Right: $H_{\rm train}$, $H'_{\rm test}$ overlaps. A large fraction of each Hessian's outlier energy is lost in low-outlier and bulk spaces of the other, indicating poor alignment.

Computing the trace in (10) is prohibitively expensive for large A, B, and so we resort to Hutchinson trace estimation, which approximates $\operatorname{tr}[X]$ with the average of $v^{\top}Xv$ for several samples of $v \sim \mathcal{N}(0, I)$. To ensure the trace is positive, instead of approximating (10), we use

$$\bar{\operatorname{tr}}\left[G_{A,\lambda_1}G_{B,\lambda_2}\right] = \bar{\operatorname{tr}}\left[G_{A,\lambda_1}^{1/2}G_{B,\lambda_2}G_{A,\lambda_1}^{1/2}\right] = \mathbb{E}_v \|G_{B,\lambda_2}^{1/2}G_{A,\lambda_1}^{1/2}v\|^2. \tag{11}$$

The KPM proceeds by approximating the kernel functions $G_{A,\lambda_1}^{1/2}$, $G_{B,\lambda_2}^{1/2}$ using truncated Chebyshev series. (Kernels and approximation degree must be chosen with care so that the truncated series doesn't allow "leakage" of extremely numerous near-0 eigenspaces; see Appendix F.3.) Thus (11) can be evaluated in terms of the vectors $T_i\left(B\right)T_j\left(A\right)v$, where T_k is the k^{th} Chebyshev polynomial. These vectors in turn can be generated efficiently via Chebyshev recurrences using only matrix-vector products.

Using these algorithms, we examined the effect of test-set class imbalance on train-test Hessian alignment. A CIFAR10-trained ResNet-20 was obtained from Chen (top-1 test set accuracy: 92.6%). 5000 train and 5000 test examples were randomly selected to define train and test Hessians. Their spectra, estimated using subspace iteration and the Lanczos algorithm, are shown in Fig. 5a (train in cyan; test in red). Non-outlier eigenspaces were grouped into a single bulk space for ease of visualization. Overlaps between train and test spaces obtained by subspace iteration are shown in the purple scatter, exhibiting strong alignment (bulk overlaps shown in Fig. 9). A class-imbalanced test Hessian H'_{test} was defined by selecting only test images with class labels 0, 1, or 2. This leads to significant misalignment with the leading eigenspaces of H_{train} (Fig. 5b), explaining the deleterious effects of class imbalance directly in terms of train-test loss geometry.

4 DISCUSSION

We show how, within a two-loss geometric framework, overlaps occupy a central role linking optimization geometry, random matrix theory, and practical machine learning phenomena. We derive novel theoretical tools for computing overlaps, illustrate through several examples how spectra set curvatures, while eigenvector overlaps route variance into error - unifying covariate shift and multiple descent - and develop scalable estimators for overlap analysis in large models. A natural next step is to track Hessian overlaps through training time (and across phases like plateaus and edge-of-stability) in large networks. Another direction ripe for exploration is overlap-aware regularizers and data/architecture interventions that encourage strong eigenvector alignment between train and test Hessians.

REFERENCES

- Ryan P. Adams, Jeffrey Pennington, Matthew J. Johnson, Jamie Smith, Yaniv Ovadia, Brian Patton, and James Saunderson. Estimating the spectral density of large implicit matrices, 2018. URL https://arxiv.org/abs/1802.03451.
- Florent Benaych-Georges and Raj Rao Nadakuditi. The eigenvalues and eigenvectors of finite, low rank perturbations of large random matrices. *Advances in Mathematics*, 227(1):494–521, 2011. ISSN 0001-8708. doi: https://doi.org/10.1016/j.aim.2011.02.007. URL https://www.sciencedirect.com/science/article/pii/S0001870811000570.
- Joël Bun, Jean-Philippe Bouchaud, and Marc Potters. Cleaning large correlation matrices: Tools from random matrix theory. *Physics Reports*, 666:1–109, January 2017. ISSN 0370-1573. doi: 10. 1016/j.physrep.2016.10.005. URL http://dx.doi.org/10.1016/j.physrep.2016.10.005.
- Lin Chen and Song Mei. Spectral multiplicity entails sample-wise multiple descent, 2022. URL https://openreview.net/forum?id=qaQ8kUBYhEK.
- Lin Chen, Yifei Min, Mikhail Belkin, and Amin Karbasi. Multiple descent: Design your own generalization curve, 2021. URL https://arxiv.org/abs/2008.01036.
- Yaofo Chen. Pytorch cifar models. https://github.com/chenyaofo/pytorch-cifar-models. Accessed: 2025-5-17.
- Jeremy Cohen, Simran Kaur, Yuanzhi Li, J Zico Kolter, and Ameet Talwalkar. Gradient descent on neural networks typically occurs at the edge of stability. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=jh-rTtvkGeM.
- Pierre Foret, Ariel Kleiner, Hossein Mobahi, and Behnam Neyshabur. Sharpness-aware minimization for efficiently improving generalization. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=6TmlmposlrM.
- Behrooz Ghorbani, Shankar Krishnan, and Ying Xiao. An investigation into neural net optimization via hessian eigenvalue density. In Kamalika Chaudhuri and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 2232–2241. PMLR, 09–15 Jun 2019. URL https://proceedings.mlr.press/v97/ghorbani19b.html.
- Gene H Golub and Gerard Meurant. *Matrices, moments and quadrature with applications*. Princeton Series in Applied Mathematics. Princeton University Press, Princeton, NJ, December 2009.
- Sobhan Hemati, Guojun Zhang, Amir Estiri, and Xi Chen. Understanding hessian alignment for domain generalization, 2023. URL https://arxiv.org/abs/2308.11778.
- Iain M. Johnstone. On the distribution of the largest eigenvalue in principal components analysis. *The Annals of Statistics*, 29(2), April 2001. ISSN 0090-5364. doi: 10.1214/aos/1009210544. URL http://dx.doi.org/10.1214/aos/1009210544.
- Ryo Karakida, Shotaro Akaho, and Shun-ichi Amari. Universal statistics of fisher information in deep neural networks: Mean field approach. In Kamalika Chaudhuri and Masashi Sugiyama (eds.), *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, volume 89 of *Proceedings of Machine Learning Research*, pp. 1032–1041. PMLR, 16–18 Apr 2019. URL https://proceedings.mlr.press/v89/karakida19a.html.
- Nitish Shirish Keskar, Dheevatsa Mudigere, Jorge Nocedal, Mikhail Smelyanskiy, and Ping Tak Peter Tang. On large-batch training for deep learning: Generalization gap and sharp minima. In *International Conference on Learning Representations*, 2017. URL https://openreview.net/forum?id=HloyRlygg.
- Itamar D. Landau, Gabriel C. Mel, and Surya Ganguli. Singular vectors of sums of rectangular random matrices and optimal estimation of high-rank signals: The extensive spike model. *Phys. Rev. E*, 108:054129, Nov 2023. doi: 10.1103/PhysRevE.108.054129. URL https://link.aps.org/doi/10.1103/PhysRevE.108.054129.

- Yue Li and Yuting Wei. Minimum ℓ₁-norm interpolators: Precise asymptotics and multiple descent,
 2021. URL https://arxiv.org/abs/2110.09502.
- Zhenyu Liao and Michael W. Mahoney. Hessian eigenspectra of more realistic nonlinear models.

 In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan (eds.), Advances in Neural Information Processing Systems, 2021. URL https://openreview.net/forum?id=o-RYNVOlxA8.
 - Lin Lin, Yousef Saad, and Chao Yang. Approximating spectral densities of large matrices. *SIAM Review*, 58(1):34–65, 2016. doi: 10.1137/130934283. URL https://doi.org/10.1137/130934283.
 - Gabriel Mel and Surya Ganguli. A theory of high dimensional regression with arbitrary correlations between input features and target functions: sample complexity, multiple descent curves and a hierarchy of phase transitions. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 7578–7587. PMLR, 18–24 Jul 2021. URL https://proceedings.mlr.press/v139/mel21a.html.
 - Gabriel Mel and Jeffrey Pennington. Anisotropic random feature regression in high dimensions. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=JfaWawZ8BmX.
 - Xuran Meng, Jianfeng Yao, and Yuan Cao. Multiple descent in the multiple random feature model, 2023. URL https://arxiv.org/abs/2208.09897.
 - James A. Mingo and Roland Speicher. Free Probability and Random Matrices. Springer New York, 2017. ISBN 9781493969425. doi: 10.1007/978-1-4939-6942-5. URL http://dx.doi. org/10.1007/978-1-4939-6942-5.
 - Boaz Nadler. Finite sample approximation results for principal component analysis: A matrix perturbation approach. *The Annals of Statistics*, 36(6), December 2008. ISSN 0090-5364. doi: 10.1214/08-aos618. URL http://dx.doi.org/10.1214/08-AOS618.
 - Vardan Papyan. The full spectrum of deepnet hessians at scale: Dynamics with sgd training and sample size, 2019. URL https://arxiv.org/abs/1811.07062.
 - Vardan Papyan. Traces of class/cross-class structure pervade deep learning spectra. *Journal of Machine Learning Research*, 21(252):1-64, 2020. URL http://jmlr.org/papers/v21/20-933.html.
 - Debashis Paul. Asymptotics of sample eigenstructure for a large dimensional spiked covariance model. *Statistica Sinica*, 17(4):1617–1642, 2007. ISSN 10170405, 19968507. URL http://www.jstor.org/stable/24307692.
 - Jeffrey Pennington and Yasaman Bahri. Geometry of neural network loss surfaces via random matrix theory. In Doina Precup and Yee Whye Teh (eds.), *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 2798–2806. PMLR, 06–11 Aug 2017. URL https://proceedings.mlr.press/v70/pennington17a.html.
 - Jeffrey Pennington and Pratik Worah. The spectrum of the fisher information matrix of a single-hidden-layer neural network. In *Neural Information Processing Systems*, 2018. URL https://api.semanticscholar.org/CorpusID:53410991.
 - Jeffrey Pennington and Pratik Worah. Nonlinear random matrix theory for deep learning. *Journal of Statistical Mechanics: Theory and Experiment*, 2019(12):124005, December 2019. ISSN 1742-5468. doi: 10.1088/1742-5468/ab3bc3. URL http://dx.doi.org/10.1088/1742-5468/ab3bc3.
- Marc Potters and Jean-Philippe Bouchaud. A First Course in Random Matrix Theory: for Physicists, Engineers and Data Scientists. Cambridge University Press, November 2020. ISBN 9781108488082. doi: 10.1017/9781108768900. URL http://dx.doi.org/10.1017/9781108768900.

Alexandre Rame, Corentin Dancette, and Matthieu Cord. Fishr: Invariant gradient variances for out-of-distribution generalization. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepes-vari, Gang Niu, and Sivan Sabato (eds.), *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 18347–18377. PMLR, 17–23 Jul 2022. URL https://proceedings.mlr.press/v162/rame22a.html.

Levent Sagun, Leon Bottou, and Yann LeCun. Eigenvalues of the hessian in deep learning: Singularity and beyond, 2017. URL https://openreview.net/forum?id=B186cP9gx.

Adepu Ravi Sankar, Yash Khasbage, Rahul Vigneswaran, and Vineeth N Balasubramanian. A deeper look at the hessian eigenspectrum of deep neural networks and its applications to regularization. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(11):9481–9488, May 2021. doi: 10.1609/aaai.v35i11.17142. URL https://ojs.aaai.org/index.php/AAAI/article/view/17142.

Sidak Pal Singh, Aurelien Lucchi, Thomas Hofmann, and Bernhard Schölkopf. Phenomenology of double descent in finite-width neural networks. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=lTqGXfn9Tv.

Nilesh Tripuraneni, Ben Adlam, and Jeffrey Pennington. Covariate shift in high-dimensional random feature regression, 2021. URL https://arxiv.org/abs/2111.08234.

