

PAPER

Risk and cross validation in ridge regression with correlated samples^{*}

To cite this article: Alexander Atanasov *et al* *J. Stat. Mech.* (2026) 094001

View the [article online](#) for updates and enhancements.

You may also like

- [A solvable model of learning generative diffusion: theory and insights](#)
Hugo Cui, Cengiz Pehlevan and Yue M Lu
- [Selecting the Regularization Parameter in the Distribution of Relaxation Times](#)
Adeleke Maradesa, Baptiste Py, Ting Hei Wan et al.
- [Hyperparameter estimation using a resolution matrix for Bayesian sensing](#)
Tatsu Kuwatani, Hideitsu Hino, Kenji Nagata et al.

Risk and cross validation in ridge regression with correlated samples *

Alexander Atanasov^{1,2,6} , Jacob A Zavatore-Veth^{2,3,6} 
and Cengiz Pehlevan^{2,4,5,**} 

¹ Department of Physics, Harvard University, Cambridge, MA, United States of America

² Center for Brain Science, Harvard University, Cambridge, MA, United States of America

³ Society of Fellows, Harvard University, Cambridge, MA, United States of America

⁴ John A. Paulson School of Engineering and Applied Sciences, Harvard University, Cambridge, MA, United States of America

⁵ Kempner Institute for the Study of Natural and Artificial Intelligence, Harvard University, Cambridge, MA, United States of America

E-mail: cpehlevan@seas.harvard.edu and jzavatoreveth@fas.harvard.edu

Received 4 June 2026

Accepted for publication 23 July 2026

Published 4 September 2026



Online at stacks.iop.org/JSTAT/2026/094001
<https://doi.org/10.1088/1742-5468/ae9378>

Abstract. Recent years have seen substantial advances in our understanding of high-dimensional ridge regression, but existing theories assume that training examples are independent. By leveraging techniques from random matrix theory and free probability, we provide sharp asymptotics for the in- and out-of-sample risks of ridge regression when the data points have arbitrary correlations. We demonstrate that in this setting, the generalized cross validation estimator (GCV) fails to correctly predict the out-of-sample risk. However, in the case where the noise residuals have the same correlations as the data points, one can modify the GCV to yield an efficiently-computable unbiased estimator that concentrates in the high-dimensional limit, which we dub CorrGCV. We further

⁶AA and JZ-V contributed equally to this work.

*This article is an updated version of: Alexander Atanasov, Jacob A Zavatore-Veth and Cengiz Pehlevan 2025 Risk and cross validation in ridge regression with correlated samples *42nd Int. Conf. on Machine Learning*.

** Authors to whom any correspondence should be addressed.

extend our asymptotic analysis to the case where the test point has nontrivial correlations with the training set, a setting often encountered in time series forecasting. Assuming knowledge of the correlation structure of the time series, this again yields an extension of the GCV estimator, and sharply characterizes the degree to which such test points yield an overly optimistic prediction of long-time risk. We validate the predictions of our theory across a variety of high-dimensional data.

Keywords: Statistical inference, Machine Learning

Contents

1. Introduction	3
2. Setup and notation	5
3. Deterministic equivalences	7
3.1. Weak deterministic equivalents	8
3.2. One-point strong deterministic equivalents	9
3.3. Two-point strong deterministic equivalents	10
4. Predicting an uncorrelated test set	10
4.1. Warm-up: linear regression without correlations	10
4.2. Correlated data with identically correlated noise	11
4.3. Mismatched correlations and OOD generalization	13
4.4. Effect on double descent	14
5. Algorithmic implementation	15
6. Testing on correlated data	17
7. Conclusion	17
Acknowledgment	19
Appendix A. Algorithmic implementation of the CorrGCV	19
A.1. Code for implementation	19
A.2. Estimating the S -transform from data	19
Appendix B. Review of free probability	21
B.1. Definition of freedom	21
B.2. The R and S transforms	21
B.3. Strong deterministic equivalence	23
Appendix C. Deferred proofs of main results	23
C.1. Warm-up: 1-point deterministic equivalents	23
C.2. 2-point deterministic equivalents	26
C.3. Duality relations and derivatives	29

C.4. Correlated samples and noise, uncorrelated test point	31
C.5. Correlated test point	33
Appendix D. Bias-variance decompositions	35
Appendix E. Scaling analysis	36
E.1. Review of optimal rates	36
E.2. Strong correlations do not affect scaling exponents	37
Appendix F. Double descent analysis	40
F.1. Bounds on renormalized ridges	41
F.2. Ridgeless limit	42
F.2.1. Effective ridge enhancement	43
F.3. Mismatched correlations	44
F.4. Further comments on the effect of mismatched correlations	45
Appendix G. Asymptotics of previously-proposed extensions of the GCV	46
G.1. The GCV	46
G.2. The Altman estimator	47
G.3. The GCCV estimator of Carmack <i>et al</i>	47
Appendix H. Weighted risks and the MMSE estimator	49
Appendix I. S-transforms for certain Toeplitz covariance matrices	50
I.1. Nearest-neighbor correlations	50
I.2. Exponential correlations	52
Appendix J. Further experiments	52
J.1. Exponential correlations	52
J.2. Nearest-neighbor correlations	52
J.3. Power-law correlations	53
Appendix K. Experimental details	59
References	60

1. Introduction

Statistics classically assumes that one has access to independent and identically distributed (i.i.d.) samples. However, this fundamental assumption is often violated when one considers data sampled from a time series—e.g. in the case of financial, climate, or neuroscience data [1–3]—rendering results obtained under the i.i.d. assumption inapplicable. In particular, estimators of the out-of-sample prediction risk based on cross-validation assume independence [4–10]. To enable accurate prediction of risk for regression from time series data, it is imperative that correlations be taken into account.

In the paradigmatic setting of ridge regression, a significant body of recent research has aimed to provide sharp asymptotic characterizations of the out-of-sample risk when the in-sample risk is estimated using i.i.d. samples [11–21]. The crucial feature of these high-dimensional asymptotics is that they allow both the dimensionality of the covariates and the number of examples to be large⁶. This body of work reveals two broad principles: First, ridge regression has a spectral bias towards learning functions aligned with eigenfunctions of the feature covariance [14, 17–19]. Second, ridge regression displays *Gaussian universality* in high dimension, i.e. the out-of-sample risk for a given dataset will under general conditions be asymptotically identical to that for a Gaussian dataset with matched first and second moments [11, 23, 24]. Importantly, nearly all of these works assume the training examples are i.i.d.. A rare exception is a recent paper by Bigot *et al* [25], who allow for *independent but non-identically distributed Gaussian data* with per-example modulation of total variance.

These sharp asymptotics are *omniscient* risk estimates: they assume one has access to the true joint distribution of covariates and labels from which the training set is sampled. If one wants to perform hyperparameter tuning in practice—for ridge regression this of course means tuning the ridge parameter—a non-omniscient estimate is required. A standard approach to this problem is the method of generalized cross-validation, which dates back at least to the 1970 s [4–6]. The cross validation estimator (GCV) estimates the out-of-sample risk by applying a multiplicative correction to the in-sample risk, which itself can be estimated from the data. Recent works have shown that the GCV is asymptotically exact in high dimensions: the limiting GCV estimate coincides with the omniscient asymptotic [11, 16, 20]. Though classical cross-validation and the GCV estimator assume i.i.d. training data points, several works have aimed to extend these methods to cases in which there are correlations [26–31]. However, most of these works focus on correlations only in the label noise, and none of them show that their estimators are asymptotically exact in high dimensions. Indeed, we will show that none of the previously-proposed corrections to the GCV are asymptotically exact in high dimensions, and require further modification to accurately predict out-of-sample risk.

Here, we fill this gap in understanding by providing a detailed asymptotic characterization of ridge regression with correlated samples. We first show that even under relatively mild sample-sample correlations previously-obtained omniscient risk asymptotics assuming i.i.d. data are not predictive. Correspondingly, the ordinary GCV estimator fails to accurately estimate the out-of-sample risk. We then compute sharp high-dimensional asymptotics for the out-of-sample risk when the training examples are drawn from a general matrix Gaussian with anisotropic correlations across both features and samples. When the test point is uncorrelated with the training data and the label noise has the same correlation structure as the covariates, we show that there exists a corrected GCV estimator. We term this estimator the **CorrGCV**. Unlike previous attempts to correct the original GCV, CorrGCV is asymptotically exact (figure 1)⁷.

⁶ In contrast, classical asymptotics assume that the dimension is small relative to the number of examples [22].

⁷ Code to compute the CorrGCV and reproduce all of our experiments is publicly available on GitHub at <https://github.com/Pehlevan-Group/S-transform>.

Risk and cross validation in ridge regression with correlated samples




Figure 1. Empirical risk $\hat{R}$, out-of-sample risk R , and fine-grained bias-variance decompositions for ridge regression with structured features and correlated examples. Theory is plotted in solid lines. Experiments with error bars over 10 dataset repetitions are plotted as markers. The data points are exponentially correlated as $\mathbb{E}[\mathbf{x}_t \cdot \mathbf{x}_s] \propto e^{-|t-s|/\xi}$. Left: Weak correlations, $\xi = 10^{-2}$. Here, the generalized cross validation method (orchid) as well as its other proposed corrections in the presence of correlations (pink, purple) all agree and are overlaid. Right: Strong correlations, $\xi = 10^2$. Here, we see that the naive estimates of the GCV proposed in prior works fail in this setting. They either underestimate (purple) or overestimate (pink) the out-of-sample risk. We define the naive GCVs in the text, and connect them with prior proposals in appendix G. By contrast, our proposed estimator, CorrGCV, correctly predicts the out-of-sample risk in all settings.

In deriving the CorrGCV, we uncover an interesting duality between train-test covariate shift and a mismatch in the covariate and noise correlations. Finally, we extend all of these results to the case when the test point is correlated with the training points. Focusing on the setting of time series regression, this gives us a sharp characterization of how accuracy depends on prediction horizon, and makes precise the notion that testing on near-horizon data gives an overly optimistic picture of long-term forecast accuracy. In all, our results both advance the theoretical understanding of ridge regression in high dimensions and provide application to time series data.

2. Setup and notation

We begin by briefly introducing the setup of our work, and fixing notation. We consider ridge regression on a dataset $\mathcal{D} = \{\mathbf{x}_t, y_t\}_{t=1}^T$ of T data points, with N -dimensional covariates $\mathbf{x}_t \in \mathbb{R}^N$ and scalar labels y_t . We minimize the mean-squared error over this dataset, with a ridge penalty:

$$L(\mathbf{w}) = \frac{1}{T} \sum_{t=1}^T (y_t - \mathbf{x}_t^\top \mathbf{w})^2 + \lambda \|\mathbf{w}\|^2. \quad (1)$$

We write $\hat{\mathbf{w}} = \operatorname{argmin}_{\mathbf{w}} L(\mathbf{w})$. As long as $\lambda > 0$ the solution is unique. We will consider both the underparameterized case of classical statistics where $T > N$ as well as the

modern overparameterized setting where $N > T$ [11]. We define the overparameterization ratio $q \equiv N/T$. All of our results will hold exactly in the limit of $N, T \rightarrow \infty$ with q fixed.

We now state our statistical assumptions on the data. Defining the design matrix $\mathbf{X} \in \mathbb{R}^{T \times N}$ such that $\mathbf{X}_{ti} = [\mathbf{x}_t]_i$ and collecting the labels into a vector $\mathbf{y} \in \mathbb{R}^T$, we assume the labels are generated from a deterministic linear ‘teacher’ $\bar{\mathbf{w}} \in \mathbb{S}^{N-1}$ plus noise $\boldsymbol{\epsilon} \in \mathbb{R}^T$ as:

$$\mathbf{y} = \mathbf{X}\bar{\mathbf{w}} + \boldsymbol{\epsilon}. \quad (2)$$

We assume that $\mathbf{X}$ is Gaussian, with zero mean and

$$\mathbb{E}[x_{i,t}x_{j,s}] = \Sigma_{ij}K_{ts} \quad (3)$$

for a feature-feature covariance $\Sigma \in \mathbb{R}^{N \times N}$ and a sample-sample correlation $\mathbf{K} \in \mathbb{R}^{T \times T}$. We have the representation $\mathbf{X} = \mathbf{K}^{1/2}\mathbf{Z}\Sigma^{1/2}$ for $\mathbf{Z} \in \mathbb{R}^{T \times N}$ a matrix with i.i.d. standard Gaussian elements $Z_{ti} \sim \mathcal{N}(0,1)$, where $\Sigma^{1/2}$ denotes the principal square root of the positive-definite symmetric matrix Σ . We assume that the noise $\boldsymbol{\epsilon}$ is independent of $\mathbf{X}$, and is Gaussian with mean zero and covariance

$$\mathbb{E}[\epsilon_t\epsilon_s] = \sigma_\epsilon^2 K'_{ts}. \quad (4)$$

Without loss of generality, we will always assume the normalization $\frac{1}{T}\text{Tr}(\mathbf{K}) = 1$, as the overall scale can be absorbed into Σ . We will often assume that the data are statistically stationary, in which case $K_{tt} = 1$ for all t ; we will explicitly highlight when we impose this condition. Similarly, we normalize $\frac{1}{T}\text{Tr}(\mathbf{K}') = 1$, as the scale here can be absorbed into σ_ϵ .

We contrast this with prior treatments of linear regression in the proportional limit where $\mathbf{K}, \mathbf{K}'$ were chosen to be the identity matrix [11–21], and where the original GCV estimator [4] applies. In section 4.2 we consider the case where $\mathbf{K} = \mathbf{K}'$. There, we show that the GCV estimator [4] has a natural analogue. In section 4.3 we consider the more general case and show that there is an obstruction to a GCV estimator.

We define the **empirical covariance** $\hat{\Sigma} \equiv \frac{1}{T}\mathbf{X}^\top \mathbf{X}$ and **kernel Gram matrix** $\hat{\mathbf{K}} \equiv \frac{1}{T}\mathbf{X}\mathbf{X}^\top$. Writing $\mathbf{y} \in \mathbb{R}^T$ as the vector of training labels and $\boldsymbol{\epsilon} \in \mathbb{R}^T$ as the vector of label noises, we have:

$$\hat{\mathbf{w}} = (\hat{\Sigma} + \lambda)^{-1} \frac{\mathbf{X}^\top \mathbf{y}}{T} = \hat{\Sigma}(\hat{\Sigma} + \lambda)^{-1} \bar{\mathbf{w}} + (\hat{\Sigma} + \lambda)^{-1} \frac{\mathbf{X}^\top \boldsymbol{\epsilon}}{T}. \quad (5)$$

We also let $\hat{\mathbf{y}} = \mathbf{X}\hat{\mathbf{w}}$ be the vector of predictions on the training set. We are interested in analytically characterizing the in-sample and out-of-sample risks. The in-sample risk is defined as:

$$\hat{R}_{\text{in}}(\hat{\mathbf{w}}) \equiv \frac{1}{T} \|\mathbf{y} - \hat{\mathbf{y}}\|^2. \quad (6)$$

For the out-of-sample risk, we consider a held-out test point $\mathbf{x} \sim \mathcal{N}(0, \Sigma')$ and label noise $\epsilon \sim \mathcal{N}(0, \sigma_\epsilon^2)$. If $\Sigma = \Sigma'$, the test point is drawn from the same marginal distribution as a single example $\mathbf{x}_t, \epsilon_t$, discarding the time-dependent scale factor K_{tt} in the marginal.

More generally, though, we can allow distribution shifts with $\Sigma' \neq \Sigma$. This leads to a general expression for the out-of-sample risk:

$$R_{\text{out}}(\hat{\mathbf{w}}) \equiv \mathbb{E}_{\mathbf{x}, \epsilon} (\mathbf{x}^\top \bar{\mathbf{w}} + \epsilon - \mathbf{x}^\top \hat{\mathbf{w}})^2 = \underbrace{(\bar{\mathbf{w}} - \hat{\mathbf{w}})^\top \Sigma' (\bar{\mathbf{w}} - \hat{\mathbf{w}})}_{R_g} + \sigma_\epsilon^2. \quad (7)$$

Here we have identified the generalization error R_g as the excess risk, i.e. R_{out} minus the Bayes error σ_ϵ^2 . We study the case where the test point and label noise are drawn independently of the training set in section 4, and extend this result to the more general case in which $\mathbf{x}, \epsilon$ have nontrivial correlations with the training set in section 6.

To motivate the random matrix theory analysis in the sequel, we note that substituting the explicit form of $\hat{\mathbf{w}}$ from (5) into the generalization error gives

$$R_g \simeq \underbrace{\lambda^2 \bar{\mathbf{w}}^\top (\hat{\Sigma} + \lambda)^{-1} \Sigma' (\hat{\Sigma} + \lambda)^{-1} \bar{\mathbf{w}}}_{\text{Signal}} + \underbrace{\sigma_\epsilon^2 \frac{1}{T^2} \text{Tr}[\Sigma' (\hat{\Sigma} + \lambda)^{-1} \mathbf{X}^\top \mathbf{K}' \mathbf{X} (\hat{\Sigma} + \lambda)^{-1}]}_{\text{Noise}}, \quad (8)$$

where we highlight that the risk is self-averaging with respect to ϵ at large T , and splits into two contributions that can be conceptually identified as signal and noise, because the noise term arises from the contribution of ϵ to $\hat{\mathbf{w}}$. This sets up the goals of the subsequent analysis: we want to understand how the random matrices $(\hat{\Sigma} + \lambda)^{-1} \Sigma' (\hat{\Sigma} + \lambda)^{-1}$ and $\frac{1}{T} (\hat{\Sigma} + \lambda)^{-1} \mathbf{X}^\top \mathbf{K}' \mathbf{X} (\hat{\Sigma} + \lambda)^{-1}$ behave in the high-dimensional limit, in the sense of their normalized traces against suitable test matrices. This information is provided by the so-called strong deterministic equivalents we review and prove below.