Shashanka Ubaru, Yousef Saad, and Abd-Krim Seghouane. Fast estimation of approximate matrix ranks using spectral densities. *Neural Computation*, 29(5):1317–1351, 05 2017. ISSN 0899-7667. doi: 10.1162/NECO_a_00951. URL https://doi.org/10.1162/NECO_a_00951.

Zhewei Yao, Amir Gholami, Kurt Keutzer, and Michael W. Mahoney. Pyhessian: Neural networks through the lens of the hessian. 2020 IEEE International Conference on Big Data (Big Data), pp. 581–590, 2019. URL https://api.semanticscholar.org/CorpusID: 209376531.

A STATEMENT ON LARGE LANGUAGE MODEL USE

Large language models were used to polish writing and in conjunction with other tools to discover relevant published work.

B EIGENVECTOR OVERLAP FUNCTION

To connect the finite-d decomposition to random-matrix and free-probability tools, we now express eigenvector overlaps in a kernelized trace form amenable to free-probabilistic methods. Let X,Y be symmetric $d \times d$ matrices with eigendecompositions

$$X = \sum_{i=1}^{d} \lambda_i^X u_i u_i^\top, \qquad Y = \sum_{j=1}^{d} \lambda_j^Y v_j v_j^\top.$$

For bounded functions f, g,

$$\bar{\text{tr}}[f(X)g(Y)] = \frac{1}{d^2} \sum_{i=1}^{d} \sum_{j=1}^{d} f(\lambda_i^X) g(\lambda_j^Y) \left[d(u_i^\top v_j)^2 \right].$$
(12)

If f and g are sharply peaked around λ_1 and λ_2 , the sum concentrates on overlaps between eigenvectors with eigenvalues near (λ_1, λ_2) .

A convenient choice is the Poisson kernel

$$K(x; \mu, \sigma) := \frac{1}{\pi} \frac{\sigma}{(x - \mu)^2 + \sigma^2},$$

with center μ and width $\sigma > 0$. We define the overlap function

$$O(\lambda_1, \lambda_2) := \lim_{\sigma_1, \sigma_2 \to 0} \frac{\overline{\operatorname{tr}}[K(X; \lambda_1, \sigma_1) K(Y; \lambda_2, \sigma_2)]}{\overline{\operatorname{tr}}[K(X; \lambda_1, \sigma_1)] \overline{\operatorname{tr}}[K(Y; \lambda_2, \sigma_2)]}.$$
(13)

The denominator normalizes the total weight in the sum (12) to one, so $O(\lambda_1, \lambda_2)$ is the weighted average of the (scaled) squared overlaps $d(u_i^{\mathsf{T}} v_j)^2$ over eigenpairs near (λ_1, λ_2) . Now (12) can be rewritten

$$\bar{\operatorname{tr}}\left[f\left(X\right)g\left(Y\right)\right] = \iint f\left(\lambda_{1}\right)g\left(\lambda_{2}\right)O\left(\lambda_{1},\lambda_{2}\right)d\mu_{X}\left(\lambda_{1}\right)d\mu_{Y}\left(\lambda_{2}\right). \tag{14}$$

In fact, another way to define the overlap function is to write $\mu_{X,Y}$ for the measure taking $f,g\mapsto \bar{\operatorname{tr}}\left[f\left(X\right)g\left(Y\right)\right]$ and then defining $O\left(\lambda_{1},\lambda_{2}\right)$ to be the function making (14) hold, ie. $O=\frac{d\mu_{X,Y}}{d\mu_{X}\otimes\mu_{Y}}$.

B.1 Free transfer law

 Here we prove the following free transfer law for overlap functions:

Proposition 2. Let \tilde{B} combine B with a source of noise X that is free from A, B. Then

$$O_{A,\tilde{B}}(a,\tilde{b}) = \int O_{A,B}(a,b)O_{B,\tilde{B}}(b,\tilde{b}) \,\mu_B\left(db\right). \tag{15}$$

Note that despite its simple and appealing form, this relationship does not hold for general triples of matrices A,B,C - it suffices to check on finite dimensional matrices with simple spectra. While the coefficients of C's eigenvectors in the eigenbases of A,B do follow a change of basis law resembling formula (15), recall that the overlap function encodes the squared coefficients rather than the raw coefficients themselves. This relationship holds specifically because of the freeness relationship we've assumed.

To prove (15), start by noting that all of the spectral and overlap information for two matrices A, B is contained in the measure on \mathbb{R}^2 corresponding to the functional

$$\mu_{A,B}:f,g\mapsto \tau\left[f\left(A\right) g\left(B\right) \right] .$$

For example, setting $g \to 1$ and $f(x) = x^m$ gives access to all moments of A, and therefore to its spectrum, and similarly for B, while the overlap function is precisely $O_{A,B} = \frac{d\mu_{A,B}}{d\mu_A\otimes\mu_B}$, the Radon-Nikodym derivative of the joint measure with respect to the marginals.

For the remainder of this section, we will work in an abstract free probability space rather than with concrete matrices. Let (\mathcal{M}, τ) be a W^* -probability space (in our application this corresponds to the space of $d \times d$ matrices with $\tau = \mathbb{E}[\bar{\mathbf{r}}]$). See Mingo & Speicher (2017) for details.

Let $A, B, \tilde{B}, X \in \mathcal{M}$ be random variables and consider the problem of determining the overlap function $O_{A,\tilde{B}}$, where $\tilde{B} := F(B,X)$, where X is a source of noise that is free from A,B. As noted above, all of the overlap information for the three possible pairs of variables A,B,\tilde{B} is contained in the joint measures

$$\mu_{A,\tilde{B}},\mu_{A,B},\mu_{B,\tilde{B}},$$

supported on (some subset of) \mathbb{R}^2 . We are free now to treat these measures as ordinary probability measures of two scalar random variables. Denote by $\langle \cdot \rangle$ these scalar expectations. We now make use of the following two identities:

$$\langle f(x) g(y) \rangle_{(x,y) \sim \mu_{X,Y}} = \tau [f(X) g(Y)],$$
 (16)

and that for each g in a suitably broad class of functions (eg. at minimum all Poisson kernels), there is another function L_g depending linearly on g encoding the "expectation of $g\left(\tilde{B}\right)=g\left(F\left(B,X\right)\right)$ over X, conditioned on B", ie. such that

$$\tau \left[f(A,B) g(\tilde{B}) \right] = \tau \left[f(A,B) L_g(B) \right], \tag{17}$$

for all bounded borel f. This is Proposition 3, proved below using operator-valued free probability.

Combining these, we have $\left\langle f\left(a\right)g\left(\tilde{b}\right)\right\rangle _{\left(a,\tilde{b}\right)\sim\mu_{A,B}}=\tau\left[f\left(A\right)L_{g}\left(B\right)\right]$. Writing the right hand side as a scalar expectation and then conditioning on b,

$$\left\langle f\left(a\right)g\left(\tilde{b}\right)\right\rangle _{\left(a,\tilde{b}\right)\sim\mu_{A,B}} = \left\langle \left\langle f\left(a\right)\right\rangle _{a\sim\mu_{A\mid B=b}} L_{g}\left(b\right)\right\rangle _{b\sim\mu_{B}}$$
$$= \left\langle Q\left(b\right)L_{g}\left(b\right)\right\rangle _{b\sim\mu_{B}},$$

where $Q\left(b\right):=\left\langle f\left(a\right)\right\rangle _{a\sim\mu_{A|B=b}}$. Applying (16) and (17) again,

$$\begin{split} \left\langle f\left(a\right)g\left(\tilde{b}\right)\right\rangle_{\left(a,\tilde{b}\right)\sim\mu_{A,B}} &= \tau\left[Q\left(B\right)g\left(\tilde{B}\right)\right] \\ &= \left\langle Q\left(b\right)g\left(\tilde{b}\right)\right\rangle_{\left(b,\tilde{b}\right)\sim\mu_{B,\tilde{B}}} \\ &= \left\langle \left\langle f(a)\right\rangle_{a\sim\mu_{A|B=b}} \left\langle g(\tilde{b})\right\rangle_{\tilde{b}\sim\mu_{\tilde{B}|B=b}}\right\rangle_{bown}, \end{split}$$

which shows that a, \tilde{b} are independent conditioned on b:

$$\mu_{A,\tilde{B}} = \int \mu_{A|B=b} \mu_{\tilde{B}|B=b} d\mu_{B} (b).$$

Applying, for example, classical ϵ -gaussian smoothing to the measures of A, B, \tilde{B} , we can assume that $\mu_{X,Y} \ll \mu_X \otimes \mu_Y$ for any two of the three. Thus we are free to form the Radon-Nikodym derivative $\frac{d\mu_{X,Y}}{d\mu_X \otimes \mu_Y}$, which corresponds to the gaussian-smoothed overlap function $O_{X,Y;\epsilon}(x,y)$. Since $O_{X,Y;\epsilon}(x,y) \mu_X(dx) = \mu_{X|Y=y}(dx)$, we have that for any bounded measurable function

$$\begin{split} \int \phi \left(a, \tilde{b} \right) \, d\mu_{A,\tilde{B}} &= \iint \phi \left(a, \tilde{b} \right) \left(\int \mu_{A|B=b} \left(da \right) \mu_{\tilde{B}|B=b} \left(d\tilde{b} \right) \mu_{B} \left(db \right) \right) \\ &= \iint \phi \left(a, \tilde{b} \right) \int \left(O_{A,B;\epsilon} \left(a, b \right) \mu_{A} \left(da \right) \right) \left(O_{B,\tilde{B};\epsilon} \left(b, \tilde{b} \right) \mu_{\tilde{B}} \left(d\tilde{b} \right) \right) \mu_{B} \left(d\tilde{b} \right) \\ &= \iint \phi \left(a, \tilde{b} \right) \left(\int O_{A,B;\epsilon} \left(a, b \right) O_{B,\tilde{B};\epsilon} \left(b, \tilde{b} \right) \mu_{B} \left(db \right) \right) \mu_{A} \left(da \right) \mu_{\tilde{B}} \left(d\tilde{b} \right), \end{split}$$

so the last quantity is exactly the Radon-Nikodym density $O_{A\ \tilde{B}\cdot\epsilon}$:

$$O_{A,\tilde{B};\epsilon}\left(a,\tilde{b}\right) = \int O_{A,B;\epsilon}\left(a,b\right) O_{B,\tilde{B};\epsilon}\left(b,\tilde{b}\right) \mu_{B}\left(db\right).$$

Taking the smoothing to 0, one obtains the transfer law for the overlap function. Note that one may have to interpret the $O_{X,Y}$ as distributions (eg involving δ kernels) in case of degenerate overlap between two of the matrices (eg. A = B).

Proposition 3. Let A, B be free from X, and let $\tilde{B} = F(B, X)$ be a rational function of B, X. Then there is a linear operator L on functions such that for arbitrary bounded borel H,

$$\tau\left[H\left(A,B\right)g\left(\tilde{B}\right)\right]=\tau\left[H\left(A,B\right)L\left[g\right]\left(B\right)\right].$$

Proof. Let \mathbf{B} be a $k \times k$ linearization of \tilde{B} such that $\left[\mathbf{B}^{-1}\right]_{k,1} = g\left(F\left(B,X\right)\right)$. Decompose \mathbf{B} into its X-dependent and X-independent parts:

$$\mathbf{B} = X\mathbf{B}_X + \mathbf{B}_0.$$

Now form the linearization matrix

$$\mathbf{L}\left(M\right) := \left(\begin{array}{cc} -\mathbf{B} & 0\\ M\mathbf{e}_{k}^{\top} & -1 \end{array}\right),\,$$

so that

$$\mathbf{L}^{-1}(M) := \begin{pmatrix} -\mathbf{B}^{-1} & 0 \\ -M\mathbf{e}_k^{\mathsf{T}}\mathbf{B}^{-1} & -1 \end{pmatrix}.$$

In particular, $\mathbf{L}_{k,1}^{-1}\left(M\right)=-g\left(F\left(B,X\right)\right)$ and $\mathbf{L}_{k+1,1}^{-1}\left(M\right)=-Mg\left(F\left(B,X\right)\right)$. So now

$$[g_{\mathbf{L}(H(A,B))}(0)]_{k+1,1} = \tau \left[\left[(-\mathbf{L}(H(A,B)))^{-1} \right]_{k+1,1} \right]$$
$$= \tau \left[H(A,B) g(F(B,X)) \right].$$

As with **B**, decompose **L** into X-dependent and X-independent parts:

$$\mathbf{L}(M) = X\mathbf{L}_X + \mathbf{L}_0.$$

Assuming M is a function of A, B only, these two parts are operator-free (ie. with amalgamation over the space of complex-entried $(k+1) \times (k+1)$ matrices). By the additive subordination law,

$$g_{\mathbf{L}(M)}(0) = g_{\mathbf{L}_0}\left(-\mathcal{R}_{X\mathbf{L}_X}\left(g_{\mathbf{L}(M)}(0)\right)\right).$$

From the linearization, $g_{\mathbf{L}(M)}(0)$ is block lower triangular, and $X\mathbf{L}_X$ only has nonzero components in the upper left block. This implies that $\mathcal{R}_{X\mathbf{L}_X}\left(g_{\mathbf{L}(M)}(0)\right)$ is also confined to the upper left block, and that this entry is simply $\mathcal{R}_{-X\mathbf{B}_X}\left(\mathcal{E}\mathbf{B}^{-1}\right)$, where $\mathcal{E}:=\tau\otimes \mathrm{id}$ is the expectation functional of the operator-space. These imply

$$\begin{split} \tau\left[H\left(A,B\right)g\left(F\left(B,X\right)\right)\right] &= \left[g_{\mathbf{L}\left(H\left(A,B\right)\right)}\left(0\right)\right]_{k+1,k} \\ &= \left[g_{\mathbf{L}_{0}}\left(-\mathcal{R}_{X\mathbf{L}_{X}}\left(g_{\mathbf{L}\left(H\left(A,B\right)\right)}\left(0\right)\right)\right)\right]_{k+1,k} \\ &= \tau\left[\left(\begin{array}{cc} \mathbf{B}_{0} - \mathcal{R}_{-X\mathbf{B}_{X}}\left(\mathcal{E}\mathbf{B}^{-1}\right) & 0 \\ -H\left(A,B\right)\mathbf{e}_{k}^{\top} & 1 \end{array}\right)^{-1}\right]_{k+1,k} \\ &= \tau\left[\left(\begin{array}{cc} \mathbf{B}_{0} - \mathcal{R}_{-X\mathbf{B}_{X}}\left(\mathcal{E}\mathbf{B}^{-1}\right) & 0 \\ -H\left(A,B\right)\mathbf{e}_{k}^{\top} & 1 \end{array}\right)^{-1}\right]_{k+1,k} \\ &= \tau\left[H\left(A,B\right)\left(\mathbf{B}_{0} - \mathcal{R}_{-X\mathbf{B}_{X}}\left(\mathcal{E}\mathbf{B}^{-1}\right)\right)_{k,1}^{-1}\right]. \end{split}$$

The second factor can be regarded simply as a function of B (\mathbf{B}_0 is the non-X part of \mathbf{B} and $\mathcal{E}\mathbf{B}^{-1}: B \mapsto \mathcal{E}\mathbf{B}^{-1}(B,X) \in M_k(\mathbb{C})$. This proves that there is some operator L[g] as in the statement of the proposition. L must obviously be linear in g, completing the proof.

C TWO-LOSS GEOMETRY IN ANISOTROPIC RIDGE REGRESSION

We consider ridge regression on multivariate gaussian input data with train and test covariances $\Sigma_{\rm train}, \Sigma_{\rm test}$ and with linear ground truth, $y(x) = \frac{1}{\sqrt{d}} w_* \cdot x$. We will assume for simplicity that $w_* \sim \mathcal{N}(0, I_d)$. A noisy training set is generated by sampling inputs as follows. The training set consists of

$$x_i \sim \mathcal{N}\left(0, \Sigma_{\text{train}}\right), \quad y_i = y\left(x_i\right) + \xi_i, \quad \xi_i \sim \mathcal{N}\left(0, \sigma_{\xi}^2\right),$$

for $i=1,\ldots,m$. Using the framing of section 3.1.1, we will regard the noise ξ as perturbing a noiseless training objective. The train and test losses are formulated as follows:

$$J_{\text{train}}(w,\xi) := \frac{1}{2} \left[\frac{1}{m} \sum_{i=1}^{m} \left(y_i - \frac{1}{\sqrt{d}} w \cdot x_i \right)^2 \right] + \frac{\lambda}{2d} \|w\|^2$$
$$J_{\text{test}}(w) := \frac{1}{2} \left[\mathbb{E}_x \left(y(x) - \frac{1}{\sqrt{d}} w \cdot x \right)^2 \right].$$

Note 1) we keep J_{train} 's dependence on the perturbation ξ explicit, and 2) the noise is not included in the test loss (and when it is, after averaging, it changes the loss only by an additive constant). Finally, let us write $\hat{w}(\xi) := \operatorname{argmin}_w J_{\text{train}}(\hat{w}, \xi)$ for the learned weights, $\hat{y}(x) := \frac{1}{\sqrt{d}}\hat{w} \cdot x$ for the learned model, and $H_{\text{train}} := d \nabla^2 J_{\text{train}}$ and $H_{\text{test}} := d \nabla^2 J_{\text{test}}$ for the train and test Hessians; these scalings are chosen to keep spectra O(1).

High-dimensional ridge regression has been studied extensively, so rather than rederiving published formulas, as much as possible, we restrict attention to the novel focus of this article: overlap decompositions. We treat the label noise as a fluctuation of the training objective, and derive exact asymptotic formulas for the needed train-test spectra and overlap functions that describing the effect of the fluctuation on test error.

All formulas are obtainable from the general trace formula stated in the following two propositions, which we prove in Appendix D.

Proposition 4. The equation

$$r = \left(1 - q \int \frac{t}{z - tr} d\mu_{\Sigma_{\text{train}}}(t)\right)^{-1},\tag{18}$$

has a unique solution r satisfying $r \in \mathbb{H}^{\pm}$ for $z \in \mathbb{H}^{\mp}$ and satisfying 0 < r < 1 for z < 0. This defines a holomorphic function r(z) on all of $\mathbb{C}\backslash\mathbb{R}^{\geq 0}$ that is obtainable for each z by fixed point iteration of the right hand side of (18) from an arbitrary initial point r_0 satisfying $r_0 \in \mathbb{H}^{\pm}$ for $z \in \mathbb{H}^{\mp}$ and $0 < r_0 < 1$ for z < 0.

Proposition 5. Let

$$t_f(z) := \bar{tr} \left[f\left(\Sigma_{\text{test}}, \Sigma_{\text{train}} \right) \left(zI - A \right)^{-1} \right], \tag{19}$$

for bounded function f and complex scalar $z \in \mathbb{C} \backslash \mathbb{R}^{\geq 0}$. As $m, d \to \infty$ with q = d/m fixed,

$$t_f(z) \to t\bar{r} \left[f\left(\Sigma_{\text{test}}, \Sigma_{\text{train}}\right) \left(zI - r\left(z\right) \Sigma_{\text{train}}\right)^{-1} \right],$$

where r(z) is the solution of the self-consistent equation (18).

Thus to perform the calculation we simply express all quantities in terms of traces of the form $t_f(z)$, and then apply Propositions 4,5.

C.1 TRAIN-TEST HESSIAN OVERLAP FUNCTION

Since we will be interested primarily in the ridgeless limit $\lambda \to 0$, and since the effect of nonzero λ is simply to shift the spectrum of $H_{\text{train}} := \frac{1}{m} X^\top X + \lambda I$, we will omit λ in the computation of $O_{H_{\text{test}},H_{\text{train}}}(\lambda_1,\lambda_2)$, and will write $H_{\text{train}} = A = \frac{1}{m} X^\top X$ from now on unless explicitly stated otherwise. Thus, we are interested in the overlap function of the matrices

$$H_{\text{test}} = \Sigma_{\text{test}}, \qquad H_{\text{train}} = \Sigma_{\text{train}}^{1/2} \frac{1}{m} Z^{\top} Z \Sigma_{\text{train}}^{1/2}.$$

The computation is simplified significantly by taking advantage of asymptotic freeness of $\frac{1}{m}Z^{\top}Z$ from $\Sigma_{\text{train}}, \Sigma_{\text{test}}$. By Proposition 2, we have asymptotically

$$O_{H_{\text{test}},H_{\text{train}}}(\lambda_{\text{te}},\lambda_{\text{tr}}) = O_{\Sigma_{\text{test}},H_{\text{train}}}(\lambda_{\text{te}},\lambda_{\text{tr}})$$

$$= \int O_{\Sigma_{\text{test}},\Sigma_{\text{train}}}(\lambda_{\text{te}},\lambda) O_{\Sigma_{\text{train}},H_{\text{train}}}(\lambda,\lambda_{\text{tr}}) d\mu_{\Sigma_{\text{train}}}(\lambda). \quad (20)$$

In particular this shows, somewhat intuitively, that the overlap function of the train/test Hessians will itself depend on the overlap function of the population covariance matrices. Eq. (20) shows the dependence is quite simple: One simply composes the overlap kernels taking an $H_{\rm train} = \Sigma_{\rm train}^{1/2} \frac{1}{m} Z^{\top} Z \Sigma_{\rm train}^{1/2}$ eigenspace to a $\Sigma_{\rm train}$ one, and taking a $\Sigma_{\rm train}$ eigenspace to a $\Sigma_{\rm test}$ one. The overlap of the train and test population covariances, $O_{\Sigma_{\rm test},\Sigma_{\rm train}}$, is part of the input data of the problem and is therefore known. As for the second factor, $O_{\Sigma_{\rm train},H_{\rm train}}$, this is simply the overlap function of the population and sample covariance matrices for an anisotropic gaussian sample. Formulas for this quantity are known (see, eg. Potters & Bouchaud (2020)). To keep the presentation self-contained we quickly derive an expression using operator-valued free probability.