We note in passing that one could instead consider an estimator based on a weighted loss, i.e. $\hat{\mathbf{w}}_M = \arg \min_{\mathbf{w}} L_M(\mathbf{w})$ for

$$L_M(\mathbf{w}) = \frac{1}{T} (\mathbf{X} \mathbf{w} - \mathbf{y})^\top \mathbf{M} (\mathbf{X} \mathbf{w} - \mathbf{y}) + \lambda |\mathbf{w}|^2 \quad (9)$$

for some positive-definite matrix $\mathbf{M}$ [1, 22]. Indeed, the Bayesian minimum mean squared error estimator is equivalent to minimizing this loss with $\mathbf{M} = (\mathbf{K}')^{-1}$ (appendix H). However, under our assumptions on the data this is equivalent to considering an isotropic loss under the mapping $\mathbf{K} \leftarrow \mathbf{M}^{1/2} \mathbf{K} \mathbf{M}^{1/2}$ and $\mathbf{K}' \leftarrow \mathbf{M}^{1/2} \mathbf{K}' \mathbf{M}^{1/2}$. As a result, our asymptotics for R_g apply immediately to general choices of $\mathbf{M}$.

3. Deterministic equivalences

We now introduce the key tools of our analysis: deterministic equivalents for the sample covariance. We first review standard aspects of the free probability approach to random matrices, and then state the required deterministic equivalents. Our presentation of free probability follows Potters and Bouchaud [32].

3.1. Weak deterministic equivalents

We define the first and second **degrees of freedom** of a matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$ as

$$\begin{aligned} \text{df}_{\mathbf{A}}^1(\lambda) &= \frac{1}{N} \text{Tr}[\mathbf{A}(\mathbf{A} + \lambda)^{-1}], \\ \text{df}_{\mathbf{A}}^2(\lambda) &= \frac{1}{N} \text{Tr}[\mathbf{A}^2(\mathbf{A} + \lambda)^{-2}]. \end{aligned} \quad (10)$$

When it is clear from context, we will write these as df_1, df_2 . The S -transform $S(\text{df})$ is a formal function of a variable df defined as

$$S_{\mathbf{A}}(\text{df}) \equiv \frac{1 - \text{df}}{\text{df} \text{df}_{\mathbf{A}}^{-1}(\text{df})}. \quad (11)$$

Here $\text{df}_{\mathbf{A}}^{-1}$ is the functional inverse of the equation $\text{df} = \text{df}_{\mathbf{A}}^1(\lambda)$, i.e. $\text{df}_{\mathbf{A}}^{-1}(\text{df}_{\mathbf{A}}^1(\lambda)) = \lambda$. We will use the symbol df to highlight its role as a formal variable in the S -transform, whereas we will call df_1 the actual value of the degrees of freedom for a given matrix $\mathbf{A}$ at a ridge λ . Our definition of $S_{\mathbf{A}}$ differs by a sign from the common definition of the S transform in terms of the t -function $t_{\mathbf{A}}(\lambda) = -\text{df}_{\mathbf{A}}^1(-\lambda)$. (11) implies that:

$$\text{df}_{\mathbf{A}}^1(\lambda) = \frac{1}{1 + \lambda S_{\mathbf{A}}(\text{df}_{\mathbf{A}}^1(\lambda))}. \quad (12)$$

The S -transform plays a crucial role in high-dimensional random matrix theory and free probability because for two symmetric matrices $\mathbf{A}, \mathbf{B}$ that are **free** of one another, taking the symmetrized product $\mathbf{A} * \mathbf{B} \equiv \mathbf{A}^{1/2} \mathbf{B} \mathbf{A}^{1/2}$ with $\mathbf{A}^{1/2}$ the principal matrix square root yields:

$$S_{\mathbf{A} * \mathbf{B}}(\text{df}) = S_{\mathbf{A}}(\text{df}) S_{\mathbf{B}}(\text{df}). \quad (13)$$

This implies the following **subordination relation**:

$$\begin{aligned} \text{df}_{\mathbf{A} * \mathbf{B}}^1(\lambda) &= \frac{1}{1 + \lambda S_{\mathbf{A} * \mathbf{B}}(\text{df}_{\mathbf{A} * \mathbf{B}}^1(\lambda))} = \frac{1}{1 + \lambda S_{\mathbf{A}}(\text{df}_{\mathbf{A} * \mathbf{B}}^1(\lambda)) S_{\mathbf{B}}(\text{df}_{\mathbf{A} * \mathbf{B}}^1(\lambda))} \\ &= \text{df}_{\mathbf{A}}^1(\lambda S_{\mathbf{B}}(\text{df}_{\mathbf{A} * \mathbf{B}}^1(\lambda))). \end{aligned} \quad (14)$$

We define what it means for two random variables to be free, and give a longer discussion on subordination relations and R and S transforms in appendix B. Often, one views $\mathbf{B}$ as a source of multiplicative noise, and $\mathbf{A}$ as deterministic. Then $\mathbf{A} * \mathbf{B}$ is a random matrix. The above equation thus relates the degrees of freedom of a random matrix to the degrees of freedom of a deterministic one. For this reason it is known as (weak) **deterministic equivalence**.

We now specialize to the Wishart matrices $\hat{\Sigma} = \frac{1}{T} \mathbf{X}^{\top} \mathbf{X}$ and $\hat{\mathbf{K}} = \frac{1}{T} \mathbf{X} \mathbf{X}^{\top}$ we consider here, though a subset of the subsequent results extend to more general ensembles (see appendix B and [16, 32]). Define the shorthand $\text{df}_1 \equiv \text{df}_{\Sigma}^1(\kappa)$, $\tilde{\text{df}}_1 \equiv \text{df}_{\mathbf{K}}^1(\tilde{\kappa})$, $\text{df}_2 \equiv \text{df}_{\Sigma}^2(\kappa)$, and $\tilde{\text{df}}_2 \equiv \text{df}_{\mathbf{K}}^2(\tilde{\kappa})$. Then, we define κ and $\tilde{\kappa}$ via:

$$\kappa = \lambda S(\text{df}_1), \quad \tilde{\kappa} = \lambda \tilde{S}(\tilde{\text{df}}_1). \quad (15)$$

As the multiplicative noise is a *structured* Wishart matrix, the S -transforms appearing in (15) are

$$\begin{aligned} S(\mathrm{df}_1) &= S_{\frac{1}{T}\mathbf{Z}^\top\mathbf{Z}}(\mathrm{df}_1)S_{\mathbf{K}}\left(\frac{N}{T}\mathrm{df}_1\right), \\ \tilde{S}(\tilde{\mathrm{df}}_1) &= S_{\frac{1}{T}\mathbf{Z}\mathbf{Z}^\top}(\tilde{\mathrm{df}}_1)S_{\Sigma}\left(\frac{T}{N}\tilde{\mathrm{df}}_1\right). \end{aligned} \quad (16)$$

For uncorrelated data ($\mathbf{K} = \mathbf{I}$), $S_{\mathbf{K}} = 1$. Weak deterministic equivalents for Wishart matrices with general correlation structure have long been a subject of study in random matrix theory due to their applications to covariance matrix estimation [32–35].

3.2. One-point strong deterministic equivalents

The equivalence (14) extends to a class of ‘equalities’ of matrices known as strong deterministic equivalents:

Definition 3.1 (Strong deterministic equivalence). For two sequences of (possibly random) matrices $\mathbf{A}$ and $\mathbf{B}$ indexed by their common size N , we say that $\mathbf{A}$ and $\mathbf{B}$ are **deterministically equivalent** and write $\mathbf{A} \simeq \mathbf{B}$ if $\mathrm{Tr}(\mathbf{A}\mathbf{M})/\mathrm{Tr}(\mathbf{B}\mathbf{M}) \rightarrow 1$ in probability as $N \rightarrow \infty$ for any sequence of test matrices $\mathbf{M}$ of bounded spectral norm⁸.

For the Wishart matrices $\hat{\Sigma}$ and $\hat{\mathbf{K}}$, we have:

Lemma 3.1. Define κ and $\tilde{\kappa}$ as in (15). Then,

$$\begin{aligned} \hat{\Sigma}(\hat{\Sigma} + \lambda)^{-1} &\simeq \Sigma(\Sigma + \kappa)^{-1}, \\ \hat{\mathbf{K}}(\hat{\mathbf{K}} + \lambda)^{-1} &\simeq \mathbf{K}(\mathbf{K} + \tilde{\kappa})^{-1}. \end{aligned} \quad (17)$$

Proof. We derive (17) using a diagrammatic argument in appendix C.1; see Atanasov *et al* [32] Potters and Bouchaud [37] and Bun *et al* [16] for alternative proofs. We note that one could in principle compute finite-size corrections to this asymptotic result using the same diagrammatic approach, but we will not do so here. $\square$

Using the one-point deterministic equivalents, we see that κ and $\tilde{\kappa}$ are related by the identity

$$q\mathrm{df}_1 \equiv q\mathrm{df}_{\Sigma}^1(\kappa) \simeq q\mathrm{df}_{\Sigma}^1(\lambda) = \mathrm{df}_{\hat{\mathbf{K}}}^1(\lambda) \simeq \mathrm{df}_{\mathbf{K}}(\tilde{\kappa}) \equiv \tilde{\mathrm{df}}_1. \quad (18)$$

In addition, we show in appendix C that κ and $\tilde{\kappa}$ satisfy the *duality relation*:

$$\frac{\kappa\tilde{\kappa}}{\lambda} = \frac{1}{\tilde{\mathrm{df}}_1}. \quad (19)$$

In physical terms, we can interpret (17) as a renormalization effect: the effect of the random fluctuations in $\mathbf{B}$ can be absorbed into a **renormalized ridge** κ [16]. This renormalization generates implicit regularization. Even in the limit of zero regularization, one can have $\lim_{\lambda \rightarrow 0} \kappa > 0$ (appendix F).

⁸ Our convention here follows Bach [36].

3.3. Two-point strong deterministic equivalents

As (17) involves a trace against a single test matrix, we refer to it as a ‘one-point’ equivalent. As we can see from (8), we also require ‘two-point’ equivalents involving pairs of resolvents:

Lemma 3.2. *Let Σ' be an $N \times N$ test matrix. Then,*

$$(\hat{\Sigma} + \lambda)^{-1} \Sigma' (\hat{\Sigma} + \lambda)^{-1} \simeq S^2 (\Sigma + \kappa)^{-1} \Sigma' (\Sigma + \kappa)^{-1} + S^2 (\Sigma + \kappa)^{-2} \Sigma \frac{\gamma_{\Sigma, \Sigma'}}{1 - \gamma}, \quad (20)$$

where we define

$$\gamma \equiv \frac{\text{df}_2 \tilde{\text{df}}_2}{\text{df}_1 \tilde{\text{df}}_1}, \quad \text{and} \quad \gamma_{\Sigma, \Sigma'} \equiv \frac{\text{df}_{\Sigma, \Sigma'}^2 \text{df}_{\mathbf{K}}^2}{\text{df}_1 \tilde{\text{df}}_1} \quad (21)$$

for

$$\text{df}_{\Sigma, \Sigma'}^2 \equiv \frac{1}{N} \text{Tr} [\Sigma \Sigma' (\Sigma + \kappa)^{-2}]. \quad (22)$$

Proof. See appendix C.2. □

When $\mathbf{K} = \mathbf{I}$, this recovers two-point equivalents proved in previous works [36, 38]. The first term is a ‘disconnected’ component coming from separately averaging the two resolvents, while the second term is a ‘connected’ component coming from averaging them together. These two contributions correspond to the bias-variance decomposition of the estimator over the training data. An analogous deterministic equivalent holds for $\mathbf{K}'$; we state this explicitly in appendix C.

We also derive the final two-point equivalent appearing in (8):

Lemma 3.3. *Let $\mathbf{K}'$ be a $T \times T$ test matrix. Then,*

$$\frac{1}{T} (\hat{\Sigma} + \lambda)^{-1} \mathbf{X}^\top \mathbf{K}' \mathbf{X} (\hat{\Sigma} + \lambda)^{-1} \simeq \lambda^2 S^2 \tilde{S}^2 (\Sigma + \kappa)^{-2} \Sigma \frac{\text{df}_{\mathbf{K}, \mathbf{K}'}^2}{1 - \gamma}, \quad (23)$$

where we let $\text{df}_{\mathbf{K}, \mathbf{K}'}^2 \equiv \frac{1}{T} \text{Tr} [\mathbf{K} \mathbf{K}' (\mathbf{K} + \tilde{\kappa})^{-2}]$.

Proof. See appendix C.2. □

4. Predicting an uncorrelated test set

4.1. Warm-up: linear regression without correlations

We begin by reviewing the known results in the case where the data points are assumed to be drawn i.i.d. from a distribution with covariance Σ . This is case of $\mathbf{K} = \mathbf{I}$ introduced above. In this case, we require only κ and not $\tilde{\kappa}$, which satisfies the simplified equation

$$\kappa = \lambda S_{\frac{1}{T} \mathbf{Z}^\top \mathbf{Z}} (\text{df}_1) = \frac{\lambda}{1 - q \text{df}_1}, \quad (24)$$

where $\text{df}_1 \equiv \text{df}_{\Sigma}^1(\kappa) \simeq \text{df}_{\Sigma}^1(\lambda)$. Here we have used that the S -transform of a Wishart matrix is $S_{\frac{1}{T}} \mathbf{Z}^T \mathbf{Z} = (1 - q\text{df}_1)^{-1}$. At this point, we note that depending on whether one picks $\text{df}_1 = \text{df}_{\Sigma}^1(\kappa)$ or $\text{df}_1 = \text{df}_{\Sigma}^1(\lambda)$, this either gives a self-consistent equation given omniscient knowledge of Σ or a way to estimate κ from the data alone, namely by computing $\text{df}_{\Sigma}^1(\lambda)$. Then, one has the following deterministic equivalents:

Theorem 4.1. *For uncorrelated data ($\mathbf{K} = \mathbf{I}$), one has*

$$R_g \simeq \frac{\kappa^2}{1 - \gamma} \bar{\mathbf{w}}^T \Sigma (\Sigma + \kappa)^{-2} \bar{\mathbf{w}} + \frac{\gamma}{1 - \gamma} \sigma_{\epsilon}^2, \quad (25)$$

where $\gamma = q\text{df}_2$, and

$$\hat{R}_{in} \simeq \frac{\lambda^2}{1 - \gamma} \bar{\mathbf{w}}^T \Sigma (\Sigma + \kappa)^{-2} \bar{\mathbf{w}} + \frac{\lambda^2}{\kappa^2} \frac{1}{1 - \gamma} \sigma_{\epsilon}^2. \quad (26)$$

Proof. This result is well-known [11, 12, 14–16, 19]. It follows as a special case of the proof in appendix C.4. $\square$

This yields the **generalized-cross-validation estimator** or **GCV** [4, 5]:

$$R_{\text{out}} \simeq \frac{\kappa^2}{\lambda^2} \hat{R}_{in} = S^2 \hat{R}_{in}. \quad (27)$$

Since S depends only on df_1 , which can be estimated solely from the data, the right hand side of this equation yields a way to estimate the out-of-sample error from the training error alone. The GCV estimator has also appeared in recent literature as the **kernel alignment risk estimator** or KARE [20]. The fact that the GCV estimator is directly related to the S -transform in the free probability regime $N, T \rightarrow \infty$ was first pointed out in our previous work [16].

4.2. Correlated data with identically correlated noise

We now state the results for general correlated data with matched correlations between the covariates and noise ($\mathbf{K} = \mathbf{K}'$). Using the deterministic equivalents introduced in section 3, we obtain:

Theorem 4.2. *Assume $\mathbf{K} = \mathbf{K}'$. Then, we have*

$$R_g \simeq \frac{\kappa^2}{1 - \gamma} \bar{\mathbf{w}}^T \Sigma (\Sigma + \kappa)^{-2} \bar{\mathbf{w}} + \frac{\gamma}{1 - \gamma} \sigma_{\epsilon}^2, \quad (28)$$

where

$$\gamma \equiv \frac{\text{df}_2 \tilde{\text{df}}_2}{\text{df}_1 \tilde{\text{df}}_1}. \quad (29)$$

Similarly, we obtain:

$$\hat{R}_{in} \simeq \frac{\tilde{\text{df}}_1 - \tilde{\text{df}}_2}{S \tilde{\text{df}}_1} \frac{\kappa^2}{1 - \gamma} \bar{\mathbf{w}}^T \Sigma (\Sigma + \kappa)^{-2} \bar{\mathbf{w}} + \frac{\tilde{\text{df}}_1 - \tilde{\text{df}}_2}{S \tilde{\text{df}}_1} \frac{1}{1 - \gamma} \sigma_{\epsilon}^2. \quad (30)$$

Risk and cross validation in ridge regression with correlated samples




Figure 2. Estimating the optimal ridge parameter for exponential correlations using the CorrGCV. The setup here is as in figure 1. We see that only the CorrGCV accurately predicts the out-of-sample risk, and thus is the only estimator that allows one to correctly pinpoint the optimal ridge parameter λ .

Proof. See appendix C.4. □

This yields an extension of the GCV estimator to correlated data. We call this the **CorrGCV** for **correlated generalized cross-validation**:

$$R_{\text{out}} = S(\text{df}_1) \frac{\tilde{\text{df}}_1}{\tilde{\text{df}}_1 - \tilde{\text{df}}_2} \hat{R}_{\text{in}}. \quad (31)$$

This estimator is unbiased and asymptotically exact in the proportional limit of $T, N \rightarrow \infty$. Moreover, it concentrates over draws of the dataset for T sufficiently large.

Like the ordinary GCV, the CorrGCV can be estimated from the training data alone. We give the explicit algorithm in appendix A. In figures 1–4 and 6, and appendices G and J we compare this estimator with other previously proposed estimators of out-of-sample risk in correlated data. In particular, we compare against the ordinary GCV

$$\text{Naive GCV}_1 = \frac{1}{(1 - q\text{df}_1)^2}, \quad (32)$$

which does not take into account the sample-sample correlations, and the estimator proposed by Altman [26],

$$\text{Naive GCV}_2 = S^2 = (\kappa/\lambda)^2, \quad (33)$$

which was designed to account for sample-sample correlations in the label noise but not in the covariates. Further comparisons to the estimator of Carmack *et al* [28] are given in appendix G. When the samples are correlated, only the CorrGCV accurately predicts the out-of-sample risk. As a result, it is the only correction to the GCV that allows accurate tuning of the ridge parameter λ (figure 2).

An important contribution from the theory of high-dimensional regression with uncorrelated samples is a characterization of how power-law decays in covariance

Risk and cross validation in ridge regression with correlated samples




Figure 3. Power law scalings for data with (a) exponential correlations with $\xi = 10^2$ and (b) power law correlations $\mathbb{E}\mathbf{x}_t^\top \mathbf{x}_{t+\tau} \propto \tau^{-\chi}$ with $\chi = 0.3$. In both cases, the correlations of the data do not affect the scaling of the generalization error as a function of T , which generally goes as $T^{-2\alpha\min(r,1)}$, as derived in prior works. Although other estimators correctly predict the rate of decay, only the CorrGCV correctly recovers the exact risk.

eigenspectra—as occur for a variety of real data—give rise to power-law decays in out-of-sample risk R_g as a function of sample size T [16, 19, 39–41]. Thus, it is natural to ask whether correlations between samples can alter these scaling laws. We show in appendix E that the correlation structure of the stationary processes we consider cannot change the decay rates. Concretely, assume that Σ has eigenvalues $\lambda_k \sim k^{-\alpha}$ for an exponent α known as the **capacity**. Further, the signal $\bar{\mathbf{w}}$ has that $\lambda_k \bar{w}_k^2 \sim k^{-(2\alpha r + 1)}$ for an exponent r known as the **source**. For correlated data, the scaling is unchanged from prior predictions of optimal rates in [19, 39–41], namely that $R_g \sim T^{-2\alpha\min(r,1)}$. We illustrate this phenomenon in figure 3.

4.3. Mismatched correlations and OOD generalization

We generalize the above result to a setting where ϵ does not have the same correlation structure as $\mathbf{K}$. That is, $\mathbb{E}[\epsilon_t \epsilon_s] = \sigma_\epsilon^2 K'_{ts}$. In fact, we can consider an even more general case: namely when the covariance of the test point Σ' is also different from the covariance of the training set Σ . This is the case of out-of-distribution generalization under covariate shift. These two different mismatches exhibit a surprising duality. In general, we have

Theorem 4.3. *Consider a covariate-shifted setting with test covariance Σ' not necessarily equal to Σ , and noise covariance $\mathbf{K}'$ that may not match $\mathbf{K}$. Then, we have*

$$R_g \simeq \underbrace{\kappa^2 \bar{\mathbf{w}}^\top (\Sigma + \kappa)^{-1} \Sigma' (\Sigma + \kappa)^{-1} \bar{\mathbf{w}}}_{\text{Bias}^2} + \underbrace{\kappa^2 \frac{\gamma_{\Sigma, \Sigma'}}{1 - \gamma} \bar{\mathbf{w}}^\top \Sigma (\Sigma + \kappa)^{-2} \bar{\mathbf{w}}}_{\text{Var}_X} + \underbrace{\frac{\gamma_{\Sigma, \Sigma', \mathbf{K}, \mathbf{K}'}}{1 - \gamma} \sigma_\epsilon^2}_{\text{Var}_{X_\epsilon}}. \quad (34)$$

Here, we have highlighted how the risk naturally splits into three terms given by respective bias and variance components resulting from the covariates $\mathbf{X}$ and the noise ϵ , as

outlined in appendix D [16]. We have also defined:

$$\gamma_{\Sigma, \Sigma'} \equiv \frac{\text{df}_{\Sigma, \Sigma'}^2 \text{df}_{\mathbf{K}}^2}{\text{df}_1 \tilde{\text{df}}_1}, \quad \gamma_{\Sigma, \Sigma', \mathbf{K}, \mathbf{K}'} \equiv \frac{\text{df}_{\Sigma, \Sigma'}^2 \text{df}_{\mathbf{K}, \mathbf{K}'}^2}{\text{df}_1 \tilde{\text{df}}_1}, \quad (35)$$

$$\begin{aligned} \text{df}_{\Sigma, \Sigma'}^2 &\equiv \frac{1}{N} \text{Tr} [\Sigma \Sigma' (\Sigma + \kappa)^{-2}], \\ \text{df}_{\mathbf{K}, \mathbf{K}'}^2 &\equiv \frac{1}{T} \text{Tr} [\mathbf{K} \mathbf{K}' (\mathbf{K} + \tilde{\kappa})^{-2}]. \end{aligned} \quad (36)$$

Similarly, for the training error, we have

$$\hat{R}_{in} \simeq \frac{\tilde{\text{df}}_1 - \tilde{\text{df}}_2}{S \tilde{\text{df}}_1} \frac{\kappa^2}{1 - \gamma} \bar{\mathbf{w}}^\top \Sigma (\Sigma + \kappa)^{-2} \bar{\mathbf{w}} + \sigma_\epsilon^2 \tilde{\kappa} \left[\frac{1}{T} \text{Tr} \mathbf{K}' (\mathbf{K} + \tilde{\kappa})^{-1} - \frac{\text{df}_1 - \text{df}_2}{1 - \gamma} \frac{\text{df}_{\mathbf{K}, \mathbf{K}'}^2}{\text{df}_1} \right]. \quad (37)$$

Proof. See appendix C.4. $\square$

We validate these formulae across several experiments (figure 4, appendix J). As a special case, setting $\mathbf{K} = \mathbf{I}_T$ recovers previously-obtained asymptotics for the risk of ridge regression under covariate shift [42–44]. In general, the training error is proportional to the generalization error only in the setting where $\Sigma = \Sigma'$ and $\mathbf{K} = \mathbf{K}'$. Thus, it is only in this setting that a GCV exists⁹.

4.4. Effect on double descent

It is interesting to consider how correlation structure affects double descent. From equations (15) and (16) one has in the structured case that:

$$\kappa = \frac{\lambda}{1 - q \text{df}_1} S_{\mathbf{K}}(q \text{df}_1); \quad (38)$$

compare with (24) in the uncorrelated setting. In the ridgeless limit in the overparameterized regime $q > 1$, κ will be such that $q \text{df}_{\Sigma}^1(\kappa) = 1$, which is unchanged from the uncorrelated setting (appendix F). The double descent behavior in that case is unchanged by correlations. At finite ridge, we see that training on correlated data is similar to replacing $\lambda \rightarrow \lambda S_{\mathbf{K}}(q \text{df}_1)$ with uncorrelated data. Near the double descent peak, this expands to leading order in $q - 1$ as $\lambda S_{\mathbf{K}}(1)$. In appendix F, we show that $S_{\mathbf{K}} \geq 1$ pointwise. Thus, correlations will enhance the ridge near the interpolation threshold, meaning that the double-descent peak should be mollified. For strong exponential correlations $K_{ts} = e^{-|t-s|/\xi}$, $\lambda S_{\mathbf{K}}(1)$ is approximately $\lambda \xi$. This leads to the less sharp double descent peak in figure 4(b). Thus, correlations mollify the double descent phenomenon. We do not say that they regularize the effect, because the risk still explodes in the ridgeless $q \rightarrow 1$ limit.

⁹ A full enumeration of model settings where the condition $\mathbf{K} = \mathbf{K}'$ holds is outside the scope of the present work. We note in passing that one case in which it does hold is if the true feature vector generating the labels is only partially-observed; this can generate an effective label noise with matched sample-wise correlations [11].

Risk and cross validation in ridge regression with correlated samples




Figure 4. Precise asymptotics for double descent in linear regression with unstructured data across various correlations. We choose an exponential correlation with correlation length ξ and vary ξ . (a) Weakly correlated data and noise, giving rise to the traditional double descent curve as analyzed in [11, 45]. All GCV-related estimators agree and correctly estimate the out-of-sample risk. (b) Strongly correlated data with matched noise correlations. The double descent peak is *mollified*. (c) Strongly correlated data but uncorrelated noise. The double descent peak is *exacerbated*. This mismatch in correlations violates the assumptions of the CorrGCV, and thus no GCV can asymptotically match it without knowledge of the noise level σ_ϵ . Across all settings the theory curves (solid lines) find excellent agreement with the experiments (solid markers with error bars over 10 different datasets).

Noise correlations affect the generalization error only if the covariates are also correlated. If the covariates are uncorrelated, namely $\mathbf{K} = \mathbf{I}_T$ then $\mathbf{K}'$ enters the risk (34) only through $\frac{1}{T}\text{Tr}\mathbf{K}'$, which is 1 by definition. Conversely, we can consider the case where the covariates are correlated but the noise is uncorrelated, i.e. $\mathbf{K}' = \mathbf{I}_T$. In the overparameterized regime, in the ridgeless limit this yields an error equal to that with uncorrelated data multiplied by $\frac{1}{T}\text{Tr}(\mathbf{K}^{-1})$, which is greater than one so long as $\mathbf{K} \neq \mathbf{I}_T$ (appendix F). This implies that having uncorrelated noise is generally worse than having noise with matched correlations. We illustrate this behavior in figure 4(c), where this effect is visible as a strong magnification of double-descent.

5. Algorithmic implementation

In this section, we give an algorithm to compute the CorrGCV given knowledge of the function $S_{\mathbf{K}}(q\text{df}_1)$. We discuss how to estimate $S_{\mathbf{K}}$ from data in appendix A.2. The diagram for this calculation is given in figure 5. In what follows, we will use the notation $b \leftarrow a$ to indicate variable assignment of b given that a has been computed. We also use $a \equiv b$ to highlight that the variable a is shorthand for b . This allows one to easily track the causal chain of how the relevant variables are estimated.

First, given $\hat{\Sigma}$ alone, we can calculate κ as long as the functional form of $S_{\mathbf{K}}$ is known. We obtain:

$$\text{df}_1 \leftarrow \text{df}_{\hat{\Sigma}}^1(\lambda), \quad \kappa \leftarrow \lambda \frac{S_{\mathbf{K}}(q\text{df}_1)}{1 - q\text{df}_1}. \quad (39)$$




Figure 5. A graphical representation of the program needed to obtain the CorrGCV empirically from a given dataset. The asymmetry of the diagram arises from the fact that we estimate κ first rather than $\tilde{\kappa}$. This is because it is more reasonable to assume a good estimate of the correlations $\mathbf{K}$, which often have properties such as stationarity that improve the estimation process, compared to estimating Σ . As a result, it is easier to use either an exact form or a differentiable interpolation of $S_{\mathbf{K}}(\text{df})$ rather than S_{Σ} .

At this point, the ordinary GCV is directly given by $S^2 \hat{R}_{\text{in}}$. Given a functional form of $S_{\mathbf{K}}$, this gives κ as a differentiable program of λ . Derivatives can be efficiently evaluated with autograd, for example using the JAX library [46]. This will be important for the next steps.

Second, by applying the duality relation (see appendix C), we obtain estimates for $\tilde{\text{df}}_1 \equiv \text{df}_{\mathbf{K}}^1(\lambda)$ and $\tilde{\kappa}$:

$$\tilde{\text{df}}_1 \leftarrow q\text{df}_1, \quad \tilde{\kappa} \leftarrow \frac{\lambda}{\kappa q\text{df}_1}. \quad (40)$$

Third, we leverage autograd to estimate $\text{df}_2 \equiv \text{df}_{\Sigma}^2$ by applying another duality relationship:

$$\text{df}_2 \leftarrow \text{df}_1 + \frac{\partial_{\lambda} \text{df}_1}{\partial_{\lambda} \log \kappa}. \quad (41)$$

Fourth, by using this and applying autograd again, we obtain an estimate for $\tilde{\text{df}}_2 \equiv \text{df}_{\mathbf{K}}^2$:

$$\tilde{\text{df}}_2 \leftarrow \tilde{\text{df}}_1 - q \frac{\partial_{\lambda} \log \kappa}{\partial_{\lambda} \log \tilde{\kappa}} (\text{df}_1 - \text{df}_2). \quad (42)$$

Finally, from this we obtain the CorrGCV estimator:

$$E_{\text{CorrGCV}} \leftarrow \hat{R}_{\text{in}} S \frac{\tilde{\text{df}}_1}{\tilde{\text{df}}_1 - \tilde{\text{df}}_2}. \quad (43)$$

6. Testing on correlated data

Finally, we consider the setting where the data point $\mathbf{x}$ on which we test on has a nontrivial correlation $\mathbf{k}$ with each $\mathbf{x}_t$. That is, $\mathbb{E}[\mathbf{x}_t \cdot \mathbf{x}] \propto k_t$. For simplicity, we assume that the covariates and noise are identically correlated. Such a setting arises naturally in the case of forecasting time series, where one trains a model on a window of data and then aims to predict at a future time within the correlation time of the process. This correlation introduces a multiplicative correction to the asymptotic risk:

Theorem 6.1. *Denote by $R_{\text{out}}^{\mathbf{k}}$ the out-of-sample risk when the test point has correlation $[\mathbf{k}]_t$ with data point $\mathbf{x}_t$. In this case, the test point $\mathbf{x}$ is conditionally Gaussian with*

$$\mathbb{E}[\mathbf{x}|\mathbf{X}] = \mathbf{X}^\top \boldsymbol{\alpha}, \quad \text{Var}[\mathbf{x}|\mathbf{X}] = (1 - \rho)\boldsymbol{\Sigma}, \quad (44)$$

where $\boldsymbol{\alpha} = \mathbf{K}^{-1}\mathbf{k}$ and $\rho = \mathbf{k}^\top \mathbf{K}^{-1}\mathbf{k}$. Assume that $\boldsymbol{\epsilon}_t$ has the same correlation as $\mathbf{x}_t$, namely $\mathbf{K}$. Then,

$$R_{\text{out}}^{\mathbf{k}} \simeq R_{\text{out}}^{\mathbf{k}=0} [1 - \rho + \tilde{\kappa}^2 \boldsymbol{\alpha}^\top \mathbf{K} (\mathbf{K} + \tilde{\kappa})^{-2} \boldsymbol{\alpha}]. \quad (45)$$

Proof. See appendix C.5. □

We note that because $\tilde{\kappa} > 0$, the spectrum of $\tilde{\kappa}^2(\mathbf{K} + \tilde{\kappa})^{-2}$ is strictly bounded between zero and one. Thus, the last term in brackets is strictly bounded from above by ρ , meaning that $R_{\text{out}}^{\mathbf{k}}$ is always less than $R_{\text{out}}^{\mathbf{k}=0}$. This precisely quantifies the over-optimism of cross-validation when the held-out samples are nearby in time.

The train error is not affected by the correlated test point. Consequently, this yields an extension of the CorrGCV:

$$R_{\text{out}} \simeq \frac{S\tilde{\text{df}}_1}{\tilde{\text{df}}_1 - \tilde{\text{df}}_2} [1 - \rho + \tilde{\kappa}^2 \boldsymbol{\alpha}^\top \mathbf{K} (\mathbf{K} + \tilde{\kappa})^{-2} \boldsymbol{\alpha}] \hat{R}_{\text{in}}. \quad (46)$$

In practice, we find that the simple approximation $R_{\text{out}}^{\mathbf{k}} \approx (1 - \rho)R_{\text{out}}^{\mathbf{k}=0}$ works very well (figure 6).

7. Conclusion

In this paper, we have provided a comprehensive characterization of the asymptotic risk of high-dimensional ridge regression with correlated Gaussian datapoints. We expect our asymptotics to extend to non-Gaussian data thanks to the universality of ridge regression [23, 24, 47]. Our results show that previously-proposed extensions of the GCV estimator to non-i.i.d. data are asymptotically biased, and immediately give the formula for the corrected GCV. This correction factor requires one to estimate more fine-grained spectral statistics of the sample-sample covariance than the usual GCV. However, as we have shown in the figures and discussed in appendix A, obtaining an excellent estimate appears relatively straightforward.

Risk and cross validation in ridge regression with correlated samples




Figure 6. Strongly exponentially correlated data as in figure 1(b), with $\mathbb{E}\mathbf{x}_t\mathbf{x}_{t+\tau} \sim e^{-\tau/\xi}$, $\xi = 10^2$. (a) Comparison of the out-of-sample risk for testing on a correlated point that is $\tau = 5$ time points in the future from the training data. We emphasize that the CorrGCV here is the CorrGCV for testing on an uncorrelated test point, which is pessimistic relative to the risk for the correlated point at $\tau = 5$. (b) Plot of the out-of-sample risk as a function of the horizon time that we test out. We see that testing on closer times is predictably more optimistic. We see that $(1 - \rho)$ times the CorrGCV provides a good approximation of the optimism of testing on a correlated point. Given a knowledge of the correlation structure of $\mathbf{K}$, this can be efficiently computed for stationary time series data.

Though our results are quite general, they are not without limitations, and there are several opportunities for further theoretical inquiry. First, our matrix-Gaussian model for the training data (3) could be relaxed to allow for sample-dependent covariance between features. On the technical side, to our knowledge Gaussian universality of ridge regression for non-independent data has not been rigorously established, though we expect it to hold [23, 24, 47]. One might also want to establish dimension-free deterministic equivalents with relative error bounds, as in Misiakiewicz and Saeed [47]. However, these are largely technical, rather than conceptual, limitations.

After the completion of this work, we became aware of the contemporaneous work of Luo *et al* [48], who derived an asymptotically unbiased non-omniscient risk estimator under the assumption of right rotation-invariance. Concretely, they assume that the design matrix $\mathbf{X}$ satisfies $\mathbf{X}\mathbf{O} \stackrel{d}{=} \mathbf{X}$ for any orthogonal matrix $\mathbf{O}$. Their results thus allow for non-Gaussian data, but do not allow for correlations between features, and are therefore complementary to our own.