Following Appendix B, the train-test Hessian overlap function can be computed via

$$O_{\Sigma_{\mathrm{train}},H_{\mathrm{train}}}\left(\lambda_{1},\lambda_{2}\right):=\lim_{\sigma\rightarrow0}\frac{\bar{\mathrm{tr}}\big[K(\Sigma_{\mathrm{train}};\lambda_{1},\sigma_{1})\,K(H_{\mathrm{train}};\lambda_{2},\sigma_{2})\big]}{\bar{\mathrm{tr}}\big[K(\Sigma_{\mathrm{train}};\lambda_{1},\sigma_{1})\big]\,\bar{\mathrm{tr}}\big[K(H_{\mathrm{train}};\lambda_{2},\sigma_{2})\big]},$$

where

$$K(x; \mu, \sigma) := \frac{1}{\pi} \frac{\sigma}{(x - \mu)^2 + \sigma^2},$$

is the Poisson kernel with center μ and width σ . Moving the first trace into the numerator trace and cancelling a factor of π from top and bottom, we find that computing $O_{H_{\text{test}},H_{\text{train}}}(\lambda_1,\lambda_2)$ requires the numerator and denominator traces

$$\bar{\operatorname{tr}}\left[h\left(\Sigma_{\mathrm{train}}\right)\frac{\sigma_{2}}{\left(A-\lambda_{2}I\right)^{2}+\sigma_{2}^{2}I}\right],\qquad \bar{\operatorname{tr}}\left[\frac{\sigma_{2}}{\left(A-\lambda_{2}I\right)^{2}+\sigma_{2}^{2}I}\right],$$

where $h(\Sigma_{\text{train}}) := K(\Sigma_{\text{train}}; \lambda_1, \sigma_1) / \bar{\text{tr}}[K(\Sigma_{\text{train}}; \lambda_1, \sigma_1)].$

Using the definition of $t_f(z)$, and the resolvent form of the Poisson kernel,

$$K(x; \mu, \sigma) = -\frac{1}{\pi} \operatorname{Im} (\mu + i\sigma - x)^{-1},$$

these become

$$-\operatorname{Im} t_h (\lambda_2 + i\sigma_2), \quad -\operatorname{Im} t_1 (\lambda_2 + i\sigma_2).$$

Proposition 5 implies

$$t_h(z) \to \bar{\text{tr}} \left[h\left(\Sigma_{\text{train}}\right) \left(zI - r\left(z\right) \Sigma_{\text{train}}\right)^{-1} \right]$$
 (21)

$$t_1(z) \to \bar{\text{tr}} \left[\left(zI - r(z) \Sigma_{\text{train}} \right)^{-1} \right],$$
 (22)

SO

$$O_{\Sigma_{\text{train}},H_{\text{train}}}(\lambda_{1},\lambda_{2}) \to \lim_{\sigma_{1}\to 0} \lim_{z\to\lambda_{2}^{+i}} \frac{\operatorname{Im} \bar{\text{tr}}\left[h\left(\Sigma_{\text{train}}\right)\left(zI-r\left(z\right)\Sigma_{\text{train}}\right)^{-1}\right]}{\operatorname{Im} \bar{\text{tr}}\left[\left(zI-r\left(z\right)\Sigma_{\text{train}}\right)^{-1}\right]},$$

where $\lim_{z\to\lambda_2^{+i}}$ is shorthand for $\lim_{\sigma\to 0}$ with $z=\lambda_2+i\sigma$. Taking $\sigma_1\to 0$ sends $h\left(\Sigma_{\rm train}\right)$ to a delta function and collapses the trace in the numerator to the λ_1 eigenspace of $\Sigma_{\rm train}$, so

$$O_{\Sigma_{\text{train}}, H_{\text{train}}}(\lambda_1, \lambda_2) \to \lim_{z \to \lambda_2^{+i}} \frac{\text{Im} \frac{1}{z - r(z)\lambda_1}}{\int \text{Im} \frac{1}{z - r(z)\lambda} d\mu_{\Sigma_{\text{train}}}(\lambda)}.$$

Composing with $O_{\Sigma_{\text{test}},\Sigma_{\text{train}}}$ yields the overlap function $O_{\Sigma_{\text{test}},H_{\text{train}}} = O_{H_{\text{test}},H_{\text{train}}}$.

C.2 Overlap decomposition of ΔJ

Trace integrals are written in terms of the spectra and overlaps of the matrices involved. To explicitly determine the spectral density of H_{train} , note that it can be written in terms of the trace in equation (22),

$$\rho_{H_{\text{train}}}(\lambda_{\text{tr}}) = \lim_{\sigma \to 0} \bar{\text{tr}} [K(H_{\text{train}}; \lambda_{\text{tr}}, \sigma)].$$

Using the same approach as above, we have the following for the σ -Poisson-smoothed spectral density of H_{train} :

$$\rho_{H_{\text{train}};\sigma}(\lambda_{\text{tr}}) = -\frac{1}{\pi} \text{Im} \, t_1 \left(\lambda_{\text{tr}} + i\sigma \right)$$

$$\rightarrow -\frac{1}{\pi} \int \text{Im} \frac{1}{z - r(z) \lambda} d\mu_{\Sigma_{\text{train}}}(\lambda) .$$

Collecting the results of the previous section and the fluctuation formula (7),

$$\Delta J_{\text{test}} = \frac{\sigma^2 q}{2} \iiint \underbrace{\lambda_{\text{te}}}_{\text{test curvature}} \underbrace{\frac{\lambda_{\text{tr}}}{(\lambda_{\text{tr}} + \lambda)^2}}_{\text{train variance}} \underbrace{O_{H_{\text{train}}, H_{\text{test}}}(\lambda_{\text{te}}, \lambda_{\text{tr}})}_{\text{eigenspace overlap}} \mu_{H_{\text{test}}}(d\lambda_{\text{te}}) \mu_{H_{\text{train}}}(d\lambda_{\text{tr}}),$$
(23)

where

$$O_{H_{\text{test}},H_{\text{train}}}(\lambda_{\text{te}},\lambda_{\text{tr}}) = \int O_{\Sigma_{\text{test}},\Sigma_{\text{train}}}(\lambda_{\text{te}},\lambda) O_{\Sigma_{\text{train}},H_{\text{train}}}(\lambda,\lambda_{\text{tr}}) d\mu_{\Sigma_{\text{train}}}(\lambda), \qquad (24)$$

with

$$O_{\Sigma_{\text{train}}, H_{\text{train}}}(\lambda_1, \lambda_2) \to \lim_{z \to \lambda_2^{+i}} \frac{\operatorname{Im} \frac{1}{z - r(z)\lambda_1}}{\int \operatorname{Im} \frac{1}{z - r(z)\lambda} d\mu_{\Sigma_{\text{train}}}(\lambda)}$$

This provides a complete decomposition of the test loss fluctuation in terms of spectra and overlaps of the train and test Hessian.

C.3 COVARIATE SHIFT

Formulas (23) and (24) show the effect of covariate shift in train/test sets is naturally expressed in terms of the overlap function $O_{\Sigma_{\text{test}},\Sigma_{\text{train}}}(x,y)$ of the two population covariances. (Note that there are two levels of overlap decomposition - that of the test loss fluctuation (23), depending on the train-test Hessian overlap function, and that of the train-test Hessian overlap function itself, which unsurprisingly depends on the overlap of $\Sigma_{\text{test}},\Sigma_{\text{train}}$.)

We can equivalently start from explicit expressions for the fluctuation. We show in Appendix C.4 that

$$\Delta J = -\frac{1}{2} q \sigma_{\xi}^2 \frac{d}{d\lambda} \lambda t_{\text{id}} \left(-\lambda \right).$$

Adopting the notation $\tilde{\lambda} := \frac{\lambda}{r(-\lambda)}$,

$$\lambda t_{\rm id} \left(-\lambda \right) \to -\tilde{\lambda} \bar{\rm tr} \left[\Sigma_{\rm test} \left(\tilde{\lambda} I + \Sigma_{\rm train} \right)^{-1} \right]$$
$$\frac{d}{d\lambda} \lambda t_{\rm id} \left(-\lambda \right) \to -\tilde{\lambda}' \bar{\rm tr} \left[\Sigma_{\rm test} \Sigma_{\rm train} \left(\tilde{\lambda} I + \Sigma_{\rm train} \right)^{-2} \right],$$

meaning

$$\Delta J \rightarrow \frac{1}{2} q \sigma_{\xi}^2 \tilde{\lambda}' \bar{\mathrm{tr}} \left[\Sigma_{\mathrm{test}} \Sigma_{\mathrm{train}} \left(\tilde{\lambda} I + \Sigma_{\mathrm{train}} \right)^{-2} \right],$$

so

$$\Delta J \rightarrow \frac{1}{2} q \sigma_{\xi}^2 \tilde{\lambda}' \int \lambda_{\rm te} \frac{\lambda_{\rm tr}}{\left(\tilde{\lambda} + \lambda_{\rm tr}\right)^2} O_{\Sigma_{\rm test}, \Sigma_{\rm train}} \left(\lambda_{\rm te}, \lambda_{\rm tr}\right) d\mu_{\Sigma_{\rm test}} \left(\lambda_{\rm te}\right) d\mu_{\Sigma_{\rm train}} \left(\lambda_{\rm tr}\right),$$

which parallels (23) but averages out the random training inputs and label noise to express ΔJ purely in terms of the known objects $\Sigma_{\rm train}, \Sigma_{\rm test}$. This expression shows that label noise leads to large increases in test loss when a direction of large training variance (small eigenvalue $\lambda_{\rm tr}$ of $\Sigma_{\rm train}$) and a direction of large test curvature (large eigenvalue $\lambda_{\rm te}$ of $\Sigma_{\rm test}$) experience significant overlap (large $O_{\Sigma_{\rm test},\Sigma_{\rm train}}$ ($\lambda_{\rm te},\lambda_{\rm tr}$)).

C.4 EXPLICIT FORMULAS FOR TEST LOSS, FLUCTUATION

Here we derive explicit expressions for the full test loss and test loss fluctuation under general covariate shift. Since these formulas and generalizations of them are already published, this section is mostly for internal reference - especially for calculation of theoretical loss curves in Figs. 1 and 2.

Let X have rows x_i^{\top} and ξ have components ξ_i . J_{train} can be written

$$J_{\text{train}}(w,\xi) := \frac{1}{2m} \| \frac{1}{\sqrt{d}} X w_* + \xi - \frac{1}{\sqrt{d}} X w \|^2 + \frac{\lambda}{2d} w^\top w.$$

Differentiating, we find

$$\nabla J_{\text{train}}(w,\xi) = \frac{1}{d} H_{\text{train}} w - \frac{1}{\sqrt{d}} \frac{1}{m} X^{\top} \left(\frac{1}{\sqrt{d}} X w_* + \xi \right)$$
$$H_{\text{train}} := d \nabla^2 J_{\text{train}} = \frac{1}{m} X^{\top} X + \lambda I.$$

Similarly,

$$J_{\text{test}}(w) = \frac{1}{2d} (w - w_*) \Sigma_{\text{test}} (w - w_*).$$

$$H_{\text{test}} := d \nabla^2 J_{\text{test}} = \Sigma_{\text{test}}.$$

Solving $0 = \nabla_w J_{\text{train}}$ yields

$$\hat{w} = H_{\text{train}}^{-1} \left(\frac{1}{m} X^{\top} X \right) w_* + \left(H_{\text{train}} \right)^{-1} \frac{\sqrt{d}}{m} X^{\top} \xi.$$

Substituting into J_{test} yields

$$J_{\mathrm{test}}\left(\hat{w}\left(\xi\right)\right) = \frac{1}{2}\bar{\mathrm{tr}}\left[\Sigma_{\mathrm{test}}\frac{q\sigma_{\xi}^{2}\frac{1}{m}X^{\top}X + \lambda^{2}I}{\left(\frac{1}{m}X^{\top}X + \lambda I\right)^{2}}\right].$$

Since $(A + \lambda I)^{-2} = -\frac{d}{d\lambda} (A + \lambda I)^{-1}$, we can write

$$J_{\text{test}}\left(\hat{w}\left(\xi\right)\right) = -\frac{1}{2} \left(q\sigma_{\xi}^{2} \frac{d}{d\lambda} \lambda - \lambda^{2} \frac{d}{d\lambda}\right) t_{\text{id}}\left(-\lambda\right) \tag{25}$$

$$= -\frac{1}{2} \left(q \sigma_{\xi}^{2} \left(t_{id} \left(-\lambda \right) - \lambda t_{id}^{\prime} \left(-\lambda \right) \right) + \lambda^{2} t_{id}^{\prime} \left(-\lambda \right) \right) \tag{26}$$

Proposition 5 implies

$$\begin{split} t_{\mathrm{id}}\left(z\right) &\to \bar{\mathrm{tr}}\left[\Sigma_{\mathrm{test}}\left(zI - r\left(z\right)\Sigma_{\mathrm{train}}\right)^{-1}\right] \\ t_{\mathrm{id}}'\left(z\right) &\to -\bar{\mathrm{tr}}\left[\Sigma_{\mathrm{test}}\left(zI - r\left(z\right)\Sigma_{\mathrm{train}}\right)^{-2}\left(I - r'\left(z\right)\Sigma_{\mathrm{train}}\right)\right], \end{split}$$

which fully specifies J_{test} $(\hat{w}(\xi))$. The fluctuation is easily gotten by setting $\sigma_{\xi}^2 \to 0$ and subtracting from J_{test} $(\hat{w}(\xi))$.