We conclude by commenting briefly on applications of our results. First and foremost, we anticipate that our asymptotically exact correction to the GCV should be of some utility in time series regression settings, whether in finance [1], neuroscience [3], or elsewhere. Moreover, our results might be of use in the study of in-context learning in

language models, where most theoretical investigations neglect for the sake of analytical convenience the rich correlations present in language [49].

Acknowledgment

AA is grateful to Holden Leslie-Bole and Jacob Prince for useful conversations. We also thank Benjamin Ruben for helpful comments on a previous version of this manuscript. JAZV was supported by the Office of the Director of the National Institutes of Health under Award Number DP5OD037354. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health. JAZV is further supported by a Junior Fellowship from the Harvard Society of Fellows. CP is supported by NSF grant DMS-2134157, NSF CAREER Award IIS-2239780, DARPA grant DIAL-FP-038, a Sloan Research Fellowship, and The William F. Milton Fund from Harvard University. This work has been made possible in part by a gift from the Chan Zuckerberg Initiative Foundation to establish the Kempner Institute for the Study of Natural and Artificial Intelligence.

Author contributions

AA and JAZV conceived the project, performed research, and wrote the manuscript. CP supervised the project and edited the manuscript.

Appendix A. Algorithmic implementation of the CorrGCV

A.1. Code for implementation

In Code Block 7 we give the code for implementing the CorrGCV estimator in JAX given a ridge λ , training error $\hat{R}_{\text{in}}$, overparameterization ratio q , empirical covariance matrix $\hat{\Sigma}$, and function $S_{\mathbf{K}}(\text{df})$ given by the S -transform of the correlation structure.

A.2. Estimating the S -transform from data

If the functional form of the correlations is not exactly known, there are several options.

1. From empirical data, one can fit a parametric form of $\mathbf{K}$ for which the S -transform $S_{\mathbf{K}}$ can be calculated exactly. We give some examples of this in appendix I; a practically-relevant example is an exponential $e^{-|t-t'|/\xi}$ [32].
2. For a time series with the assumption of stationarity, one can estimate the autocorrelation $A_{\tau} = K_{t,t+\tau}$ as a function of lag τ much more reliably than the full matrix $K_{t,t'}$, as it depends on only $\mathcal{O}(T)$ parameters rather than $\mathcal{O}(T^2)$. This, in turn, can enable more reliable parametric fitting.

Alternatively, one can use this to form a more reliable direct estimate for $\mathbf{K}$. Given such an estimate, which we will call $\check{\mathbf{K}}$ to distinguish it from the plug-in estimate $\hat{\mathbf{K}}$, we can use a solver (such as a bisection method) to evaluate $\text{df}_{\check{\mathbf{K}}}^{-1}(t)$,

```

1 # R_in: in-sample risk
2 # lamb: ridge parameter
3 # q: N/T
4 # eigs_Sh: eigenvalues of empirical covariance Sigma hat
5 # S_K: sample-side S-transform S_K(df)
6
7 from jax import grad
8 import jax.numpy as jnp
9
10 df1_fn_Sh = lambda l: jnp.mean(eigs_Sh / (1 + eigs_Sh))
11 d_df1_fn_Sh = grad(df1_fn_Sh)
12
13 df1_est = df1_fn_Sh(lamb)
14 df1_Kh_est = q * df1_est
15
16 # Naive GCV ignoring sample correlations
17 R_GCV1 = R_in / (1 - q * df1_est) ** 2
18
19 # Full S-transform factor for the empirical covariance
20 S_full_fn = lambda l: S_K(q * df1_fn_Sh(l)) / (1 - q * df1_fn_Sh(l))
21 S_est = S_full_fn(lamb)
22
23 # Altman / naive correlated-noise correction
24 R_GCV2 = R_in * S_est ** 2
25
26 # Renormalized ridges
27 kappa1_fn = lambda l: 1 * S_full_fn(l)
28 kappa2_fn = lambda l: 1 / (q * df1_fn_Sh(l) * kappa1_fn(l))
29
30 d_log_kappa1_fn = grad(lambda l: jnp.log(kappa1_fn(l)))
31 d_log_kappa2_fn = grad(lambda l: jnp.log(kappa2_fn(l)))
32
33 # Estimate df_2 for Sigma
34 df2_fn_Sh = lambda l: df1_fn_Sh(l) + d_df1_fn_Sh(l) / d_log_kappa1_fn(l)
35 df2_Sh_est = df2_fn_Sh(lamb)
36
37 # Estimate df_2 for K
38 df2_Kh_est = df1_Kh_est - q * (df1_est - df2_Sh_est) * (
39     d_log_kappa1_fn(lamb) / d_log_kappa2_fn(lamb)
40 )
41
42 R_CorrGCV = R_in * S_est * df1_Kh_est / (df1_Kh_est - df2_Kh_est)

```

Figure 7. Python code for computing CorrGCV estimator given relevant parameters, assuming S_K is known.

and hence estimate

$$S_K(t) \approx S_{\tilde{K}}(t) = \frac{1-t}{t \, df_{\tilde{K}}^{-1}(t)}. \quad (\text{A1})$$

This is the procedure we use for the power-law correlation experiments, where we know $\mathbf{K}$ exactly but cannot evaluate S_K analytically.

We can then evaluate κ_1 as follows:

3. • First, sample a fine grid of λ .
 - From the data, we estimate $df_1 = df_{\Sigma}^1(\lambda)$ for each lambda, leading to a fine grid of df_1 and $\tilde{df}_1 = qdf_1$
 - Given the bisection-based algorithm to evaluate $df_{\tilde{K}}^{-1}$. We can then evaluate $S_{\tilde{K}}(\tilde{df})$ on the corresponding fine grid of $\tilde{df}_1 = qdf_1$ using the approximation stated above, which gives us an estimate of $S_K(\tilde{df})$
 - Using the definition $\kappa = \frac{S_K(qdf_1)}{1-qdf_1}$, we obtain estimates of κ as a function of λ .
 - We now apply an appropriate interpolation method that we can differentiate through and guarantee that as the grid gets finer, both the function and its derivatives converge to the true κ and $\partial_{\lambda}\kappa$.

The last bullet point requires further discussion. In practice, because κ can vary over several orders of magnitude, and because estimating the CorrGCV

requires logarithmic derivatives, it is often best to interpolate $\log \kappa$ against $\log \lambda$. A good interpolator compatible with autograd can be a Gaussian process with appropriately chosen covariance, a piece-wise smooth spline function, polynomial interpolator, or a wavelet based method that converges in the appropriate Sobolev norm. We used a simple degree 5 polynomial interpolation of $S_{\mathbf{K}}$ to generate the power law correlation learning curves in figure 3(b).

4. If one instead tries to estimate $S_{\mathbf{K}}$ from the empirical Gram matrix $\hat{\mathbf{K}} = \frac{1}{T} \mathbf{X} \mathbf{X}^\top$, this matrix contains information both about the sample correlations $\mathbf{K}$ and the feature correlations Σ . For general anisotropic Σ , we must therefore deconvolve S_{Σ} :

$$S_{\mathbf{K}}(t) = (q - t) \frac{S_{\hat{\mathbf{K}}}(t)}{S_{\Sigma}(t/q)}. \quad (\text{A2})$$

If we can do so (for instance, if we know that we have isotropic features such that $S_{\Sigma} = 1$), then we can proceed as before.

Appendix B. Review of free probability

B.1. Definition of freedom

We now define what it means for a set of n random matrices $\{\mathbf{A}_i\}_{i=1}^n$ to be jointly (asymptotically) free. All of these matrices are $N \times N$ and we consider the limit of $N \rightarrow \infty$. This is the limit in which free probability theory applies for random matrices. There are several texts on this rich subject. See, for example [50–52].

We say that a polynomial $p(\mathbf{A}_i)$ has the mean zero property if

$$\frac{1}{N} \text{Tr}[p(\mathbf{A}_i)] \simeq 0. \quad (\text{B1})$$

Above, $\simeq 0$ means that this converges to zero in probability as $N \rightarrow \infty$.

Take a set of m polynomials $\{p_k\}_{k=1}^m$, each with the mean zero property. Further, take a labeling $\{i_k\}_{k=1}^m$ with each $i_k \in \{1, \dots, n\}$ so that $i_k \neq i_{k+1}$ for all k . The $\{\mathbf{A}_i\}_{i=1}^n$ are jointly asymptotically free if and only if

$$\frac{1}{N} \text{Tr}[p_1(\mathbf{A}_{i_1}) \cdots p_m(\mathbf{A}_{i_m})] \simeq 0 \quad (\text{B2})$$

for any m , any labeling $\{i_k\}$, and any set of polynomials $\{p_k\}_{k=1}^m$ satisfying the properties above.

For our purposes we only require that the $\frac{1}{T} \mathbf{Z}^\top \mathbf{Z}$ and $\frac{1}{T} \mathbf{Z} \mathbf{Z}^\top$ matrices are free of any fixed deterministic matrix, specifically Σ and $\mathbf{K}$.

B.2. The R and S transforms

We review here basic properties of the R and S transforms of free probability [51, 53]. The recent book by Potters and Bouchaud [32] provides an accessible introduction to these techniques in the context of random matrix theory.

Given a $N \times N$ symmetric random matrix $\mathbf{A}$, one considers the **resolvent**, also known as the **Stieltjes Transform** of the eigenvalue distribution:

$$g_{\mathbf{A}}(\lambda) \equiv \frac{1}{N} \text{Tr} [(\mathbf{A} + \lambda)^{-1}]. \quad (\text{B3})$$

The eigenvalue density of $\mathbf{A}$, $\rho_{\mathbf{A}}$, can be obtained from the **inverse Stieltjes transform**:

$$\rho_{\mathbf{A}}(\lambda) = \lim_{\epsilon \rightarrow 0^+} \frac{1}{\pi} \text{Im} g_{\mathbf{A}}(-\lambda - i\epsilon). \quad (\text{B4})$$

Similarly, one defines the degrees of freedom $\text{df}_{\mathbf{A}}^1(\lambda)$ as

$$\text{df}_{\mathbf{A}}^1(\lambda) = \frac{1}{N} \text{Tr} [\mathbf{A}(\mathbf{A} + \lambda)^{-1}], \quad \text{df}_{\mathbf{A}}^2(\lambda) = \frac{1}{N} \text{Tr} [\mathbf{A}^2(\mathbf{A} + \lambda)^{-2}]. \quad (\text{B5})$$

One has the relationship:

$$g_{\mathbf{A}}(\lambda) = \frac{1 - \text{df}_{\mathbf{A}}^1(\lambda)}{\lambda}. \quad (\text{B6})$$

We let $g_{\mathbf{A}}^{-1}, \text{df}_{\mathbf{A}}^{-1}$ denote the respective functional inverses. The R and S transforms are respectively functions of formal variables g, df given by:

$$R_{\mathbf{A}}(g) \equiv g_{\mathbf{A}}^{-1}(g) - \frac{1}{g}, \quad S_{\mathbf{A}}(\text{df}) \equiv \frac{1 - \text{df}}{\text{df} \text{df}_{\mathbf{A}}^{-1}(\text{df})}. \quad (\text{B7})$$

They have the property that for any two matrices $\mathbf{A}, \mathbf{B}$ that are free of one another:

$$R_{\mathbf{A}+\mathbf{B}}(g) = R_{\mathbf{A}}(g) + R_{\mathbf{B}}(g), \quad S_{\mathbf{A}*\mathbf{B}}(\text{df}) = S_{\mathbf{A}}(\text{df})S_{\mathbf{B}}(\text{df}). \quad (\text{B8})$$

From these two respective properties, together with the definitions (B7) we obtain the **subordination relations**, also known as (weak) deterministic equivalence:

$$g_{\mathbf{A}+\mathbf{B}}(\lambda) = g_{\mathbf{A}}(\lambda + R_{\mathbf{B}}), \quad \text{df}_{\mathbf{A}*\mathbf{B}}(\lambda) = \text{df}_{\mathbf{A}}(\lambda S_{\mathbf{B}}). \quad (\text{B9})$$

Here, the additive and multiplicative renormalizations of λ , given by $R_{\mathbf{B}}$ and $S_{\mathbf{B}}$ respectively, can each be evaluated in two different ways. The first way is from the original noisy matrices $\mathbf{A} + \mathbf{B}$ and $\mathbf{A} * \mathbf{B}$. This is the analogue of the empirical estimate of S , κ discussed in the text.,

$$R_{\mathbf{B}} \equiv R_{\mathbf{B}}(g_{\mathbf{A}+\mathbf{B}}(\lambda)), \quad S_{\mathbf{B}} \equiv S_{\mathbf{B}}(\text{df}_{\mathbf{A}*\mathbf{B}}^1(\lambda)). \quad (\text{B10})$$

The second way is from the clean matrix $\mathbf{A}$ itself. This is what gives the omniscient estimate of the renormalized ridges. Writing κ_+, κ_* for $\lambda + R_{\mathbf{B}}, \lambda S_{\mathbf{B}}$ respectively, we obtain the self-consistent equations:

$$\kappa_+ = \lambda + R_{\mathbf{B}}(g_{\mathbf{A}}(\kappa_+)), \quad \kappa_* = \lambda S_{\mathbf{B}}(\text{df}_{\mathbf{A}}^1(\kappa_*)). \quad (\text{B11})$$

B.3. Strong deterministic equivalence

The deterministic equivalences in (B9) extend to the matrices themselves. Taking $\mathbf{A}$ deterministic and $\mathbf{B}$ random and free of $\mathbf{A}$, we have:

$$(\mathbf{A} + \mathbf{B} + \lambda)^{-1} \simeq (\mathbf{A} + \kappa_+)^{-1}, \quad \mathbf{A}^* \mathbf{B} (\mathbf{A}^* \mathbf{B} + \lambda)^{-1} \simeq \mathbf{A} (\mathbf{A} + \kappa_*)^{-1}. \quad (\text{B12})$$

Here, for two matrices, we use the relation $\simeq$ to denote that the traces of these quantities against any test matrix of bounded spectral norm will converge in probability to the same quantity as $N \rightarrow \infty$.

The above two formulas can be derived using replica theory [32, 37], from diagrammatics [16, 54], or from cavity arguments [36].

In this paper, we only require the properties of the S -transform. The above deterministic equivalences are called ‘one point’ equivalences, as they only involve a single matrix inverse. We will derive two-point equivalents in the sequel.

Appendix C. Deferred proofs of main results

C.1. Warm-up: 1-point deterministic equivalents


In this section, we prove lemma 3.1 using a diagrammatic argument based on Atanasov *et al* [16], specialized to the specific case of Wishart noise matrices $\mathbf{B}$. The argument is diagrammatic, appealing to the fact that at large N , free random matrices obey non-crossing properties in their diagrammatics [55]. The argument holds rigorously only in the $N \rightarrow \infty$ limit. However, it can rigorously be shown to be the leading term in an asymptotic series in $1/N$ by leveraging known properties of the higher-order corrections [56]. Higher-order terms in this series could also, in principle, be computed.

In what follows, for the sake of notational levity we absorb a factor of $1/\sqrt{T}$ into each insertion of $\mathbf{Z}$ or $\mathbf{Z}^\top$. Because each pair of $\mathbf{Z}, \mathbf{Z}^\top$ insertions will necessarily be averaged over via Wick contraction, this is the same as associating an factor of $1/T$ with each Wick contraction.

$$\begin{aligned} \mathbb{E}_{\mathbf{Z}} \hat{\Sigma} (\hat{\Sigma} + \lambda)^{-1} &= (-\lambda)^{-1} \text{---} \Sigma^{1/2} \mathbf{Z}^\top \text{---} K \text{---} \mathbf{Z} \text{---} \Sigma^{1/2} \\ &+ (-\lambda)^{-2} \text{---} \Sigma^{1/2} \mathbf{Z}^\top \text{---} K \text{---} \mathbf{Z} \text{---} \Sigma \text{---} \mathbf{Z}^\top \text{---} K \text{---} \mathbf{Z} \text{---} \Sigma^{1/2} \\ &+ (-\lambda)^{-2} \text{---} \Sigma^{1/2} \mathbf{Z}^\top \text{---} K \text{---} \mathbf{Z} \text{---} \Sigma \text{---} \mathbf{Z}^\top \text{---} K \text{---} \mathbf{Z} \text{---} \Sigma^{1/2} + \dots \end{aligned} \quad (\text{C1})$$

Such a suppression of crossing diagrams is a defining characteristic of free probability. Because each loop contributes a trace, while each insertion of $\mathbf{Z}^\top \mathbf{Z}$ contributes a factor

of $1/T$, in order for a diagram to give an order 1 contribution we will need to have as many loops as there are pairs of $\mathbf{Z}^\top, \mathbf{Z}$. Consequently, crossing diagrams such as the following are suppressed as $N, T \rightarrow \infty$:



Because crossing diagrams do not contribute, one can observe the following pattern. Any diagram that appears will be a link of averages from one $\mathbf{Z}^\top$ to some later $\mathbf{Z}$ that creates an arc. Beneath that arc, all averages can only be between the matrices within the arc, by non-crossing. As such, we can expand:

$$\begin{aligned} \mathbb{E} \hat{\Sigma}(\hat{\Sigma} + \lambda)^{-1} &= (-\lambda)^{-1} \text{---} \overset{\text{arc}}{\Sigma^{1/2} \text{---} \Sigma^{1/2}} \text{---} \\ &+ (-\lambda)^{-2} \text{---} \overset{\text{arc}}{\Sigma^{1/2} \text{---} \Sigma^{1/2}} \text{---} \Sigma \text{---} \overset{\text{arc}}{\Sigma^{1/2} \text{---} \Sigma^{1/2}} \text{---} + \dots \quad (\text{C2}) \end{aligned}$$

The terms in the dashed lines are pre-emptively denoted by $1/S$. We make two observations.

- Because the isotropic matrix $\mathbf{Z}$ is right-invariant to rotations in $\mathbb{R}^N$, $1/S$ must be a deterministic matrix that is invariant under rotation. The only such matrices are constants times the identity.
- Because the isotropic matrix $\mathbf{Z}$ is also left-invariant to rotations in $\mathbb{R}^T$, $1/S$ is a rotationally invariant scalar functional of the product of matrices beneath the arc. The only such rotationally invariant scalar functional is any constant multiple of the trace.

The fact that the shaded parts are scalars immediately implies we can resume this as:

$$\mathbb{E}_{\mathbf{Z}} \hat{\Sigma}(\hat{\Sigma} + \lambda)^{-1} = \sum_{n=1}^{\infty} \Sigma^n \frac{1}{(-\lambda S)^n} = \Sigma(\Sigma + \lambda S)^{-1}. \quad (\text{C3})$$

An immediate consequence is:

$$\mathbb{E}_{\mathbf{Z}}(\hat{\Sigma} + \lambda)^{-1} = \frac{1}{\lambda} \left[1 - \mathbb{E}_{\mathbf{Z}} \hat{\Sigma}(\hat{\Sigma} + \lambda)^{-1} \right] = S(\Sigma + \lambda S)^{-1}. \quad (\text{C4})$$

$$\frac{1}{S} = \text{Diagram 1} + \text{Diagram 2} + \dots \quad (\text{C5})$$

$$\frac{1}{S} = \mathbb{E}_{\mathbf{Z}} \lambda \frac{1}{T} \text{Tr} \left[\mathbf{K}(\hat{\mathbf{K}} + \lambda)^{-1} \right]. \quad (\text{C6})$$

By a slight adjustment of the above diagrammatic argument, swapping the role of Σ and $\mathbf{K}$, we have an analogue of (C4):

We now define

Note that at this stage we have shown the equivalence:

Putting (C7) back into (C6) yields:

We recognize as the **duality relation**. We can also get an explicit expression for S in terms of S_K and df by defining the formal function of df

and plugging in $\mathbf{A} = \mathbf{K}$, $\text{df} = \text{df}_K^1(\tilde{\kappa}) = \tilde{\text{df}}_1$ into this:

Using $\tilde{\text{df}}_1 = q\text{df}_1$ this yields the desired relationship:

$$\kappa = \lambda \frac{S_K(q\text{df})}{1 - q\text{df}}. \quad (\text{C13})$$

Thus, given knowledge of λ, df_1 and S_K , one can calculate κ exactly. For a discussion of how to obtain S_K from a given correlated dataset, see A.2.

C.2. 2-point deterministic equivalents


In this section, we extend this diagrammatic technique to calculate the necessary 2-point deterministic equivalents stated in lemmas 3.2 and 3.3. Some of these were derived using leave-one-out (cavity) arguments in recent papers of Bach [36, 38] in the case of $K = I$. The remaining results that we derive are to our knowledge novel. As we discuss in Atanasov *et al* [57], they have useful applications to high-dimensional regression problems beyond just those studied in the context of this paper.

As in the prior section, we absorb a factor of $1/\sqrt{T}$ into each Z or $Z^\top$ insertion, or equivalently associate a factor of $1/T$ with each Wick contraction.

We first prove lemma 3.2, which concerns the evaluation of

$$(\hat{\Sigma} + \lambda)^{-1} \Sigma' (\hat{\Sigma} + \lambda)^{-1}, \quad (\text{C14})$$

for arbitrary matrix Σ' between the resolvents. In performing this average, we note that the types of diagrams that appear split into two classes. The first are those with no arcs over Σ' :




$$(\text{C15})$$

Because no arcs connect the resolvent on the left with the resolvent on the right, we can take the averages of the left and right resolvent separately, and obtain:

$$S^2(\Sigma + \kappa)^{-1} \Sigma' (\Sigma + \kappa)^{-1}. \quad (\text{C16})$$


The second class of term has arcs connecting the two resolvents on both sides. An example of such a diagram is:



$$(\text{C17})$$

Any term in the second class will have some number of arcs over Σ' . Note that this will necessarily be an even number $2n$, as they will alternate between averaging a $Z^\top$ on the left with a Z on the right and vice versa. This will give $2n$ loops that are traced over. There will be n loops involving K matrices, which we will call K loops, and n

loops involving Σ matrices which we will call Σ loops. The above term has $n = 1$. An example of a Σ loop is:



$$(C18)$$

Here, on each side we denote by $\Sigma + \dots$ the series:

$$\Sigma + (-\lambda)^{-1} \Sigma Z^T Z \Sigma + (-\lambda)^{-2} \Sigma Z^T Z \Sigma Z^T Z \Sigma + \dots = \lambda \Sigma^{1/2} (\hat{\Sigma} + \lambda)^{-1} \Sigma^{1/2}. \quad (C19)$$

We recognize this as a resolvent. Because we have explicitly accounted for arcs connecting terms in the left resolvent with the terms on the right resolvent, within each loop we can treat the left and right resolvent average separately. This average is given in (C4), and similarly for $\mathbf{K}$. This yields

$$\Sigma\text{-loop} = \lambda^2 \frac{1}{T} \text{Tr} \left[\Sigma^{1/2} \mathbb{E}_Z \left[(\hat{\Sigma} + \lambda)^{-1} \right] \Sigma \mathbb{E}_Z \left[(\hat{\Sigma} + \lambda)^{-1} \right] \Sigma^{1/2} \right] = q \kappa^2 \text{df}_2, \quad (C20)$$

$$\mathbf{K}\text{-loop} = \lambda^2 \frac{1}{T} \text{Tr} \left[\mathbf{K}^{1/2} \mathbb{E}_Z \left[(\hat{\mathbf{K}} + \lambda)^{-1} \right] \mathbf{K} \mathbb{E}_Z \left[(\hat{\mathbf{K}} + \lambda)^{-1} \right] \mathbf{K}^{1/2} \right] = \tilde{\kappa}^2 \tilde{\text{df}}_2. \quad (C21)$$

The innermost loop will have an insertion of Σ' in between the two resolvents, each of which is separately averaged, yielding:

$$\frac{S^2}{T} \text{Tr} [\Sigma^{1/2} (\Sigma + \kappa)^{-1} \Sigma' (\Sigma + \kappa)^{-1} \Sigma^{1/2}] \simeq S^2 \frac{1}{T} \text{Tr} [\Sigma \Sigma' (\Sigma + \kappa)^{-2}]. \quad (C22)$$

We adopt the shorthand

$$\begin{aligned} \text{df}_{\Sigma, \Sigma'}^2 &\equiv \frac{1}{N} \text{Tr} [\Sigma \Sigma' (\Sigma + \kappa)^{-2}], \\ \text{df}_{\mathbf{K}, \mathbf{K}'}^2 &\equiv \frac{1}{T} \text{Tr} [\mathbf{K} \mathbf{K}' (\mathbf{K} + \tilde{\kappa})^{-2}], \end{aligned} \quad (C23)$$

to denote this term by $q S^2 \text{df}_{\Sigma, \Sigma'}^2$. This must necessarily be followed by a $\mathbf{K}$ loop. By making use of the duality relationship, we can write this joint contribution as:

$$\frac{\kappa^2}{\lambda^2} q \text{df}_{\Sigma, \Sigma'}^2 \tilde{\kappa}^2 \tilde{\text{df}}_2 = \frac{\text{df}_{\Sigma, \Sigma'}^2 \tilde{\text{df}}_2}{\text{df}_1 \tilde{\text{df}}_1} \equiv \gamma_{\Sigma, \Sigma'}. \quad (C24)$$

Between this innermost Σ' and $\mathbf{K}$ loop and the outside, there can be an arbitrary number of pairs of closed $\mathbf{K}$ and Σ loops in between. Again applying the duality relation,

we see each pair contributes:

$$\frac{1}{\lambda^2} q \kappa^2 \text{df}_2 \tilde{\kappa}^2 \tilde{\text{df}}_2 = \frac{\text{df}_2 \tilde{\text{df}}_2}{\text{df}_1 \tilde{\text{df}}_1} \equiv \gamma. \quad (\text{C25})$$

Here we divide by λ^2 because each $\mathbf{Z}\mathbf{Z}^\top$ pair introduces a factor of $(-\lambda)^{-1}$. This gives an interpretation of γ as the contribution of a pair of Σ and $\mathbf{K}$ loops. It remains to sum over n to get the final contribution from all of the loops:

$$\gamma_{\Sigma, \Sigma'} \sum_{n=0}^{\infty} \gamma^n = \frac{\gamma_{\Sigma, \Sigma'}}{1 - \gamma}. \quad (\text{C26})$$

Finally, outside of the loops, we can perform the average of the left resolvent and the right resolvent separately. This gives the desired relation:

$$(\hat{\Sigma} + \lambda)^{-1} \Sigma' (\hat{\Sigma} + \lambda)^{-1} \simeq S^2 (\Sigma + \kappa)^{-1} \Sigma' (\Sigma + \kappa)^{-1} + S^2 (\Sigma + \kappa)^{-2} \Sigma \frac{\gamma_{\Sigma, \Sigma'}}{1 - \gamma}. \quad (\text{C27})$$

When $\mathbf{K} = \mathbf{I}$, this recovers the earlier result of [36]. By applying the same argument with minimal modifications, one obtains the deterministic equivalence for kernel resolvents, namely:

$$(\hat{\mathbf{K}} + \lambda)^{-1} \mathbf{K}' (\hat{\mathbf{K}} + \lambda)^{-1} \simeq \tilde{S}^2 (\mathbf{K} + \tilde{\kappa})^{-1} \mathbf{K}' (\mathbf{K} + \tilde{\kappa})^{-1} + \tilde{S}^2 (\mathbf{K} + \tilde{\kappa})^{-2} \mathbf{K} \frac{\gamma_{\mathbf{K}, \mathbf{K}'}}{1 - \gamma}. \quad (\text{C28})$$


Here $\gamma_{\mathbf{K}\mathbf{K}} \equiv \frac{\text{df}_\Sigma^2 \text{df}_{\mathbf{K}, \mathbf{K}'}}{\text{df}_1 \text{df}_1}$.

We next prove lemma 3.3, which gives an equivalent for

$$(\hat{\Sigma} + \lambda)^{-1} \frac{1}{T} \mathbf{X}^\top \mathbf{K}' \mathbf{X} (\hat{\Sigma} + \lambda)^{-1} = (\hat{\Sigma} + \lambda)^{-1} \Sigma^{1/2} \mathbf{Z}^\top \mathbf{K}^{1/2} \mathbf{K}' \mathbf{K}^{1/2} \mathbf{Z} \Sigma^{1/2} (\hat{\Sigma} + \lambda)^{-1}. \quad (\text{C29})$$

This time, because there are an odd number of $\mathbf{Z}$ s on both the left and the right side of $\mathbf{K}'$ in all diagrams, the disconnected terms which average each side separately will vanish. We are left with just the connected term.

Again, we sum over the $\mathbf{K}$ and Σ loops. The first term (which involves no Σ loops) is:



$$\Sigma^{1/2} \mathbf{Z}^\top \mathbf{K}^{1/2} \mathbf{K}' \mathbf{K}^{1/2} \mathbf{Z} \Sigma^{1/2} \quad (\text{C30})$$

The innermost loop will be a $\mathbf{K}$ loop involving an insertion of $\mathbf{K}'$ in the middle. As in the calculation of the Σ' loop in the prior case, this contributes

$$\tilde{S}^2 \frac{1}{T} \text{Tr}[\mathbf{K} \mathbf{K}' (\mathbf{K} + \tilde{\kappa})^{-2}] \equiv \tilde{S}^2 \text{df}_{\mathbf{K}, \mathbf{K}'}. \quad (\text{C31})$$

Between this innermost loop and the outer resolvents, there must be an even number $2n \geq 0$ of alternating $\Sigma, \mathbf{K}$ loops. The contributions of such loops are unchanged. Each pair gives a factor of γ . We thus get:

$$(\hat{\Sigma} + \lambda)^{-1} \frac{1}{T} \mathbf{X}^\top \mathbf{K}' \mathbf{X} (\hat{\Sigma} + \lambda)^{-1} \simeq \lambda^2 S^2 \tilde{S}^2 (\Sigma + \kappa)^{-2} \Sigma \frac{\text{df}_{\mathbf{K}, \mathbf{K}'}^2}{1 - \gamma}. \quad (\text{C32})$$

Contracting this with an arbitrary matrix Σ' and multiplying by $\frac{1}{T}$ gives the more symmetric form:

$$\frac{1}{T^2} \text{Tr} \left[\Sigma' (\hat{\Sigma} + \lambda)^{-1} \mathbf{X}^\top \mathbf{K}' \mathbf{X} (\hat{\Sigma} + \lambda)^{-1} \right] \simeq \frac{\text{df}_{\Sigma, \Sigma'}^2 \text{df}_{\mathbf{K}, \mathbf{K}'}^2}{\text{df}_1 \tilde{\text{df}}_1} \frac{1}{1 - \gamma}. \quad (\text{C33})$$

We can further simplify this by adopting the shorthand $\gamma_{\Sigma, \Sigma', \mathbf{K}, \mathbf{K}'} = \frac{\text{df}_{\Sigma, \Sigma'}^2 \text{df}_{\mathbf{K}, \mathbf{K}'}^2}{\text{df}_1 \tilde{\text{df}}_1}$.

This result can alternatively be derived using a straightforward but tedious and significantly less conceptually illuminating argument based on a one-point deterministic equivalent for a ‘sourced’ resolvent. Concretely, this argument starts by writing

$$\frac{1}{T} (\hat{\Sigma} + \lambda)^{-1} \mathbf{X}^\top \mathbf{K}' \mathbf{X} (\hat{\Sigma} + \lambda)^{-1} = - \frac{\partial}{\partial J} [\Sigma^{1/2} \mathbf{Z}^\top (\mathbf{K} + J \mathbf{K}^{1/2} \mathbf{K}' \mathbf{K}^{1/2}) \mathbf{Z} \Sigma^{1/2} + \lambda]^{-1} \Big|_{J=0}, \quad (\text{C34})$$

and then proceeds by using a one-point deterministic equivalent before implicitly differentiating the resulting S -transforms. Then, various applications of the duality relation allow the result to be simplified into the form obtained using diagrammatics above.

C.3. Duality relations and derivatives

Here, we explicitly state all duality relationships that one can derive. In the section on 1-point deterministic equivalents, we have proven that:

$$q \text{df}_1 \equiv q \text{df}_{\Sigma}^1(\kappa) \simeq q \text{df}_{\hat{\Sigma}}^1(\lambda) = \text{df}_{\hat{\mathbf{K}}}^1(\lambda) \simeq \text{df}_{\mathbf{K}}^1(\tilde{\kappa}) \equiv \tilde{\text{df}}_1. \quad (\text{C35})$$

We have also seen the first duality relation, namely that:

$$\frac{\kappa \tilde{\kappa}}{\lambda} = \lambda S \tilde{S} = \frac{1}{q \text{df}_1} = \frac{1}{\tilde{\text{df}}_1}. \quad (\text{C36})$$

Logarithmic differentiation of this yields a second duality relation:

$$\frac{\text{d} \log \kappa}{\text{d} \log \lambda} + \frac{\text{d} \log \tilde{\kappa}}{\text{d} \log \lambda} = 1 + \frac{\text{df}_1 - \text{df}_{\hat{\Sigma}}^2(\lambda)}{\text{df}_1}. \quad (\text{C37})$$

Here we have used that

$$\frac{\text{d} \text{df}_1}{\text{d} \lambda} = \frac{\text{df}_{\hat{\Sigma}}^2(\lambda) - \text{df}_{\hat{\Sigma}}^1(\lambda)}{\lambda}. \quad (\text{C38})$$

We stress that unlike with df_1 , $\text{df}_\Sigma^2(\lambda) \neq \text{df}_\Sigma^2(\kappa) = \text{df}_2$. Since the right hand side can be estimated from the data, this gives us a way to turn an estimate of $\frac{\text{dlog } \kappa}{\text{dlog } \lambda}$ into an estimate for $\frac{\text{dlog } \tilde{\kappa}}{\text{dlog } \lambda}$ or vice-versa.

Next, we relate df_2 to $\tilde{\text{df}}_2$. We will be explicit and write all dfs with subscripts to avoid confusion. We have, by differentiating (C35):

$$q\kappa\partial_\kappa\text{df}_\Sigma^1(\kappa) = q\frac{\kappa}{\lambda}\frac{\text{d}\lambda}{\text{d}\kappa}\lambda\partial_\lambda\text{df}_\Sigma^1(\lambda) = \frac{\kappa}{\lambda}\frac{\text{d}\lambda}{\text{d}\kappa}\lambda\partial_\lambda\text{df}_K^1(\lambda) = \frac{\kappa}{\tilde{\kappa}}\frac{\text{d}\kappa}{\text{d}\tilde{\kappa}}\tilde{\kappa}\partial_{\tilde{\kappa}}\text{df}_K^1(\tilde{\kappa}). \quad (\text{C39})$$

Evaluating the left and right sides using (C38) we get:

$$q(\text{df}_\Sigma^2(\kappa) - \text{df}_\Sigma^1(\kappa)) = \frac{\text{dlog } \tilde{\kappa}}{\text{dlog } \kappa}(\text{df}_K^2(\tilde{\kappa}) - \text{df}_K^1(\tilde{\kappa})). \quad (\text{C40})$$

This yields:

$$\tilde{\text{df}}_2 \equiv \text{df}_K^2(\tilde{\kappa}) = q\text{df}_\Sigma^1(\kappa) + q\frac{\text{dlog } \kappa}{\text{dlog } \tilde{\kappa}}(\text{df}_\Sigma^2(\kappa) - \text{df}_\Sigma^1(\kappa)). \quad (\text{C41})$$

Similarly one can write an estimate of df_2 from just the data alone:

$$\text{df}_2 \equiv \text{df}_\Sigma^2(\kappa) = \text{df}_\Sigma^1(\lambda) + \frac{\text{dlog } \lambda}{\text{dlog } \kappa}(\text{df}_\Sigma^2(\lambda) - \text{df}_\Sigma^1(\lambda)) = \text{df}_\Sigma^1(\lambda) + \frac{\text{dlog } \lambda}{\text{dlog } \kappa}\lambda\partial_\lambda\text{df}_\Sigma^1(\lambda). \quad (\text{C42})$$

Plugging this back into (C41) gives an estimate of $\tilde{\text{df}}_2$ from the data alone. This is crucial in allowing the CorrGCV to be efficiently computed.