Reduction to published formulas Letting $\tilde{\lambda} := \frac{\lambda}{r(-\lambda)}$ and substituting into (18), we obtain

$$\lambda = \tilde{\lambda} - \frac{1}{\alpha} \int \frac{\tilde{\lambda}t}{\tilde{\lambda} + t} d\mu_{\Sigma_{\text{train}}}(t), \qquad (27)$$

which is eq. (8) of Mel & Ganguli (2021) for the "effective regularization".

The fluctuation in (25) is

$$\Delta J = -q\sigma_{\xi}^{2} \frac{1}{2} \frac{d}{d\lambda} \lambda t_{id} \left(-\lambda \right).$$

Since

$$\begin{split} \lambda t_{\mathrm{id}} \left(- \lambda \right) &\to \frac{\lambda}{r \left(- \lambda \right)} \bar{\mathrm{tr}} \left[\Sigma_{\mathrm{train}} \left(- \frac{\lambda}{r \left(- \lambda \right)} I - \Sigma_{\mathrm{train}} \right)^{-1} \right] \\ &= - \tilde{\lambda} \bar{\mathrm{tr}} \left[\Sigma_{\mathrm{train}} \left(\tilde{\lambda} I + \Sigma_{\mathrm{train}} \right)^{-1} \right], \end{split}$$

we get

$$\Delta J = q \sigma_\xi^2 \frac{1}{2} \tilde{\lambda}' \bar{\text{tr}} \left[\left(\frac{\Sigma_{\text{train}}}{\tilde{\lambda} I + \Sigma_{\text{train}}} \right)^2 \right].$$

The authors define $\frac{1}{\rho_f} := \frac{d\tilde{\lambda}}{d\lambda}$, so

$$\Delta J = \frac{1}{2} q \sigma_\xi^2 \frac{1}{\rho_f} \bar{\mathrm{tr}} \left[\left(\frac{\Sigma_{\mathrm{train}}}{\tilde{\lambda} I + \Sigma_{\mathrm{train}}} \right)^2 \right],$$

which matches the fluctuation term of their formula up to constant factors differing in the loss definitions. Next, the remaining term can be written

$$B = \frac{1}{2} \left(-\lambda t_{id} \left(-\lambda \right) + \lambda \frac{d}{d\lambda} \lambda t_{id} \left(-\lambda \right) \right).$$

Using (25) again,

$$B = \frac{1}{2} \left(\bar{\operatorname{tr}} \left[\frac{\tilde{\lambda} \Sigma_{\text{train}}}{\tilde{\lambda} I + \Sigma_{\text{train}}} \right] - \lambda \tilde{\lambda}' \bar{\operatorname{tr}} \left[\frac{\Sigma_{\text{train}}^2}{\left(\tilde{\lambda} I + \Sigma_{\text{train}} \right)^2} \right] \right)$$

Comparing to (27), the first term is $\alpha(\tilde{\lambda} - \lambda)$, and differentiating gives

$$\tilde{\lambda}' = \frac{\alpha}{\alpha - \text{tr}\left[\left(\frac{\Sigma_{\text{train}}}{\tilde{\lambda} + \Sigma_{\text{train}}}\right)^2\right]}.$$

Substituting and simplifying yields

$$\begin{split} B &= \frac{1}{2} \left(\alpha \tilde{\lambda} - \alpha \lambda \frac{\alpha}{\alpha - \text{tr} \left[\left(\frac{\Sigma_{\text{train}}}{\tilde{\lambda} + \Sigma_{\text{train}}} \right)^2 \right]} \right) \\ &= \frac{1}{2} \tilde{\lambda}' \left(\tilde{\lambda} \left(\alpha - \text{tr} \left[\left(\frac{\Sigma_{\text{train}}}{\tilde{\lambda} + \Sigma_{\text{train}}} \right)^2 \right] \right) - \alpha \lambda \right). \end{split}$$

Once again using equation (27), $\alpha\left(\tilde{\lambda}-\lambda\right)$ can be turned back into a trace:

$$\begin{split} B &= \frac{1}{2} \tilde{\lambda}' \left(\text{tr} \left[\frac{\tilde{\lambda} \Sigma_{\text{train}}}{\tilde{\lambda} I + \Sigma_{\text{train}}} \right] - \tilde{\lambda} \text{tr} \left[\left(\frac{\Sigma_{\text{train}}}{\tilde{\lambda} + \Sigma_{\text{train}}} \right)^2 \right] \right) \\ &= \frac{1}{2} \tilde{\lambda}' \text{tr} \left[\frac{\tilde{\lambda}^2 \Sigma_{\text{train}}}{\left(\tilde{\lambda} I + \Sigma_{\text{train}} \right)^2} \right], \end{split}$$

which is equivalent to their second term.

C.5 k-LEVEL MODEL

At several points in the main text we refer to a k-level input covariance,

$$\mu_{\Sigma_{\text{train}}} = \sum_{i=1}^{k} p_i \delta_{s_i}.$$

In this case the self-consistent equation for r (18) becomes

$$r = \left(1 - q \sum_{i=1}^{k} p_i \frac{s_i}{z - s_i r}\right)^{-1},$$

which can be written as p(r, z) = 0 for some polynomial in r, z. Similarly, the overlap function simplifies to a sum over the distinct eigenvalues of Σ_{train} :

$$O_{\Sigma_{\text{train}}, H_{\text{train}}}(\lambda_1, \lambda_2) \to \lim_{z \to \lambda_2^{+i}} \frac{\text{Im} \frac{1}{z - r(z)\lambda_1}}{\sum_{i=1}^k p_i \text{Im} \frac{1}{z - r(z)s_i}}.$$

C.5.1 SEPARATED SCALES LIMIT

We now assume the scales are widely separated: $s_{i+1} \ll s_i$. We will also work with the ridgeless formulas corresponding to $\lambda \to 0$ derived in Appendix C.6. For simplicity, assume $s_1 = 1$. We will obtain leading order formulas for $h(\alpha)$ as the ratio of successive scales is taken to 0. In Appendix C.6, h is defined and found to satisfy the following self-consistent equation (equation (30)):

$$1 = \int \frac{t}{h + \alpha t} d\mu_{\Sigma_{\text{train}}}(t)$$

The right hand side is a decreasing function of h and a decreasing function of α , we have that h is a decreasing function of α . Since the integral reduces to a sum over the k eigenvalues, and since all

terms with $t \ll h$ do not contribute at leading order in s_{i+1}/s_i , we assume that $h \approx s_1 = 1$ and neglect all lower terms, giving

 $1 = p_1 \frac{1}{h + \alpha},$

so $h = p_1 - \alpha$ and ΔJ is

$$\Delta J = \frac{\sigma_{\xi}^2}{2} \frac{\alpha \int sO\left(s,1\right) d\mu_{\Sigma_{\text{test}}}\left(s\right)}{p_1 - \alpha},$$

the integral in the numerator can be written

$$\int sO\left(s,1\right)d\mu_{\Sigma_{\mathrm{test}}}\left(s\right) = \frac{1}{p_{1}}\bar{\mathrm{tr}}\left[\Sigma_{\mathrm{test}}P_{\Sigma_{\mathrm{train}}=1}\right],$$

where $P_{\Sigma_{\text{train}}=a}$ is the projector onto Σ_{train} 's a-eigenspace. In other words, the integral is simply the normalized total overlap of Σ_{test} onto the strong training covariance space, and is equal to 1 for $\Sigma_{\text{train}}=\Sigma_{\text{test}}$.

Now let us assume that h is near the scale s_i^2 . The self-consistent equation becomes

$$1 = p_i \frac{s_i^2}{h + \alpha s_i^2} + \frac{1}{\alpha} \mu_{\Sigma_{\text{train}}} \left(\gg s_i^2 \right),$$

where $\mu_{\Sigma_{\text{train}}}$ ($\gg s_i^2$) is the total probability mass of all scales greater than s_i^2 , ie $\sum_{j=1}^{i-1} p_j$. Solving yields

$$\alpha s_i^2 \left[\frac{p_i}{\alpha - \mu_{\Sigma_{\text{train}}}} \left(\gg s_i^2 \right) - 1 \right] = h,$$

which is consistent with the assumption that $h \sim s_i^2$. Since $h \geq 0$, we only get a valid solution for $\alpha \geq \mu_{\Sigma_{\text{train}}} (\gg s_i^2)$. Substituting back into the error expression yields

$$\Delta J = \frac{\sigma_{\xi}^{2}}{2} \frac{p_{i} \alpha^{2} \int \frac{st}{(h+\alpha t)^{2}} O\left(s,t\right) d\mu_{\Sigma_{\text{train}}}\left(t\right) d\mu_{\Sigma_{\text{test}}}\left(s\right)}{\left(\mu_{\Sigma_{\text{train}}}\left(\gg s_{i+1}^{2}\right) - \alpha\right) \left(\alpha - \mu_{\Sigma_{\text{train}}}\left(\gg s_{i}^{2}\right)\right)}.$$

Evaluating the numerator generally requires a choice of $\Sigma_{\rm test}$'s behavior in the limit $s_{i+1}/s_i \to 0$, but note that the denominator has zeros at $\alpha = \mu_{\Sigma_{\rm train}} \left(\gg s_i^2 \right)$, $\mu_{\Sigma_{\rm train}} \left(\gg s_{i+1}^2 \right)$, and so the error will generically become infinite whenever α is equal to the cumulative mass of some number of top scales. As a simple special case, letting $\Sigma_{\rm test} = \Sigma_{\rm train}$, this reduces to

$$\Delta J = \frac{\sigma_{\xi}^2}{2} \frac{\left(\alpha - \mu_{\Sigma_{\text{train}}} \left(\gg s_i^2 \right) \right)^2 + p_i \mu_{\Sigma_{\text{train}}} \left(\gg s_i^2 \right)}{\left(\mu_{\Sigma_{\text{train}}} \left(\gg s_{i+1}^2 \right) - \alpha \right) \left(\alpha - \mu_{\Sigma_{\text{train}}} \left(\gg s_i^2 \right) \right)}.$$

Since under this assumption,

$$B = \frac{1}{2} h \, \bar{\text{tr}} \left[\Sigma_{\text{train}} \left(hI + \alpha \Sigma_{\text{train}} \right)^{-1} \right] = \frac{1}{2} h$$

$$= \frac{1}{2} \left\{ \alpha s_i^2 \left[\frac{p_i}{\alpha - \mu_{\Sigma_{\text{train}}} \left(\gg s_i^2 \right)} - 1 \right] \quad \mu_{\Sigma_{\text{train}}} \left(\gg s_i^2 \right) < \alpha \le \mu_{\Sigma_{\text{train}}} \left(\gg s_{i+1}^2 \right),$$

when we take $s_{i+1}/s_i \to 0$, B only contributes at the highest scale, so

$$B = \frac{1}{2}\sigma_+ \left(p_1 - \alpha\right),\,$$

where σ_+ is the relu function.

C.6 RIDGELESS LIMIT

Here we simplify our formula for the test error in the ridgeless limit. From (25), we have

$$\begin{split} \Delta J &= -\frac{1}{2}q\sigma_{\xi}^2\frac{d}{d\lambda}\lambda t_{\mathrm{id}}\left(-\lambda\right)\\ B &:= J - \Delta J = -\frac{1}{2}\lambda^2 t_{\mathrm{id}}'\left(-\lambda\right). \end{split}$$

It will also be helpful to consult t_{id} and r's explicit expressions as matrix traces (equations (19) and (34)):

$$t_{\rm id}\left(z\right) = \bar{\rm tr}\left[\Sigma_{\rm test}\left(zI - \frac{1}{m}X^{\top}X\right)^{-1}\right],$$
 (28)

$$r(z) = 1 + q \, \bar{\text{tr}} \left[(z - A)^{-1} A \right].$$
 (29)

C.6.1 OVERSAMPLED REGIME

From (28), and since for $\alpha := m/d > 1$, the limiting spectrum of $\frac{1}{m}X^{\top}X$ is bounded away from 0, $t_{\rm id}$ ($-\lambda$) is analytic as $\lambda \to 0^+$. Thus in the oversampled regime

$$\Delta J
ightarrow - rac{1}{2} q \sigma_{\xi}^2 t_{\mathrm{id}} \left(0
ight), \qquad B
ightarrow 0.$$

From Propositions 4 and 5,

$$t_{\rm id}\left(z\right) \to -\frac{1}{r\left(0\right)} \bar{\rm tr}\left[\Sigma_{\rm test} \Sigma_{\rm train}^{-1}\right], \qquad r\left(0\right) = 1 - q,$$

$$\Delta J, J_{\mathrm{test}} \to \frac{1}{2} \sigma_{\xi}^2 \frac{1}{\alpha - 1} \bar{\mathrm{tr}} \left[\Sigma_{\mathrm{test}} \Sigma_{\mathrm{train}}^{-1} \right].$$

C.6.2 UNDERSAMPLED REGIME

Now assume $\alpha < 1$. Now for $\lambda \to 0$, $t_{\rm id}$ and r's explicit expressions in (28) and (34) suggest $t_{\rm id}(-\lambda) = O\left(\lambda^{-1}\right)$ and $r\left(-\lambda\right) = O\left(\lambda\right)$. For convenience we will rewrite our formulas in terms of $h\left(z\right) := \frac{1}{q}\left(\frac{z}{r(-z)} - z\right)$. Substituting into the self-consistent equation for r (18) and simplifying gives

$$h(z) = (qh(z) + z) \int \frac{t}{qh(z) + z + t} d\mu_{\Sigma_{\text{train}}}(t).$$

Now differentiating and setting $z \to 0$, we find

$$1 = q \int \frac{t}{qh + t} d\mu_{\Sigma_{\text{train}}}(t)$$
 (30)

$$h' = \frac{\int \left(\frac{t}{qh+t}\right)^2 d\mu_{\Sigma_{\text{train}}}(t)}{1 - q \int \left(\frac{t}{qh+t}\right)^2 d\mu_{\Sigma_{\text{train}}}(t)},$$
(31)

where we've suppressed the argument of h, h'.

We now write the error expressions in terms of these

$$\begin{split} \Delta J &= -\frac{1}{2} q \sigma_{\xi}^2 \frac{d}{d\lambda} \lambda t_{\mathrm{id}} \left(-\lambda \right) \\ &= \frac{1}{2} q \sigma_{\xi}^2 \frac{d}{d\lambda} \left(q h \left(\lambda \right) + \lambda \right) \bar{\mathrm{tr}} \left[\Sigma_{\mathrm{test}} \left(\left(q h \left(\lambda \right) + \lambda \right) I + \Sigma_{\mathrm{train}} \right)^{-1} \right] \\ &\xrightarrow{\lambda \to 0} \frac{1}{2} q \sigma_{\xi}^2 \frac{\bar{\mathrm{tr}} \left[\Sigma_{\mathrm{test}} \frac{\Sigma_{\mathrm{train}}}{\left(q h I + \Sigma_{\mathrm{train}} \right)^2} \right]}{1 - q \int \left(\frac{t}{q h + t} \right)^2 d\mu_{\Sigma_{\mathrm{train}}} \left(t \right)} \\ &= \frac{\sigma_{\xi}^2}{2} \frac{q \int \frac{st}{\left(q h + t \right)^2} O \left(s, t \right) d\mu_{\Sigma_{\mathrm{train}}} \left(t \right) d\mu_{\Sigma_{\mathrm{test}}} \left(s \right)}{1 - q \int \left(\frac{t}{q h + t} \right)^2 d\mu_{\Sigma_{\mathrm{train}}} \left(t \right)}, \end{split}$$

while

$$\begin{split} B &= -\frac{1}{2} \lambda^2 t'_{\text{id}} \left(-\lambda \right) \\ &= \frac{1}{2} q h \, \bar{\text{tr}} \left[\Sigma_{\text{test}} \left(q h I + \Sigma_{\text{train}} \right)^{-1} \right] \\ &= \frac{1}{2} q h \int \frac{s}{q h + t} O\left(s, t \right) d \mu_{\Sigma_{\text{train}}} \left(t \right) d \mu_{\Sigma_{\text{test}}} \left(s \right). \end{split}$$

Finally, the total loss is just $B + \Delta J$.

D CHARACTERIZATION OF $t_f(z)$

Here we derive an asymptotically exact expression for

$$t_{F}\left(z\right):=\bar{\operatorname{tr}}\left[F\left(\Sigma_{\mathrm{test}},\Sigma_{\mathrm{train}}\right)\left(zI-A\right)^{-1}\right].$$

Let us abbreviate $F := F(\Sigma_{\text{train}}, \Sigma_{\text{test}})$. First,

$$A = \frac{1}{m} X^{\top} X = \Sigma_{\text{train}}^{1/2} \left(\frac{1}{m} Z^{\top} Z \right) \Sigma_{\text{train}}^{1/2},$$

where Z has standard normal entries, so that

$$t_F(z) = \bar{\operatorname{tr}} \left[\Sigma_{\text{train}}^{1/2} F \Sigma_{\text{train}}^{-1/2} \left(z I - \Sigma_{\text{train}} \frac{1}{m} Z^{\top} Z \right)^{-1} \right]. \tag{32}$$

Now define $B := I_5 - E_{5,2}$, where $E_{5,2}$ is a matrix whose (5,2) entry is 1 and has all other entries equal to 0, and let

It is straightforward to verify that $(B - (\Sigma + Q))^{-1}$ has as its (5,1) block exactly the matrix in (32), and so $t_F(z) = [g_{\Sigma + Q}(B)]_{5,1}$, where $g_{\Sigma + Q}$ is the operator-valued Cauchy transform of $\Sigma + Q$.

By rotational invariance, Σ , Q are asymptotically operator free, meaning we can apply the operator-valued additive subordination relation (see, eg. Mingo & Speicher (2017) Chapter 10), which yields the self-consistent equation

$$g_{\Sigma+Q}(B) = g_{\Sigma}(B - \mathcal{R}_Q(g_{\Sigma+Q}(B))).$$

The blocks of Q are standard normal matrices, and so its \mathcal{R} -transform is given by

$$\mathcal{R}_{Q}\left(M\right) = \mathcal{E}\left[QMQ\right],$$

where \mathcal{E} is the operator-valued expectation that takes normalized traces of all square blocks. Due to the large number of zeros in Q, only two entries of $\mathcal{R}_Q(M)$ are nonzero:

$$[\mathcal{R}_Q(M)]_{3,5} = M_{4,4}$$

 $[\mathcal{R}_Q(M)]_{4,4} = qM_{5,3}.$

On the other hand, by definition $g_{\Sigma}(M) = \mathcal{E}\left[\left(M - \Sigma\right)^{-1}\right]$. Substituting back into the subordination relation and writing g for $g_{\Sigma + Q}(B)$, we find

$$g = \mathcal{E} \left[\begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ -\frac{1}{z}\Xi & 1 & -\frac{1}{z}\Sigma_{\text{train}} & 0 & 0 \\ 0 & 0 & 1 & 0 & -g_{44} \\ 0 & 0 & 0 & 1 - qg_{53} & 0 \\ 0 & -1 & 0 & 0 & 1 \end{pmatrix}^{-1} \right],$$

where to simplify notation we have written $\Xi := \Sigma_{\mathrm{train}}^{1/2} F \Sigma_{\mathrm{train}}^{-1/2}$. The entries of the right side are straightforward to compute using elementary row operations. Performing just enough such operations to determine the (5,1), (4,4), and (5,3) entries, we obtain the closed system of equations

$$g_{53} = \bar{\text{tr}} \left[\Sigma_{\text{train}} (zI - g_{44} \Sigma_{\text{train}})^{-1} \right]$$

$$g_{44} = \frac{1}{1 - qg_{53}}$$

$$g_{51} = \bar{\text{tr}} \left[F (zI - g_{44} \Sigma_{\text{train}})^{-1} \right].$$

We can eliminate g_{53} entirely, giving our trace

$$t_f(z) = g_{51} = \bar{\operatorname{tr}}\left[F\left(zI - r\Sigma_{\mathrm{train}}\right)^{-1}\right],$$

in terms of the scalar $r := g_{44}$ that satisfies

$$r = \left(1 - q\,\bar{\text{tr}}\left[\Sigma_{\text{train}}\left(zI - r\Sigma_{\text{train}}\right)^{-1}\right]\right)^{-1}$$

A few remarks are in order. First, we note that we can rewrite this trace as an integral over the spectrum of Σ_{train} :

$$r = \left(1 - q \int \frac{t}{z - tr} d\mu_{\Sigma_{\text{train}}}(t)\right)^{-1}.$$
 (33)

It is helpful to compare (33) to the explicit expressions for g_{44}, g_{53} from the linearization before applying the subordination relation, which are

$$g_{44} = 1 + q \,\text{tr}\left[(z - A)^{-1} A\right]$$
 (34)

$$g_{53} = \bar{\operatorname{tr}} \left[(zI - A)^{-1} \Sigma_{\text{train}} \right]. \tag{35}$$

Thus $g_{44}(z)$ is analytic in z everywhere outside the spectrum of A, and $g_{44}(\mathbb{H}^{\pm}) \subset \mathbb{H}^{\mp}$ and $0 < g_{44}(\mathbb{R}^{<0}) < 1$ (the first inequality is gotten most easily by using $g_{53} < 0$ and $g_{44} = (1 - qg_{53})^{-1}$). In fact, these conditions along with the self-consistent equation (33) are enough to guarantee that the solution is unique, holomorphic, and coincides with g_{44} throughout all of $\mathbb{C}\backslash\mathbb{R}^{\geq 0}$.

Proposition 6. For $z \in \mathbb{C} \backslash \mathbb{R}^{\geq 0}$, there is a unique solution r(z) to (33) satisfying the conditions $r(\mathbb{H}^{\pm}) \subset \mathbb{H}^{\mp}$ and $0 < r(\mathbb{R}^{< 0}) < 1$. r(z) depends holomorphically on z and can be obtained by iteration of the right hand side of (33) from an arbitrary initial point in $\mathbb{C} \backslash \mathbb{R}^{\geq 0}$.

Proof. Assume $z \in \mathbb{H}^-$. Let f(r, z) be the map defined by the right hand side of (33):

$$f(r,z) := \left(1 - q \int \frac{t}{z - tr} d\mu_{\Sigma_{\text{train}}}(t)\right)^{-1}.$$

It is straightforward to check that $f(\cdot,z):\mathbb{H}^+\to\mathbb{H}^+$. Furthermore, no point on the boundary of \mathbb{H}^+ is a fixed point of $f(\cdot,z)$, since $f(\mathbb{R},z)\subset\mathbb{H}^+$, and $f(\infty,z)=1$. The Denjoy-Wolff theorem then guarantees that $f(\cdot,z)$ has a unique fixed point in \mathbb{H}^+ - and that this point is obtained by iteration of $f(\cdot,z)$ from an arbitrary initial point in \mathbb{H}^+ . Thus (33) together with the condition $r\in\mathbb{H}^+$ uniquely defines a function r(z) for all $z\in\mathbb{H}^-$.