We now calculate the derivative of $\kappa, \tilde{\kappa}$ on λ . We have

$$\begin{aligned} \kappa = S\lambda &\Rightarrow \frac{\text{dlog } \lambda}{\text{dlog } \kappa} = 1 + \frac{\text{dlog } 1/S}{\text{dlog } \kappa} \\ &= 1 + \frac{\text{dlog } 1/S}{\text{dlog } \text{df}_1} \frac{\text{dlog } \text{df}_1}{\text{dlog } \kappa} \\ &= 1 - \frac{\text{dlog } 1/S}{\text{dlog } \text{df}_1} \frac{\text{df}_1 - \text{df}_2}{\text{df}_1}. \end{aligned} \quad (\text{C43})$$

Using $\tilde{\text{df}}_1 = q\text{df}_1$ we write

$$S = \frac{1}{1 - \tilde{\text{df}}_1} S_K(\tilde{\text{df}}_1) = \frac{1}{\tilde{\text{df}}_1 \text{df}_K^{-1}(\tilde{\text{df}}_1)}. \quad (\text{C44})$$

Differentiating this gives:

$$\frac{\text{dlog } 1/S}{\text{dlog } \tilde{\text{df}}_1} = 1 + \frac{\tilde{\text{df}}_1}{\tilde{\kappa}\text{df}_1'(\tilde{\kappa})} = 1 - \frac{\tilde{\text{df}}_1}{\tilde{\text{df}}_1 - \tilde{\text{df}}_2} = -\frac{\tilde{\text{df}}_2}{\tilde{\text{df}}_1 - \tilde{\text{df}}_2}. \quad (\text{C45})$$

All together this is:

$$\begin{aligned}\frac{d \log \lambda}{d \log \kappa} &= 1 + \frac{\tilde{d}f_2}{\tilde{d}f_1 - \tilde{d}f_2} \frac{df_1 - df_2}{df_1} \\ &= \frac{df_1 \tilde{d}f_1 - df_2 \tilde{d}f_2}{df_1 \tilde{d}f_1 - df_1 \tilde{d}f_2} \\ &= \frac{1 - \gamma}{1 - \frac{\tilde{d}f_2}{\tilde{d}f_1}}.\end{aligned}\tag{C46}$$

This finally yields:

$$\frac{\partial \kappa}{\partial \lambda} = S \frac{1 - \frac{\tilde{d}f_2}{\tilde{d}f_1}}{1 - \gamma}.\tag{C47}$$

An analogous argument yields:

$$\frac{\partial \tilde{\kappa}}{\partial \lambda} = \tilde{S} \frac{1 - \frac{df_2}{df_1}}{1 - \gamma}.\tag{C48}$$

Note also that when $\mathbf{K} = \mathbf{I}$, $\tilde{d}f_2 = (\tilde{d}f_1)^2 = q^2 df_1^2$, yielding

$$\gamma \equiv \frac{df_2 \tilde{d}f_2}{df_1 \tilde{d}f_1} = q df_2.\tag{C49}$$

This recovers the uncorrelated ridge regression setting.

C.4. Correlated samples and noise, uncorrelated test point

By applying the deterministic equivalences proved in the prior section, we can directly obtain an exact formula for the asymptotic form of the training and generalization error for a linear model trained on a correlated dataset and evaluated on an uncorrelated test point. In fact, we can do better. The deterministic equivalence (C27) allows us to easily consider the case where the test point has a different covariance Σ' from the covariance Σ of the training set. This allows us to state the formula for generalization under covariate shift as well. In one fell swoop, we thus prove theorems 4.1–4.3 by directly proving theorem 4.3, from which the other two results follow as special cases.

We recall that

$$\hat{\mathbf{w}} = \hat{\Sigma}(\hat{\Sigma} + \lambda)^{-1} \bar{\mathbf{w}} + (\hat{\Sigma} + \lambda)^{-1} \frac{\mathbf{X}^\top \boldsymbol{\epsilon}}{T}, \quad \mathbb{E}_{\boldsymbol{\epsilon}}[\boldsymbol{\epsilon} \boldsymbol{\epsilon}^\top] = \sigma_{\boldsymbol{\epsilon}}^2 \mathbf{K}'.\tag{C50}$$

We have that the generalization error (for a test point whose distribution has covariance Σ') is:

$$\begin{aligned}R_g &= (\bar{\mathbf{w}} - \hat{\mathbf{w}})^\top \Sigma' (\bar{\mathbf{w}} - \hat{\mathbf{w}}) \\ &= \underbrace{\lambda^2 \bar{\mathbf{w}}^\top (\hat{\Sigma} + \lambda)^{-1} \Sigma' (\hat{\Sigma} + \lambda)^{-1} \bar{\mathbf{w}}}_{\text{Signal}} + \underbrace{\sigma_{\boldsymbol{\epsilon}}^2 \frac{1}{T^2} \text{Tr}[\Sigma' (\hat{\Sigma} + \lambda)^{-1} \mathbf{X}^\top \mathbf{K}' \mathbf{X} (\hat{\Sigma} + \lambda)^{-1}]}_{\text{Noise}}.\end{aligned}\tag{C51}$$

The Signal term is immediately obtained by applying the deterministic equivalence in (C27),

$$\text{Signal} = \underbrace{\kappa^2 \bar{\mathbf{w}}^\top (\boldsymbol{\Sigma} + \kappa)^{-1} \boldsymbol{\Sigma}' (\boldsymbol{\Sigma} + \kappa)^{-1} \bar{\mathbf{w}}}_{\text{Bias}^2} + \underbrace{\kappa^2 \bar{\mathbf{w}}^\top \boldsymbol{\Sigma} (\boldsymbol{\Sigma} + \kappa)^{-2} \bar{\mathbf{w}} \frac{\gamma_{\boldsymbol{\Sigma}, \boldsymbol{\Sigma}'}}{1 - \gamma}}_{\text{Var}_{\mathbf{X}}}. \quad (\text{C52})$$

Here we have explicitly delineated which parts of the signal term are due to the bias of the estimator, and which terms are $\text{Var}_{\mathbf{X}}$. The latter can be removed by bagging the estimator over different datasets. Note that in diagrammatic language the bias term corresponds exactly to the disconnected averages of the left and right resolvents separately, which makes sense, as it is the generalization obtained by first averaging the predictor over different training sets before calculating the test risk.

Similarly, by applying the deterministic equivalence (C33), we obtain:

$$\text{Noise} = \underbrace{\sigma_\epsilon^2 \frac{\gamma_{\boldsymbol{\Sigma}, \boldsymbol{\Sigma}', \mathbf{K}, \mathbf{K}'}}{1 - \gamma}}_{\text{Var}_{\mathbf{X}_\epsilon}}. \quad (\text{C53})$$

This yields

$$R_g \simeq \underbrace{\kappa^2 \bar{\mathbf{w}}^\top (\boldsymbol{\Sigma} + \kappa)^{-1} \boldsymbol{\Sigma}' (\boldsymbol{\Sigma} + \kappa)^{-1} \bar{\mathbf{w}}}_{\text{Bias}^2} + \underbrace{\kappa^2 \frac{\gamma_{\boldsymbol{\Sigma}, \boldsymbol{\Sigma}'}}{1 - \gamma} \bar{\mathbf{w}}^\top \boldsymbol{\Sigma} (\boldsymbol{\Sigma} + \kappa)^{-2} \bar{\mathbf{w}}}_{\text{Var}_{\mathbf{X}}} + \underbrace{\frac{\gamma_{\boldsymbol{\Sigma}, \boldsymbol{\Sigma}', \mathbf{K}, \mathbf{K}'}}{1 - \gamma} \sigma_\epsilon^2}_{\text{Var}_{\mathbf{X}_\epsilon}}. \quad (\text{C54})$$

Specializing to the case of $\boldsymbol{\Sigma} = \boldsymbol{\Sigma}'$, $\mathbf{K} = \mathbf{K}'$ yields the in-distribution matched-noise-correlation setting.

We next treat the training error. We have

$$\begin{aligned} \hat{R}_{\text{in}} &= \frac{1}{T} |\mathbf{y} - \hat{\mathbf{y}}|^2 = \frac{1}{T} |\mathbf{X} \bar{\mathbf{w}} + \boldsymbol{\epsilon} - \mathbf{X} \hat{\mathbf{w}}|^2 \\ &= \frac{1}{T} |\lambda \mathbf{X} (\hat{\boldsymbol{\Sigma}} + \lambda)^{-1} \bar{\mathbf{w}}|^2 + \frac{1}{T} \left| \boldsymbol{\epsilon} - \mathbf{X} (\hat{\boldsymbol{\Sigma}} + \lambda)^{-1} \frac{\mathbf{X}^\top \boldsymbol{\epsilon}}{T} \right|^2 \\ &= \underbrace{\lambda^2 \bar{\mathbf{w}}^\top \hat{\boldsymbol{\Sigma}} (\hat{\boldsymbol{\Sigma}} + \lambda)^{-2} \bar{\mathbf{w}}}_{\text{Signal}_{\text{in}}} + \underbrace{\frac{\lambda^2}{T} \text{Tr} [\mathbf{K}' (\hat{\mathbf{K}} + \lambda)^{-2}]}_{\text{Noise}_{\text{in}}}. \end{aligned} \quad (\text{C55})$$

One can evaluate the signal term by recognizing it as a derivative and applying the chain rule:

$$\begin{aligned} \text{Signal}_{\text{in}} &= -\lambda^2 \frac{d}{d\lambda} \bar{\mathbf{w}}^\top \hat{\boldsymbol{\Sigma}} (\hat{\boldsymbol{\Sigma}} + \lambda)^{-1} \bar{\mathbf{w}} \\ &= \lambda^2 \frac{d\kappa}{d\lambda} \bar{\mathbf{w}}^\top \boldsymbol{\Sigma} (\boldsymbol{\Sigma} + \kappa)^{-2} \bar{\mathbf{w}} \\ &= \lambda^2 S \frac{\tilde{\text{df}}_1 - \tilde{\text{df}}_2}{\tilde{\text{df}}_1} \frac{1}{1 - \gamma} \bar{\mathbf{w}}^\top \boldsymbol{\Sigma} (\boldsymbol{\Sigma} + \kappa)^{-2} \bar{\mathbf{w}}. \end{aligned} \quad (\text{C56})$$

In the last line, we have inserted the form of $\frac{d\kappa}{d\lambda}$ from (C47).

The noise term for generic $\mathbf{K}'$ is slightly more involved. We can write it as:

$$\begin{aligned}\text{Noise}_{\text{in}} &= -\sigma_\epsilon^2 \frac{\lambda^2}{T} \partial_\lambda \text{Tr} [\mathbf{K}'(\hat{\mathbf{K}} + \lambda)^{-1}] \\ &= -\sigma_\epsilon^2 \frac{\lambda^2}{T} \partial_\lambda \left[\frac{\tilde{\kappa}}{\lambda} \text{Tr} [\mathbf{K}'(\mathbf{K} + \tilde{\kappa})^{-1}] \right] \\ &= \sigma_\epsilon^2 \frac{\tilde{\kappa}}{T} \text{Tr} [\mathbf{K}'(\mathbf{K} + \tilde{\kappa})^{-1}] - \sigma_\epsilon^2 \frac{\lambda}{T} \frac{d\tilde{\kappa}}{d\lambda} \text{Tr} [\mathbf{K}\mathbf{K}'(\mathbf{K} + \tilde{\kappa})^{-2}].\end{aligned}\quad (\text{C57})$$

Now applying the derivative relationship (C48) we get:

$$\text{Noise}_{\text{in}} = \sigma_\epsilon^2 \tilde{\kappa} \left[\frac{1}{T} \text{Tr} \mathbf{K}'(\mathbf{K} + \tilde{\kappa})^{-1} - \frac{df_1 - df_2}{df_1} \frac{1}{1 - \gamma} df_{\mathbf{K}, \mathbf{K}'}^2 \right]. \quad (\text{C58})$$

In the case of $\mathbf{K} = \mathbf{K}'$ this simplifies to:

$$\text{Noise}_{\text{in}} = \sigma_\epsilon^2 \tilde{\kappa} \tilde{df}_1 \left[1 - \frac{df_1 - df_2}{df_1 \tilde{df}_1 - df_2 \tilde{df}_2} \tilde{df}_2 \right] = \frac{\sigma_\epsilon^2}{S} \frac{\tilde{df}_1 - \tilde{df}_2}{\tilde{df}_1} \frac{1}{1 - \gamma}. \quad (\text{C59})$$

C.5. Correlated test point

Here, we prove theorem 6.1. Specifically, our setting is one where we allow the test point $\mathbf{x}$ to have non-vanishing correlation with the training set. Similarly, we assume that the label noise on the test point, ϵ has nontrivial correlation with the label noise on the training set. We assume a general matrix-Gaussian model for the covariates:

$$\begin{pmatrix} \mathbf{x}^\top \\ \mathbf{X} \end{pmatrix} \sim \mathcal{N} \left(0, \begin{pmatrix} 1 & \mathbf{k} \\ \mathbf{k}^\top & \mathbf{K} \end{pmatrix} \otimes \Sigma \right), \quad (\text{C60})$$

where the vector $\mathbf{k} \in \mathbb{R}^T$ gives the correlation between the training points and the test point, i.e.

$$\mathbb{E}[x_i x_{tj}] = \Sigma_{ij} k_t. \quad (\text{C61})$$

Similarly for ϵ, ϵ we write:

$$\begin{pmatrix} \epsilon \\ \epsilon \end{pmatrix} \sim \mathcal{N} \left(0, \sigma_\epsilon^2 \begin{pmatrix} 1 & \mathbf{k}^\top \\ \mathbf{k} & \mathbf{K} \end{pmatrix} \right). \quad (\text{C62})$$

We further assume that these are independent of the $\mathbf{x}$ covariates. That is, $\mathbb{E}[x_{ti}\epsilon] = \mathbb{E}[x_i \epsilon_s] = 0$.

By the usual formulas for Gaussian conditioning, this generative model implies that

$$\mathbf{x} | \mathbf{X} \sim \mathcal{N}(\mathbf{X}^\top \boldsymbol{\alpha}, (1 - \rho)\Sigma) \quad (\text{C63})$$

and

$$\epsilon | \epsilon \sim \mathcal{N}(\boldsymbol{\alpha}^\top \epsilon, (1 - \rho)\sigma_\epsilon^2), \quad (\text{C64})$$

where we write

$$\rho = \mathbf{k}^\top \mathbf{K}^{-1} \mathbf{k}, \quad \boldsymbol{\alpha} = \mathbf{K}^{-1} \mathbf{k}. \quad (\text{C65})$$

Then, the out-of-sample risk is

$$R_{\text{out}} = \mathbb{E}[\mathbf{x}^\top (\bar{\mathbf{w}} - \hat{\mathbf{w}}) + \epsilon]^2. \quad (\text{C66})$$

We have the well-known closed-form expression for $\hat{\mathbf{w}}$:

$$\bar{\mathbf{w}} - \hat{\mathbf{w}} = \lambda(\hat{\boldsymbol{\Sigma}} + \lambda)^{-1} \bar{\mathbf{w}} - \frac{1}{T}(\hat{\boldsymbol{\Sigma}} + \lambda)^{-1} \mathbf{X}^\top \boldsymbol{\epsilon}. \quad (\text{C67})$$

Given that $\mathbb{E}[\mathbf{x}_t | \boldsymbol{\epsilon}_s] = 0$, we have:

$$R_{\text{out}} = \underbrace{\lambda^2 \mathbb{E}[\mathbf{x}^\top (\hat{\boldsymbol{\Sigma}} + \lambda)^{-1} \bar{\mathbf{w}}]^2}_{\text{Signal}} + \underbrace{\mathbb{E} \left[\epsilon - \frac{1}{T} \mathbf{x}^\top (\hat{\boldsymbol{\Sigma}} + \lambda)^{-1} \mathbf{X}^\top \boldsymbol{\epsilon} \right]^2}_{\text{Noise}}. \quad (\text{C68})$$

Here, we have again identified the two terms that appear as signal and noise terms. Taking an expectation over $\mathbf{x} | \mathbf{X}$, we have for the signal term

$$\begin{aligned} \text{Signal} &= (1 - \rho) \lambda^2 \mathbb{E} \bar{\mathbf{w}}^\top (\hat{\boldsymbol{\Sigma}} + \lambda)^{-1} \boldsymbol{\Sigma} (\hat{\boldsymbol{\Sigma}} + \lambda)^{-1} \bar{\mathbf{w}} \\ &\quad + \lambda^2 \mathbb{E} \bar{\mathbf{w}}^\top (\hat{\boldsymbol{\Sigma}} + \lambda)^{-1} \mathbf{X}^\top \boldsymbol{\alpha} \boldsymbol{\alpha}^\top \mathbf{X} (\hat{\boldsymbol{\Sigma}} + \lambda)^{-1} \bar{\mathbf{w}}. \end{aligned} \quad (\text{C69})$$

Leveraging the deterministic equivalences (C27) and (C32), we obtain

$$\text{Signal} = \frac{\kappa^2}{1 - \gamma} \bar{\mathbf{w}}^\top \boldsymbol{\Sigma} (\boldsymbol{\Sigma} + \kappa)^{-2} \bar{\mathbf{w}} [1 - \rho + \tilde{\kappa}^2 \boldsymbol{\alpha}^\top \mathbf{K} (\mathbf{K} + \tilde{\kappa})^{-2} \boldsymbol{\alpha}]. \quad (\text{C70})$$

Similarly, taking an expectation first over $\epsilon | \boldsymbol{\epsilon}$ and then over $\mathbf{x} | \mathbf{X}$, we have

$$\text{Noise} = (1 - \rho) \sigma_\epsilon^2 + \mathbb{E} \left[\boldsymbol{\alpha}^\top \boldsymbol{\epsilon} - \frac{1}{T} \mathbf{x}^\top (\hat{\boldsymbol{\Sigma}} + \lambda)^{-1} \mathbf{X}^\top \boldsymbol{\epsilon} \right]^2 \quad (\text{C71})$$

$$\begin{aligned} &= (1 - \rho) \left[\frac{1}{T^2} \mathbb{E} \boldsymbol{\epsilon}^\top \mathbf{X} (\hat{\boldsymbol{\Sigma}} + \lambda)^{-1} \boldsymbol{\Sigma} (\hat{\boldsymbol{\Sigma}} + \lambda)^{-1} \mathbf{X} \boldsymbol{\epsilon} + \sigma_\epsilon^2 \right] \\ &\quad + \mathbb{E} \left[\boldsymbol{\alpha}^\top \boldsymbol{\epsilon} - \boldsymbol{\alpha}^\top \hat{\mathbf{K}} (\hat{\mathbf{K}} + \lambda)^{-1} \boldsymbol{\epsilon} \right]^2. \end{aligned} \quad (\text{C72})$$

Noting that

$$\mathbf{I}_T - \hat{\mathbf{K}} (\hat{\mathbf{K}} + \lambda)^{-1} = \lambda (\hat{\mathbf{K}} + \lambda)^{-1} \quad (\text{C73})$$

this simplifies to

$$\text{Noise} = (1 - \rho) \left[\frac{1}{T^2} \mathbb{E} \boldsymbol{\epsilon}^\top \mathbf{X} (\hat{\boldsymbol{\Sigma}} + \lambda)^{-1} \boldsymbol{\Sigma} (\hat{\boldsymbol{\Sigma}} + \lambda)^{-1} \mathbf{X} \boldsymbol{\epsilon} + \sigma_\epsilon^2 \right] + \lambda^2 \mathbb{E} \left[\boldsymbol{\alpha}^\top (\hat{\mathbf{K}} + \lambda)^{-1} \boldsymbol{\epsilon} \right]^2. \quad (\text{C74})$$

Finally, taking the remaining expectation over ϵ , we have

$$\begin{aligned} \text{Noise} &= \sigma_\epsilon^2(1 - \rho) \left[\mathbb{E} \frac{1}{T^2} \text{Tr}[\mathbf{X}(\hat{\Sigma} + \lambda)^{-1} \Sigma(\hat{\Sigma} + \lambda)^{-1} \mathbf{X} \mathbf{K}] + 1 \right] \\ &\quad + \sigma_\epsilon^2 \lambda^2 \mathbb{E} \boldsymbol{\alpha}^\top (\hat{\mathbf{K}} + \lambda)^{-1} \mathbf{K} (\hat{\mathbf{K}} + \lambda)^{-1} \boldsymbol{\alpha}. \end{aligned} \quad (\text{C75})$$

Now again leveraging deterministic equivalence (C32) as well as (C28) we obtain:

$$\begin{aligned} \text{Noise} &= \sigma_\epsilon^2(1 - \rho) \left[\frac{\gamma}{1 - \gamma} + 1 \right] + \frac{\sigma_\epsilon^2}{1 - \gamma} \tilde{\kappa}^2 \boldsymbol{\alpha}^\top \mathbf{K} (\mathbf{K} + \tilde{\kappa})^{-2} \boldsymbol{\alpha} \\ &= \frac{\sigma_\epsilon^2}{1 - \gamma} [1 - \rho + \tilde{\kappa}^2 \boldsymbol{\alpha}^\top \mathbf{K} (\mathbf{K} + \tilde{\kappa})^{-2} \boldsymbol{\alpha}]. \end{aligned} \quad (\text{C76})$$

All together equations (C70) and (C76) together give the desired result:

$$R_{\text{out}}^k = R_{\text{out}}^{k=0} [1 - \rho + \tilde{\kappa}^2 \boldsymbol{\alpha}^\top \mathbf{K} (\mathbf{K} + \tilde{\kappa})^{-2} \boldsymbol{\alpha}]. \quad (\text{C77})$$

One could straightforwardly extend this to the case where the label noise correlations are different from those of the covariates with relative ease. As doing so adds complexity to the formulas without much conceptual gain, we will not do so here.

Appendix D. Bias-variance decompositions

Often in prior work, it has been convention to call the proportional to the signal $\bar{\mathbf{w}}$ the **Bias**² and the term proportional to the noise σ_ϵ^2 the **variance** component of the total risk. Strictly speaking from the perspective of a fine-grained analysis of variance, this is not true. This was pointed out in detail by [58, 59], where fine-grained bias variance decompositions were performed for random feature models.

The estimator $\hat{\mathbf{w}}$ is given by:

$$\hat{\mathbf{w}} = \hat{\Sigma}(\hat{\Sigma} + \lambda)^{-1} \bar{\mathbf{w}} + (\hat{\Sigma} + \lambda)^{-1} \frac{\mathbf{X}^\top \epsilon}{T}. \quad (\text{D1})$$

We see that the estimator is sensitive to both the draw of ϵ and the choice of $\mathbf{X}$. Even in the case of no noise, $\sigma_\epsilon^2 = 0$, the estimator would still have variance over different draws of $\mathbf{X}$. This variance can be removed if one had access to multiple datasets by **bagging** over the data. In the limit of infinite bagging, this amounts to a data average. This yields, by deterministic equivalence:

$$\mathbb{E}_{\mathbf{X}} \hat{\mathbf{w}} = \Sigma(\Sigma + \kappa)^{-1} \bar{\mathbf{w}}. \quad (\text{D2})$$

Consequently the risk of this estimator would be just the bias:

$$\text{Bias}^2 = R(\mathbb{E}_{\mathbf{X}, \epsilon} \hat{\mathbf{w}}) = \kappa^2 \bar{\mathbf{w}}^\top \Sigma(\Sigma + \kappa)^{-2} \bar{\mathbf{w}}. \quad (\text{D3})$$

Similarly, if one fixed the dataset $\mathbf{X}$ but averaged over different draws of the noise ϵ , the risk of this averaged predictor would be the same as the risk of a predictor trained

on a noiseless dataset $\sigma_\epsilon^2 = 0$. Note that this predictor still has variance over the draw of $\mathbf{X}$. We thus get that:

$$\text{Var}_{\mathbf{X}} = R(\mathbb{E}_\epsilon \hat{\mathbf{w}}) - \text{Bias}^2 = \kappa^2 \bar{\mathbf{w}}^\top \Sigma (\Sigma + \kappa)^{-2} \bar{\mathbf{w}} \frac{\gamma}{1 - \gamma}. \quad (\text{D4})$$

The remaining term is $\text{Var}_{\mathbf{X}, \epsilon}$, since it is removable either by bagging over datasets $\mathbf{X}$ or by averaging over noise ϵ . It is given by:

$$\text{Var}_{\mathbf{X}, \epsilon} = R(\hat{\mathbf{w}}) - \text{Bias}^2 - \text{Var}_{\mathbf{X}} = \frac{\gamma}{1 - \gamma} \sigma_\epsilon^2. \quad (\text{D5})$$

We stress this is true under both correlated and uncorrelated data, given that κ, γ is appropriately defined. This decomposition was computed for linear models in [14, 16]. These equations also hold for out of distribution risk, as given in (C54).

Appendix E. Scaling analysis

E.1. Review of optimal rates

The analysis of the scaling properties of linear regression under power-law decay in the covariance Σ and target $\bar{\mathbf{w}}$ along the principal components of Σ was studied in detail in [39, 60]. It has received renewed attention given that it can sharply characterize the scaling properties of kernel methods, especially the neural tangent kernel [61] on a variety of realistic datasets [19, 40, 41]. This has implications for neural scaling laws [62, 63] in the lazy regime of neural network training identified in [64].

One defines the source and capacity exponents α, r respectively by looking at the decay of the eigenvalues λ_k of Σ and components $\bar{\mathbf{w}}_k$ of $\bar{\mathbf{w}}$ along those eigendirections. Then, for many real datasets one observes the power laws:

$$\lambda_k \sim k^{-\alpha}, \quad \lambda_k \bar{\mathbf{w}}_k^2 \sim k^{-\beta}. \quad (\text{E1})$$

It is advantageous to write $\beta = 2\alpha r + 1$, as the exponent r appears naturally in the final rates. Then, in the ridgeless limit one has $q\text{df}_1 = 1$. From this one can easily obtain that

$$\kappa \sim T^{-\alpha}. \quad (\text{E2})$$

At finite ridge, more generally one has

$$\kappa \sim \max(\lambda, T^{-\alpha}). \quad (\text{E3})$$

For this reason, κ is often called the resolution or the signal capture threshold. It determines which eigencomponents are learned and which are not. Further, taking $\sigma_\epsilon = 0$, one gets the scaling:

$$R_{\text{out}} = R_{\text{g}} \sim \kappa^{2\min(r, 1)}. \quad (\text{E4})$$

If one takes the ridge to scale with the data as $\lambda \sim T^{-\ell}$ as in [41], the full equation in the noiseless setting becomes

$$R_{\text{out}} \sim T^{-2\min(\alpha, \ell)\min(r, 1)}. \quad (\text{E5})$$

These are the rates predicted in [39].

If one passes data through N linear random features before performing, one can change the resolution scaling by adding a further bottleneck:

$$\kappa \sim \max(\lambda, T^{-\alpha}, N^{-\alpha}). \quad (\text{E6})$$

This plays an important role in current models of neural scaling laws [16, 63, 65–67]. It is therefore interesting to ask whether the presence of strong correlations could also affect the scaling laws by adding a further bottleneck to (E6). We answer this in the negative.

E.2. Strong correlations do not affect scaling exponents

Here, we ask whether strongly correlated data could change the scaling in (E3). First, we note that for correlated data, as long as each new sample is linearly independent of the prior ones in $\mathbb{R}^T$, we have that $\mathbf{K}$ does not have any zero eigenvalues. In general, the $T \times T$ correlation matrix $\mathbf{K}$ will have Toeplitz structure when the correlations are stationary and will be invertible as long as each new datapoint is not a linear combination of the prior ones. It could however be that the smallest eigenvalue goes to zero as $T \rightarrow \infty$. This would lead $S_{\mathbf{K}}$ to blow up and could lead to a worse scaling law. We show that this will not happen for exponential, nearest neighbor, and power law correlations. In the latter case, we require some assumptions on the power law exponent.

Proposition E.1. *Under exponential, nearest neighbor, and power law correlations, the $T \times T$ correlation matrix $\mathbf{K}$ remains uniformly nonsingular as $T \rightarrow \infty$.*

Proof. We can study the invertibility of the limiting $T \rightarrow \infty$ matrix by applying standard results on the asymptotic behavior of Hermitian Toeplitz matrices [68]. Let the autocorrelation be $K_{t,t+\tau} = a_\tau$. The three cases of interest are exponential correlations

$$a_\tau = e^{-|\tau|/\xi} \quad (\text{E7})$$

with correlation length $\xi > 0$, nearest-neighbor correlations

$$a_\tau = \begin{cases} 1 & \tau = 0 \\ b/2 & \tau = \pm 1 \\ 0 & \text{otherwise} \end{cases} \quad (\text{E8})$$

with $b \in [0, 1)$ (see appendix I), and power-law correlations

$$a_\tau = \frac{1}{(1 + |\tau|)^\chi} \quad (\text{E9})$$

with exponent $\chi > 0$. The key object is then the *symbol* of the Toeplitz operator corresponding to the finite matrices, which is the Fourier transform of the autocorrelation:

$$w(\theta) = \sum_{k=-\infty}^{\infty} a_k e^{ik\theta} = a_0 + 2 \sum_{k=1}^{\infty} a_k \cos(k\theta), \quad (\text{E10})$$

where $\theta \in [0, 2\pi)$. For the three cases of interest, we can easily work out that

$$w(\theta) = \frac{\sinh(1/\xi)}{\cosh(1/\xi) - \cos(\theta)} \quad (\text{exponential}) \quad (\text{E11})$$

$$w(\theta) = 1 + b \cos(\theta) \quad (\text{nearest-neighbor}) \quad (\text{E12})$$

$$w(\theta) = e^{-i\theta} \text{Li}_\chi(e^{i\theta}) + e^{i\theta} \text{Li}_\chi(e^{-i\theta}) - 1 \quad (\text{power-law}), \quad (\text{E13})$$

where $\text{Li}_\chi(z)$ is the polylogarithm of order χ . All three of these are real, locally-continuous and uniformly positive functions that are symmetric about $\theta = \pi$ and minimized at that point, where they take the values

$$w(\pi) = \tanh \frac{1}{2\xi} \quad (\text{exponential}) \quad (\text{E14})$$

$$w(\pi) = 1 - b \quad (\text{nearest-neighbor}) \quad (\text{E15})$$

$$w(\pi) = 2(1 - 2^{1-\chi})\zeta(\chi) - 1 \quad (\text{power-law}), \quad (\text{E16})$$

where for power-law correlations the result for $\chi = 1$ is understood in a limiting sense¹⁰. For the exponential and nearest-neighbor cases, the location of the minimum is obvious. For power laws, a bit more work is required, starting from the symbol

$$w_\chi(\theta) = 1 + 2 \sum_{k=1}^{\infty} \frac{\cos(k\theta)}{(1+k)^\chi}. \quad (\text{E17})$$

Note that this series converges absolutely if $\chi > 1$, else it is only conditionally convergent. Using the Schwinger parametrization trick, we write

$$\frac{1}{(1+k)^\chi} = \frac{1}{\Gamma(\chi)} \int_0^\infty u^{\chi-1} e^{-(k+1)u} du, \quad (\text{E18})$$

which allows us to write the symbol as

$$w_\chi(\theta) = 1 + \frac{2}{\Gamma(\chi)} \sum_{k=1}^{\infty} \int_0^\infty u^{\chi-1} e^{-(k+1)u} \cos(k\theta) du. \quad (\text{E19})$$

If $\chi > 1$, we can interchange the sum and integral freely, else we may infinitesimally Abel-regularize the sum by multiplying each term by a factor $e^{-k\varsigma}$ and then take $\varsigma \downarrow 0$ at the end of the computation. This leads to

$$w_\chi(\theta) = 1 + \lim_{\varsigma \downarrow 0} \frac{1}{\Gamma(\chi)} \int_0^\infty u^{\chi-1} e^{-u} \left(\frac{\sinh(u+\varsigma)}{\cosh(u+\varsigma) - \cos(\theta)} - 1 \right) du. \quad (\text{E20})$$

Now, the integrand is increasing in $\cos\theta$, and is therefore minimized at $\theta = \pi$ (the potential divergence comes when θ is near zero), which shows that

$$w(\theta) \geq w(\pi) = \frac{1}{\Gamma(\chi)} \int_0^\infty u^{\chi-1} e^{-u} \frac{\sinh(u)}{\cosh(u) + 1} du = 2(1 - 2^{1-\chi})\zeta(\chi) - 1. \quad (\text{E21})$$

¹⁰ In particular, we have $\lim_{\chi \rightarrow 1} w(\pi) = \log 4 - 1 \simeq 0.386 \dots$

This proves the desired claim. We remark that the integral computation is consistent with the direct evaluation of the sum based on recognizing it as the Dirichlet η -function, which is related to the Riemann ζ -function as $\eta(s) = (1 - 2^{1-s})\zeta(s)$:

$$w_\chi(\pi) = 2 \sum_{k=1}^{\infty} \frac{(-1)^{k-1}}{k^\chi} - 1 = 2\eta(\chi) - 1 = 2(1 - 2^{1-\chi})\zeta(\chi) - 1; \quad (\text{E22})$$

the series $\sum_{k=1}^{\infty} \frac{(-1)^{k-1}}{k^\chi} = \eta(\chi)$ converges for all $\chi > 0$. That said, the integral expression has the advantage of making it obvious that $w_\chi(\pi) > 0$.

We now translate the symbol bounds $w(\theta) \geq w(\pi)$ into uniform-in- T bounds on the minimum eigenvalue:

$$\lambda_{\min}(K) \geq w(\pi). \quad (\text{E23})$$

We start by recalling the Rayleigh quotient characterization of the minimum eigenvalue of a real symmetric matrix:

$$\lambda_{\min}(K) = \min_{\|v\|=1} v^\top K v. \quad (\text{E24})$$

For any test vector $v \in \mathbb{R}^T$, we can write

$$v^\top K v = \sum_{t,t'=0}^{T-1} v_t a_{t-t'} v_{t'} = \frac{1}{2\pi} \int_0^{2\pi} w(\theta) \sum_{t,t'=0}^{T-1} v_t v_{t'} e^{-i(t-t')\theta} d\theta = \frac{1}{2\pi} \int_0^{2\pi} w(\theta) \left| \sum_{t=0}^{T-1} v_t e^{-it\theta} \right|^2 d\theta \quad (\text{E25})$$

as by definition

$$a_k = \frac{1}{2\pi} \int_0^{2\pi} w(\theta) e^{-ik\theta} d\theta. \quad (\text{E26})$$

Thus, if we have the lower bound $w(\theta) \geq w(\pi) > 0$, then we can bound

$$v^\top K v \geq w(\pi) \frac{1}{2\pi} \int_0^{2\pi} \left| \sum_{t=0}^{T-1} v_t e^{-it\theta} \right|^2 d\theta = w(\pi) \|v\|^2 = w(\pi) \quad (\text{E27})$$

for any test vector v , and therefore we conclude that

$$\lambda_{\min}(K) \geq w(\pi). \quad (\text{E28})$$

This lower bound is uniform in T .