 Now fix $z_0 \in \mathbb{H}^-$. Since $f(\cdot, z_0) : \mathbb{H}^+ \to \mathbb{H}^+$ and $f(\cdot, z_0)$ is not a Möbius transformation (it only can be if Σ_{train} is a scalar matrix), the Schwarz lemma implies $\left|\frac{\partial}{\partial r} f(r(z_0), z_0)\right| < 1$, which means

$$\left| \frac{\partial}{\partial r} \left(f\left(r, z_0 \right) - r \right) \right| = \left| \frac{\partial}{\partial r} f\left(r, z_0 \right) - 1 \right| > 0,$$

and so the implicit function theorem implies there is a holomorphic function solving (33) on some neighborhood of z_0 that coincides with $r(z_0)$ at z_0 . Since $r(z_0) \in \mathbb{H}^+$, this function must also stay in \mathbb{H}^+ in some (possibly smaller) neighborhood of z_0 , and by uniqueness of solutions to (33), this implies that it coincides with r(z) on this neighborhood. Thus $r(z) : \mathbb{H}^- \to \mathbb{H}^+$ is holomorphic at each point of \mathbb{H}^- . An identical argument proves the proposition for $z \in \mathbb{H}^+$.

Now suppose z < 0. Conjugating the right hand side of (33) by the map $x \mapsto 1/(1 - qx)$ gives a self-consistent equation satisfied by g_{53} :

$$y = \int \frac{t}{z - \frac{1}{1 - qy} t} d\mu_{\Sigma_{\text{train}}}(t).$$
 (36)

The condition 0 < r(z) < 1 implies $g_{53} < 0$. Now letting h(y,z) be the right hand side of (36), h(y,z) - y is convex in y and satisfies h(0,z) - 0 < 0 and $h(-\infty,z) - (-\infty) = \infty$, so there is a unique solution to (36) with y < 0, and thus a unique solution to (33) with 0 < r(z) < 1.

Since $\frac{\partial}{\partial u}h(y,z) > 0$, and differentiating h at the fixed point gives

$$\begin{split} \frac{\partial}{\partial y} h\left(y,z\right) &= \frac{\partial}{\partial y} \left(1 - qy\right) \int \frac{t}{z \left(1 - qy\right) - t} d\mu_{\Sigma_{\text{train}}}\left(t\right) \\ &= 1 - \frac{1}{1 - qy} + z \int \frac{\left(1 - qy\right) qt}{\left(z \left(1 - qy\right) - t\right)^2} d\mu_{\Sigma_{\text{train}}}\left(t\right) \\ &< 1, \end{split}$$

 $y\left(z\right)$ is an attracting fixed point of $h\left(\cdot,z\right)$. Since $h\left(\cdot,z\right)$ is a conjugate of $f\left(\cdot,z\right)$, the unique solution of (33) satisfying $0 < r\left(z\right) < 1$ is an attracting fixed point of $f\left(\cdot,z\right)$. This implies that there is a neighborhood of z that extends into the upper half plane whose iterates converge to $r\left(z\right)$. But since z < 0, $f\left(\mathbb{H}^+,z\right) \subset \mathbb{H}^+$, and so the Denjoy-Wolff theorem implies that all iterates of $f\left(\cdot,z\right)$ initialized in \mathbb{H}^+ converge to the same point, which therefore must be $r\left(z\right)$.

Finally, $\frac{d}{dy}\left(h\left(y\right)-y\right)=h'\left(y\right)-1<0$ implies that $y\left(z\right)$ extends holomorphically to a solution of (36) in an entire neighborhood of z. Since $h\left(y\left(w\right),w\right)-y\left(w\right)=0$ for all w in this neighborhood, at the solution point,

$$y'\left(z\right) = -\frac{\frac{\partial}{\partial z}h\left(y\left(z\right),z\right)}{\frac{\partial}{\partial y}\left(h\left(y\left(z\right),z\right) - y\left(z\right)\right)} = \frac{\int \frac{t}{\left(\frac{t}{qy-1}+z\right)^{2}}d\mu_{\Sigma_{\text{train}}}\left(t\right)}{\frac{\partial}{\partial y}\left(h\left(y\left(z\right),z\right) - y\left(z\right)\right)} < 0.$$

A negative derivative implies that for sufficiently small neighborhood U of $z, y (U \cap \mathbb{H}^-) \subset \mathbb{H}^+$.

Mapping back to r(z):=1/(1-qy(z)) yields a holomorphic function satisfying (33) in a neighborhood U of z<0 such that for $r(U\cap \mathbb{H}^-)\subset \mathbb{H}^+$. By uniqueness of solutions in the upper half plane, r must coincide with the function defined earlier on $U\cap \mathbb{H}^-$. Thus r extends holomorphically to the negative real axis.

The subordination relation implies $g_{44}(z)=r(z)$ in a neighborhood of ∞ , but both functions extend holomorphically to all of $\mathbb{C}\backslash\mathbb{R}^{\geq 0}$, implying they are equal throughout. This completes the proof of Propositions 4 and 5.

E LOCAL GEOMETRY OF MLPS

E.1 Gradient descent dynamics

The initial gradient at w_0 is $z := d\nabla J_{\text{train}}(w_0, \epsilon)$, and the Hessian is $H_{\text{train}} := d\nabla^2 J_{\text{train}}$, so the local approximation for the training loss is

$$J_{\text{train}}(w) \approx \frac{1}{d} z^{\top} (w - w_0) + \frac{1}{2d} (w - w_0)^{\top} H_{\text{train}} (w - w_0),$$

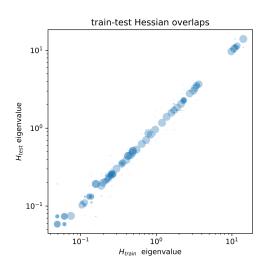


Figure 6: Eigenvector overlap function for one MLP simulation in the context of 3.3. A dot is plotted for every pair of train and test eigenvalues, with dot size and opacity representing squared overlap of the corresponding eigenvectors. Note the very strong train-test alignment indicated by the restriction of almost all overlap to the diagonal.

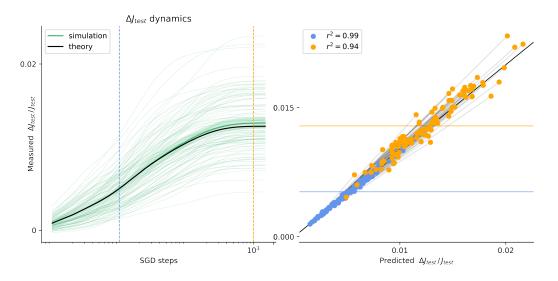


Figure 7: Learning dynamics predicted by local two-loss geometry. Left panel: Test loss trajectories in response to different label noise realizations (individual trajectories shown as thin green traces; average shown as thick green line). Noise amplitude corresponds to red dot in Fig. 4b. To reduce clutter, only theory average is shown (black trace; see (37)). Right panel: measured vs predicted relative test loss increment at two times, corresponding to the vertical blue and gold lines in the left panel. Points corresponding to the same trajectory are connected by gray lines. Horizontal blue and gold lines show means.

where we've discarded additive constants. The gradient is then

$$\nabla J_{\mathrm{train}}\left(w\right) \approx \frac{1}{d}\left(z - H_{\mathrm{train}}w_{0} + H_{\mathrm{train}}w\right),$$

so gradient descent does

$$w \to w - \eta \left(z - H_{\text{train}} w_0 + H_{\text{train}} w \right)$$

= $\left(I - \eta H_{\text{train}} \right) w - \eta \left(z - H_{\text{train}} w_0 \right)$.

Thus,

$$\Delta w_t := w_t - w_0 = \frac{\left(I - \eta H_{\text{train}}\right)^t - I}{H_{\text{train}}} z,$$

and so the test error satisfies

$$(\Delta J_{test})_t = \frac{1}{d} g^{\top} \Delta w_t + \frac{1}{2d} \Delta w_t^{\top} H_{test} \Delta w_t$$

$$= \frac{1}{d} g^{\top} \frac{(I - \eta H_{train})^t - I}{H_{train}} z + \frac{1}{2d} z \frac{(I - \eta H_{train})^t - I}{H_{train}} H_{test} \frac{(I - \eta H_{train})^t - I}{H_{train}} z,$$
(38)

with $g := d\nabla J_{\text{test}}(w_0)$ and $H_{\text{test}} := d\nabla^2 J_{\text{test}}(w_0)$.

To compute the GD trajectory for a large network, it's most efficient to precompute

$$g = \operatorname{grad}(J_{\text{train}})(w_0)$$
$$H_{\text{train}}w_0 = \operatorname{hvp}(J_{\text{train}}, w_0),$$

set $v_0 := z - H_{\text{train}} w_0$, and simply iterate

$$w \to w - \eta \left(v_0 + \text{hvp}\left(J_{\text{train}}, w\right)\right)$$

where grad, hvp compute function gradients and Hessian-vector products.

F EFFICIENT CALCULATION OF EIGENVECTOR OVERLAPS FOR LARGE MATRICES

Here we describe the Overlap-KPM algorithm, which estimates the unnormalized or total eigenvector overlaps for two matrices. That is, for self-adjoint $A, B \in \mathbb{R}^{d \times d}$, it estimates the measure

$$\mu_{A,B} := \frac{1}{d^2} \sum_{i,j=1}^{d} \delta_{\left(\lambda_i^A, \lambda_j^B\right)} \left[d \left(v_i^A \cdot v_j^B \right)^2 \right],$$

which accumulates all overlap of λ_1 A-eigenspaces with all λ_2 B-eigenspaces. To obtain the normalized overlap function treated in the main text (eg. equation (13)), one simply divides by the spectral densities of A, B at λ_1, λ_2 . The problem of estimating spectral densities for large matrices has already received significant attention (see Papyan (2019) in machine learning context), so we assume that one can practically transform back and forth between normalized and unnormalized overlap functions.

F.1 RANK DEFLATION AND NORMALIZATION

As a preprocessing step, we remove the outlier eigenspaces from each matrix obtained via subspace iteration (cf Fig. 5), and then normalize the spectra to the interval [-1, 1].

F.2 RANK DEFLATION VIA SUBSPACE ITERATION

The overlap plots of Fig. 5 were generated via Subspace Iteration (SI). SI is a well known method that generalizes power iteration, so we review it only briefly here. See Papyan (2019) for an explicit implementation in a machine learning context.

Let M be a self-adjoint operator with simple spectrum $\lambda_1, \ldots, \lambda_d$, and take $V \in \mathbb{R}^{d \times k}$ with standard normal entries and then orthonormalize the columns. SI iterates

$$V \to MV$$

 $V \to VQ$.

where Q is the Gram-Schmidt orthonormalizing upper triangular matrix of V. Informally, each application of M amplifies each i^{th} eigenspace coefficient of the columns of V by λ_i , which generically leads to exponentially greater weight on the leading eigenspaces. The orthonormalization Q prevents all eigenvectors from collapsing onto the same leading eigenvector. Since they are forced to span an k-dimensional space, they must converge to the top k eigenvectors of M. Overlaps can then be calculated direction by computing pairwise dot products of columns of V.

After convergence, outlier eigenspaces are removed from the matrices by replacing each matrix vector product $v \mapsto M(v)$ with

$$v \mapsto M_{def}(v) = M(v) - VV^{\top}v.$$

F.2.1 SPECTRUM NORMALIZATION

After removing the outlier eigenvalues, one may obtain bounds for the remaining spectrum via standard approaches (eg. the Lanczos algorithm; cf Papyan (2019)). Letting λ_{\min} , λ_{\max} denote the minimum and maximum eigenvalue (in practice, with a small amount of padding added), we the normalize the matrices to the interval [-1,1] by replacing $v\mapsto M_{def}(v)$ with

$$v\mapsto M_{norm}\left(v\right) = \frac{2}{\lambda_{\max}-\lambda_{\min}}M_{def}\left(v\right) - \left(\frac{\lambda_{\max}+\lambda_{\min}}{\lambda_{\max}-\lambda_{\min}}\right)v.$$

F.3 OVERLAP-KPM

We now assume the previous preprocessing steps have been performed and in particular that A, B's spectra lie inside [-1, 1].

First note that for kernel function G one can write the kernel-smoothed overlaps exactly as a trace:

$$\bar{\operatorname{tr}}\left[G\left(A-\lambda_{1};\sigma\right)G\left(B-\lambda_{2};\sigma\right)\right] = \frac{1}{d^{2}}\sum_{i,j=1}^{d}G\left(\lambda_{A,i}-\lambda_{1};\sigma\right)G\left(\lambda_{B,j}-\lambda_{2};\sigma\right)\left[d\left(v_{A,i}\cdot v_{B,j}\right)^{2}\right].$$
(39)

Thus the goal will be to compute such traces for each (λ_1, λ_2) for some sufficiently small fixed kernel width σ . Computing such traces directly is prohibitively expensive for very large matrices, and so a standard approach is to use Hutchinson trace estimation, ie. to average $v^{\top}Mv$ over several random samples of, say, standard normal v, since

$$\mathbb{E}_{\boldsymbol{v}}\left[\boldsymbol{v}^{\top}\boldsymbol{M}\boldsymbol{v}\right] = \operatorname{tr}\left[\boldsymbol{M}\mathbb{E}_{\boldsymbol{v}}\left[\boldsymbol{v}\boldsymbol{v}^{\top}\right]\right] = \operatorname{tr}\left[\boldsymbol{M}\right].$$

Informal experiments suggested better stability for estimation of PSD traces, so we replace the trace on the left side of (39) with

$$\bar{\operatorname{tr}}\left[G^{1/2}\left(A-\lambda_{1};\sigma\right)G\left(B-\lambda_{2};\sigma\right)G^{1/2}\left(A-\lambda_{1};\sigma\right)\right].$$

Now applying the Hutchinson trick, we sample probes v_1, \ldots, v_P and approximate

$$\begin{split} \operatorname{\bar{tr}}\left[G\left(A-\lambda_{1};\sigma\right)G\left(B-\lambda_{2};\sigma\right)\right] &\approx \frac{1}{P}\sum_{\mu=1}^{P}v_{\mu}^{\top}G^{1/2}\left(A-\lambda_{1};\sigma\right)G\left(B-\lambda_{2};\sigma\right)G^{1/2}\left(A-\lambda_{1};\sigma\right)v_{\mu} \\ &= \frac{1}{P}\sum_{\mu=1}^{P}\left\|G^{1/2}\left(B-\lambda_{2};\sigma\right)G^{1/2}\left(A-\lambda_{1};\sigma\right)v_{\mu}\right\|^{2}. \end{split}$$

To compute the summand, we generalize a standard approach known as the kernel polynomial method. Practically speaking, this entails approximating the kernel functions $G^{1/2}\left(x-\lambda;\sigma\right)$ using Chebyshev polynomials $T_{j}\left(x\right)$, which can be computed efficiently using $T_{0}\left(x\right)=1,T_{1}\left(x\right)=x$, and the recurrence

$$T_{j}(x) = 2xT_{j-1}(x) - T_{j}(x), \quad j \ge 2.$$

Letting α, β be the Chebyshev coefficients of the kernel functions,

1514
1515
$$G^{1/2}\left(x-\lambda_{1};\sigma\right)=\sum_{i=0}^{\infty}\alpha_{i}T_{i}\left(x\right)$$
1516
1517

$$G^{1/2}\left(x-\lambda_{2};\sigma\right)=\sum_{j=0}^{\infty}\beta_{j}T_{j}\left(x\right),$$

we truncate to degree K and write

$$c_{\mu} := \left\| G^{1/2} \left(B - \lambda_2; \sigma \right) G^{1/2} \left(A - \lambda_1; \sigma \right) v_{\mu} \right\|^2 \tag{40}$$

$$\approx \left\| \sum_{i,j=0}^{K} \alpha_{i} \beta_{j} T_{j} \left(B \right) T_{i} \left(A \right) v_{\mu} \right\|^{2} \tag{41}$$

$$= \sum_{i,j,k,\ell=0}^{K} \alpha_{i} \beta_{j} \alpha_{k} \beta_{\ell} v_{\mu}^{\top} T_{i} (A) T_{j} (B) T_{\ell} (B) T_{k} (A) v_{\mu}$$

$$(42)$$

$$=: \sum_{i,j,k,\ell=0}^{K} \alpha_i \beta_j \alpha_k \beta_\ell M_{i,j,k,\ell,\mu}. \tag{43}$$

Thus for m probes and order-K Chebyshev truncation, by appropriate choice of the coefficients α, β , one can approximate a general function from the $P(K+1)^4$ dot products

$$M_{i,j,k,\ell,\mu} = v_{\mu}^{\top} T_i(A) T_j(B) T_{\ell}(B) T_k(A) v_{\mu}.$$

This can be improved somewhat using the Chebyshev product identity

$$T_{j}(x) T_{\ell}(x) = \frac{1}{2} (T_{j+\ell}(x) + T_{|j-\ell|}(x)),$$
 (44)

so that

$$M_{i,j,k,\ell,\mu} = \frac{1}{2} \left(v_{\mu}^{\top} T_i(A) T_{j+\ell}(B) T_k(A) v_{\mu} + v_{\mu}^{\top} T_i(A) T_{|j-\ell|}(B) T_k(A) v_{\mu} \right),$$

and so all needed dot products can be obtained from the $P\left(K+1\right)^2\left(2K+1\right)\sim 2PK^3$ dot products

$$\tilde{M}_{i,j,k,\mu} := v_{\mu}^{\intercal} T_{i}\left(A\right) T_{j}\left(B\right) T_{k}\left(A\right) v_{\mu}, \qquad 0 \leq i, j \leq K, \quad 0 \leq j \leq 2K, \qquad 1 \leq \mu \leq P.$$

Algorithm 1 efficiently generates all such probe moments with $\sim K^2$ matrix vector products. Algorithm 1 actually stores all $2PK^2$ vectors $T_j(B)T_k(A)v_\mu$, but in practice, our implementation is significantly more memory efficient. We store all $T_k(A)v_\mu$, but as B's are added, one only needs to store the current and previous power of B. This amounts to $\sim K$ vectors in memory at once.

Once the $\tilde{M}_{i,j,k,\mu}$ are known, equation (43) is used to estimate the trace for each value of λ_1, λ_2 , yielding an approximation to the unnormalized overlap function of A, B.

Often in machine learning contexts, one or both of A,B has spectrum that is highly peaked around a particular value. For the trace in (39) to accurately reflect the overlaps at λ_1,λ_2 , the kernels more precisely, their finite K Chebyshev series - must decay sufficiently quickly away from λ_1,λ_2 to prevent the spectral spikes from overlwhelming the overlap sum. Practically speaking, this can be diagnosed by 1) forming an estimate of A,B's spectral density, eg. using the Lanczos algorithm (see Papyan (2019) for implementation in ML context), 2) forming truncated Chebyshev series for the kernels, and 3) comparing kernel decay to spike height. Insufficient decay usually requires either decreased kernel width σ , or increased Chebyshev degree K so that polynomial approximations accurately approximate the small tails needed to dampen the spectral spikes.

Algorithm 1: Overlap-KPM for Eigenvector Overlaps

```
Input: A(v), B(v) (normalized MVPs); degree K; probes P
Output: Probe moments M_{i,j,k,\mu} for 0 \le i, k \le K; 0 \le j \le 2K; 1 \le \mu \le P.

for \mu = 1 to P do

| sample probe v_{\mu} \sim \mathcal{N}(0, I_d)

v_{0,0,\mu} \leftarrow v_{\mu}; v_{0,1,\mu} \leftarrow A(v_{\mu})

for i = 2 to K do

| v_{0,i,\mu} \leftarrow A(v_{0,i-1,\mu}) - v_{0,i-2,\mu}

for k = 0 to K do

| v_{1,k,\mu} \leftarrow B(v_{0,k,\mu})

for j = 2 to 2K do

| v_{j,k,\mu} \leftarrow B(v_{j-1,k,\mu}) - v_{j-2,k,\mu}

for i = 0 to K do

| for k = 0 to k = 0
```

F.4 APPLICATION OF OVERLAP-KPM

Algorithm 1 with gaussian kernel is applied to synthetic data in Fig. 8. $A, B \in \mathbb{R}^{1000 \times 1000}$ are generated according to

$$A = W_1, B = W_2 + A^2, (45)$$

where W_1, W_2 are independent Wishart matrices with aspect ratio $\alpha = 5$. The left panel shows the ground truth gaussian-smoothed overlap function of A, B. Note the nontrivial alignment due to B's dependence on A. The right panel shows the approximation generated via Overlap-KPM, showing good qualitative match.

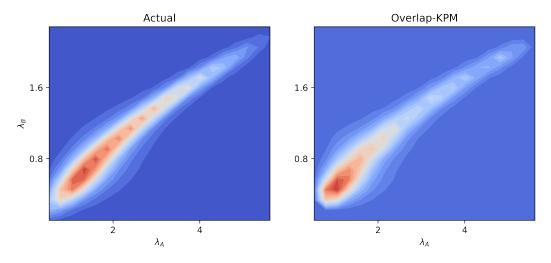


Figure 8: Overlap-KPM algorithm recovers overlaps on synthetic data. $A = Z_1^\top Z_1/m$ and $B = Z_2^\top Z_2/m + A^2$, where $Z_1, Z_2 \in \mathbb{R}^{m,d}$ are independent matrices with iid. standard normal entries. $d = 1000, \alpha := m/d = 5$. Chebyshev degree: K = 45; number of probes P = 4. Left panel shows actual eigenvector overlaps at eigenvalues λ_A, λ_B , smoothed with a gaussian kernel of with 1/3. Right panel shows approximation derived from the Overlap-KPM algorithm.

Finally, we apply Overlap-KPM to the deflated, normalized train and test Hessians of the ResNet-20 studied in 3.4 with degree K=45 and P=4 probes. Jackson smoothing is applied to gaussian kernel Chebyshev coefficients. The left panel shows the unnormalized or total overlap function

that accumulates all overlap of $\lambda_{\rm train}$ $H_{\rm train}$ -eigenspaces with all $\lambda_{\rm test}$ $H_{\rm test}$ -eigenspaces (log color scale). Due to the very large multiplicity of near-0 eigenspaces, total overlap is largest between $\lambda_{\rm train} \approx 0$ and $\lambda_{\rm test} \approx 0$ eigenspaces. Individual spectral densities for $H_{\rm train}$ ($H_{\rm test}$) are approximated by summing over rows (columns), corresponding to integrating localized kernels over all values of $\lambda_{\rm test}$ ($\lambda_{\rm train}$).

Normalized overlaps, or overlap densities - ie. the overlap function (13) studied in the text - are obtained by dividing the total overlap by the product of the individual spectral densities, $\rho_{H_{\rm train}}$ $\rho_{H_{\rm test}}$ (middle panel). (As a simple test of SNR, we confirmed that standard error of measurement estimates for $\rho_{H_{\rm train}}$ $\rho_{H_{\rm test}}$ using the P=4 independent probes were below means at each $\lambda_{\rm train}, \lambda_{\rm test}$.) Note that overlap density is highest among the larger clusters of eigenvalues, and is roughly concentrated along the diagonal, indicating rough overlap within the bulk.

Finally, we also normalize by the density of H_{test} only, effectively yielding the overlap profile of single λ_{test} -eigenvectors onto entire H_{train} eigenspaces as a function of λ_{train} . Here too one sees a general pattern of low/high λ_{test} eigenspaces overlapping with low/high λ_{train} eigenspaces.

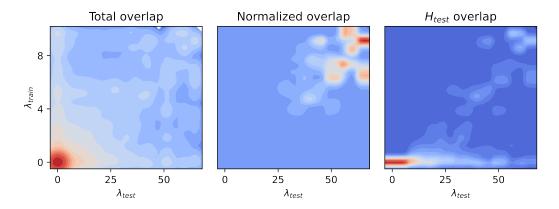


Figure 9: Overlap-KPM algorithm for overlaps of ResNet-20 train/test Hessians. Degree K=45 and P=4 probes. See text for details.