Thus, as for each case of interest $w(\pi) > 0$ within the valid range of parameters, all of the matrices of interest remain invertible, and in particular are *uniformly* non-singular; their minimum eigenvalues are bounded from below by a T -independent positive constant. For illustrative purposes, we compare (E16) to numerical computation of the minimum eigenvalue for power-law correlations in figure 8.

□




Figure 8. Verification of the minimum eigenvalue for a Toeplitz matrix with power law decay $K_{ts} = (1 + |t - s|)^{-\chi}$. We see a great match between (E16) (solid black line) and the empirics (x shapes) for matrix sizes ranging between 2×2 and 1024×1024 . In particular, this guarantees that the matrix will never be singular even at very small positive χ , corresponding to strong power law correlations. It is interesting to note that the minimum eigenvalue of a 2×2 Toeplitz matrix of this form, which is simply $\lambda_{\min} = 1 - 2^{-\chi}$, is relatively close to the minimum eigenvalue of the infinite operator; their ratio decreases from $-\frac{\log(2)}{\log(2/\pi)} \simeq 1.534\dots$ at $\chi \downarrow 0$ towards 1 as $\chi \rightarrow \infty$.

Under the assumption that $\mathbf{K}$ remains invertible in the large T limit, we have that $\lim_{\lambda \rightarrow 0} \text{df}_{\mathbf{K}}(\lambda) = 1$ at any value of T . More generally $\text{df}_{\mathbf{K}}$ is a continuous monotonically decreasing function of λ that goes from 1 when $\lambda = 0$ to 0 as $\lambda \rightarrow \infty$. As $\lambda \rightarrow \infty$ we can expand and get at linear order that $\text{df}_{\mathbf{K}} = \frac{1}{\lambda} \frac{1}{T} \text{Tr} \mathbf{K} = \frac{1}{\lambda}$. Writing:

$$\text{df}_{\mathbf{K}}(\lambda) = \frac{S_{\mathbf{K}}(\text{df}_{\mathbf{K}}(\lambda))^{-1}}{\lambda + S_{\mathbf{K}}(\text{df}_{\mathbf{K}}(\lambda))^{-1}} \quad (\text{E29})$$

implies that at small λ , $S_{\mathbf{K}}$ is bounded from above by a T -independent constant at any T . $S_{\mathbf{K}}$ is also bounded from below by 1 (see F.1). Then, writing

$$\kappa = \frac{\lambda}{1 - q \text{df}_1} S_{\mathbf{K}}(q \text{df}_1), \quad (\text{E30})$$

we see that $S_{\mathbf{K}}$ will not contribute any pole or zero that would affect the scaling properties of κ as $\lambda \rightarrow 0$. Consequently, correlated data does not affect the scaling law.

Appendix F. Double descent analysis

In this appendix, we study how double descent is affected by correlations. We also more generally study how correlations affect the key quantities of interest relative to the uncorrelated setting.

F.1. Bounds on renormalized ridges

Lemma F.1. For positive-definite $\mathbf{K}$ such that $\frac{1}{T}\text{Tr}(\mathbf{K}) = 1$ we have that $S_{\mathbf{K}}(\tilde{\text{df}}) \geq 1$ for all values of $\tilde{\text{df}} \in (0, 1]$.

Proof. We adapt an argument from Zavattone-Veth and Pehlevan [69]. We recognize that for fixed $\tilde{\kappa} > 0$ the function:

$$\rho \mapsto \frac{\rho}{\tilde{\kappa} + \rho} \quad (\text{F1})$$

is concave. Consequently, Jensen's inequality yields:

$$\text{df}_{\mathbf{K}}(\tilde{\kappa}) = \mathbb{E}_{\rho} \frac{\rho}{\tilde{\kappa} + \rho} \leq \frac{\mathbb{E}_{\rho} \rho}{\tilde{\kappa} + \mathbb{E}_{\rho} \rho} = \frac{1}{\tilde{\kappa} + 1} = \text{df}_{\mathbf{I}}(\tilde{\kappa}) \quad (\text{F2})$$

pointwise in $\tilde{\kappa}$. Here $\mathbb{E}_{\rho}$ is the expectation over the spectrum of $\mathbf{K}$. We have also used that $\mathbb{E}_{\rho}[\rho] = \frac{1}{T}\text{Tr}(\mathbf{K}) = 1$. The above inequality is strict unless $\mathbf{K} = \mathbf{I}_T$ or $\tilde{\kappa} = 0$.

We next observe that both $\text{df}_{\mathbf{K}}(\lambda)$ and $\text{df}_{\mathbf{I}}(\lambda)$ are monotonically decreasing function in λ that are equal to 1 only when $\lambda = 0$. This means that the solutions to the equations $\tilde{\text{df}} = \text{df}_{\mathbf{K}}(\lambda)$ and $\tilde{\text{df}} = \text{df}_{\mathbf{I}}(\lambda)$ are unique for all $\tilde{\text{df}} \in (0, 1]$ and is also monotonically decreasing. Consequently we have:

$$\text{df}_{\mathbf{K}}^{-1}(\tilde{\text{df}}) \leq \text{df}_{\mathbf{I}}^{-1}(\tilde{\text{df}}) = \frac{1 - \tilde{\text{df}}}{\tilde{\text{df}}}. \quad (\text{F3})$$

Upon dividing both sides by $\text{df}_{\mathbf{K}}^{-1}$ we get the desired equality $\Rightarrow S_{\mathbf{K}}(\tilde{\text{df}}) \geq 1$. This is an equality when $\mathbf{K} = \mathbf{I}_T$. When $\tilde{\kappa} = 0$ and thus $\tilde{\text{df}} = 1$, $S_{\mathbf{K}}(\tilde{\text{df}})$ may still be greater than one. $\square$

Proposition F.2. Let κ_c be the renormalized ridge λS when the data has correlation structure $\mathbf{K}$ and let κ_u be the corresponding value of κ when there is no correlation between data points. We have $\kappa_c \geq \kappa_u$. Moreover, in the ridgeless limit $\kappa_c = \kappa_u$.

Proof. In the correlated setting, we recall that the renormalized ridge κ_c solves

$$\kappa_c = \frac{\lambda S_{\mathbf{K}}(q\text{df}_{\Sigma}^1(\kappa_c))}{1 - q\text{df}_{\Sigma}^1(\kappa_c)}. \quad (\text{F4})$$

Call the numerator $\tilde{\lambda}$. From the preceding lemma we have that $\tilde{\lambda} \geq \lambda$. Thus, we have that κ_c is equivalent to the self-consistent solution for the equation

$$\kappa_c = \frac{\tilde{\lambda}}{1 - q\text{df}_{\Sigma}^1(\kappa_c)}. \quad (\text{F5})$$

But this is the same as the self-consistent equation for a linear regression problem with an explicit ridge $\tilde{\lambda} \geq \lambda$. The proposition then follows from the fact that κ is monotonic in λ in the uncorrelated ridge regression setting.

As $\lambda \rightarrow 0$, κ is entirely determined by the pole structure of $q\text{df}_{\Sigma}^1 = 1$. This is independent of any structure on $\mathbf{K}$ and so in the ridgeless limit, $\kappa_c = \kappa_u$. $\square$

Corollary F.3. Fix q . The presence of correlations either decreases or keeps constant $\text{df}_1, \tilde{\text{df}}_1$. In the ridgeless limit, they are unchanged.

Proof. Evaluating df_1 as $\text{df}_1 = \text{df}_\Sigma^1(\kappa)$ we have since $\kappa_c \geq \kappa_u$ and df_1 is monotone decreasing in κ that df_1 decreases. Since $\tilde{\text{df}}_1 = q\text{df}_1$, this also decreases. Because κ is unchanged in the ridgeless limit, the last part of the corollary follows. $\square$

Corollary F.4. Fix q . The presence of correlations either decreases or keeps constant df_2 . In the ridgeless limit, it is unchanged.

Proof. The proof is as in the prior corollary, noting that $\text{df}_2 = \text{df}_\Sigma^2(\kappa)$ is monotone decreasing in κ and $\kappa_c \geq \kappa_u$. $\square$

Proposition F.5. Let $\tilde{\kappa}_c$ be the renormalized ridge $\lambda\tilde{S}$ when the data has correlation structure $\mathbf{K}$ and let $\tilde{\kappa}_u$ be the corresponding value of $\tilde{\kappa}$ when there is no correlation between data points. We have $\tilde{\kappa}_c \leq \tilde{\kappa}_u$. Moreover, in the ridgeless limit $\tilde{\kappa}_c \leq \tilde{\kappa}_u$ still.

Proof. When $\lambda \neq 0$ we can rewrite the duality relation (C36) as

$$\frac{q}{\lambda} \kappa \text{df}_\Sigma^1(\kappa) = \frac{1}{\tilde{\kappa}}. \quad (\text{F6})$$

This separately holds true for the pairs $\kappa_u, \tilde{\kappa}_u$ and $\kappa_c, \tilde{\kappa}_c$. We note that on the left hand side $\kappa \text{df}_\Sigma^1(\kappa)$ is a sum of $\frac{\kappa}{\lambda_i + \kappa}$ over the eigenspectrum λ_i of Σ . Each term is monotone increasing in κ , and thus increases as we go from the uncorrelated κ_u to the correlated κ_c . Consequently, $\tilde{\kappa}_c \leq \tilde{\kappa}_u$.

In the ridgeless limit, by the prior corollary, $\tilde{\text{df}}$ is unchanged between correlated and uncorrelated data. Consequently, we have:

$$\frac{1}{1 + \tilde{\kappa}_u} = \text{df}_{\mathbf{I}_T}^1(\tilde{\kappa}_u) = \text{df}_{\mathbf{K}}^1(\tilde{\kappa}_c) \leq \frac{1}{1 + \tilde{\kappa}_c}. \quad (\text{F7})$$

Thus, we again have $\tilde{\kappa}_c \leq \tilde{\kappa}_u$ with equality only if $\mathbf{K} = \mathbf{I}_T$ or $\kappa_c = 0$. $\square$

F.2. Ridgeless limit

The limiting behavior of the two renormalized ridges κ and $\tilde{\kappa}$ depends on whether one is in the underparameterized regime $T > N$ ($q < 1$) or the overparameterized regime $T < N$ ($q > 1$). In the underparameterized regime, we have that $\kappa \downarrow 0$ as $\lambda \downarrow 0$, while generically $\tilde{\kappa}$ remains non-zero and solves the self-consistent equation $q = \tilde{\text{df}}_1(\tilde{\kappa})$. Conversely, in the overparameterized regime $\tilde{\kappa} \downarrow 0$ as $\lambda \downarrow 0$ while κ remains non-zero and solves $1/q = \text{df}_1(\kappa)$. We numerically illustrate these complementary limiting behaviors in appendix J.

Our task is now to determine the corresponding limits of the out-of-sample risk. First, we consider the case of matched correlations. There, in the underparameterized regime we have $\gamma \rightarrow \tilde{\text{df}}_2/q$, while in the overparameterized regime we have $\gamma \rightarrow q\text{df}_2$. As

Risk and cross validation in ridge regression with correlated samples




Figure 9. Theory curves comparing $\kappa, \tilde{\kappa}, df_1, \tilde{df}_1, df_2, \tilde{df}_2$ for correlated data $\mathbf{K} \neq \mathbf{I}$ vs uncorrelated data $\mathbf{K} = \mathbf{I}$ at $\lambda = 10^{-8}$ for $N = 100$ on isotropic data. For the correlations, we choose exponential correlations with length $\xi = 10^2$. (a) We see that when $T < N$, $\kappa, \tilde{\kappa}$ strongly agree, whereas when $T > N$, κ becomes of order the ridge while $\tilde{\kappa}$ becomes order 1. The effects of correlations on $\tilde{\kappa}$ are therefore noticeable in this limit. (b) We see that the $df_1, \tilde{df}_1$ are unchanged between correlated (solid) and uncorrelated (dashed) data when the ridge is small. (c) We see a similar behavior for df_2 but not for $\tilde{df}_2$ when $T > N$. Still, at first order near $T = N$, we see agreement of $\tilde{df}_2$ between the correlated and uncorrelated settings, as predicted by our theory.

a result, we find that

$$\lim_{\lambda \downarrow 0} R_g \simeq \begin{cases} \frac{\tilde{df}_2}{q - \tilde{df}_2} \sigma_\epsilon^2 & q < 1, \\ \frac{\kappa^2}{1 - q \tilde{df}_2} \bar{\mathbf{w}}^\top \Sigma (\Sigma + \kappa)^{-2} \bar{\mathbf{w}} + \frac{q \tilde{df}_2}{1 - q \tilde{df}_2} \sigma_\epsilon^2 & q > 1. \end{cases} \quad (\text{F8})$$

In the overparameterized ridgeless setting, since neither κ, df_1 or df_2 is modified, the generalization error is exactly identical as that for linear regression. In the underparameterized limit, we study how $\tilde{df}_2$ behaves. We can write

$$df_{\mathbf{K}}^2(\tilde{\kappa}_c) = \mathbb{E}_\rho \left(\frac{\rho}{\rho + \tilde{\kappa}_c} \right)^2 \geq \left[\mathbb{E}_\rho \left(\frac{\rho}{\rho + \tilde{\kappa}_c} \right) \right]^2 = df_{\mathbf{K}}^1(\tilde{\kappa}_c)^2 = df_{\mathbf{I}}^1(\tilde{\kappa}_u)^2 = df_{\mathbf{I}}^2(\tilde{\kappa}_u), \quad (\text{F9})$$

where the expectation is taken over the eigenspectrum of $\mathbf{K}$. Moreover, near $q = 1$ when $\tilde{\kappa}$ is small we have that at linear order:

$$df_{\mathbf{K}}^2(\tilde{\kappa}_c) = 1 - 2\tilde{\kappa}_c \frac{1}{T} \text{Tr}[\mathbf{K}^{-1}] + O(\tilde{\kappa}_c^2) = df_{\mathbf{K}}^1(\tilde{\kappa}_c)^2 + O(\tilde{\kappa}_c^2) = df_{\mathbf{I}}^2(\tilde{\kappa}_u) + O(\tilde{\kappa}_u^2). \quad (\text{F10})$$

Thus, at first order the double descent peak is unaffected by correlations in the underparameterized regime as well. We illustrate these observations in figure 9.

F.2.1. Effective ridge enhancement. If λ is finite, we have by proposition F.2 that by F.5 that $\kappa_c > \kappa_u$ and $\tilde{\kappa}_c < \tilde{\kappa}_u$. Specifically, let $\hat{\Sigma}_{\mathbf{K}}$ be the empirical covariance of the data when the data points have correlations $\mathbf{K}$. We can write the following equivalence:

$$\hat{\Sigma}_{\mathbf{K}} (\hat{\Sigma}_{\mathbf{K}} + \lambda)^{-1} \simeq \hat{\Sigma}_{\mathbf{I}} (\hat{\Sigma}_{\mathbf{I}} + \lambda S_{\mathbf{K}}(q df_1))^{-1}. \quad (\text{F11})$$

Thus adding correlations amounts to increasing the ridge. Because $\Sigma_{\mathbf{K}}$ is bounded from above from proposition E.1, in the ridgeless limit the $\lambda S_{\mathbf{K}}$ will remain zero.

Risk and cross validation in ridge regression with correlated samples




Figure 10. Theory curves comparing $\kappa, \tilde{\kappa}, df_1, \tilde{df}_1, df_2, \tilde{df}_2$ for correlated data $\mathbf{K} \neq \mathbf{I}$ vs uncorrelated data $\mathbf{K} = \mathbf{I}$ at $\lambda = 10^{-1}$ for $N = 100$ on isotropic data. For the correlations, we choose exponential correlations with length $\xi = 10^2$. (a) We see across the board that in the presence of explicit ridge, κ grows while $\tilde{\kappa}$ shrinks under correlations. (b) We see that the $df_1, \tilde{df}_1$ are always decreased by the presence of correlations when the ridge is present. (c) We see a similar behavior for df_2 . For $\tilde{df}_2$, we also see a decrease in the neighborhood of the double descent peak $T = N$. Still, at first order near $T = N$, we see agreement of $\tilde{df}_2$ between the correlated and uncorrelated settings, as predicted by our theory.

When there is explicit ridge, we have that $df_1, \tilde{df}_1, df_2$ will shrink from their uncorrelated values. Because to linear order in $\kappa, \tilde{\kappa}$, we have $\tilde{df}_2 = (\tilde{df}_1)^2$ at $q = 1$ and because correlations cause $\tilde{df}_1$ to decrease at first order in κ at this point, we get that correlations cause $\tilde{df}_2$ to decrease in a neighborhood of $q = 1$.

As a result of this, at first order in $\kappa, \tilde{\kappa}$, we can write $\gamma = qdf_2(\kappa)$. Near $q = 1$, this will also shrink in the correlated case relative to the uncorrelated case. Thus, the double descent effect will be further reduced. We illustrate these observations in figure 10.

F.3. Mismatched correlations

Now we consider the case of mismatched correlations. We consider the fully general setting in which $\Sigma' \neq \Sigma$ and $\mathbf{K}' \neq \mathbf{K}$. In the underparameterized regime, we have $df_{\Sigma\Sigma'}^2 \rightarrow \frac{1}{N}\text{Tr}(\Sigma^{-1}\Sigma')$, so $\gamma_{\Sigma\Sigma'} \rightarrow q^{-1}df_2\frac{1}{N}\text{Tr}(\Sigma^{-1}\Sigma')$. Similarly, $\gamma_{\Sigma\Sigma'\mathbf{K}\mathbf{K}'} \rightarrow q^{-1}df_{\mathbf{K}\mathbf{K}'}^2\frac{1}{N}\text{Tr}(\Sigma^{-1}\Sigma')$. In the overparameterized regime, we similarly have $df_{\mathbf{K}\mathbf{K}'}^2 \rightarrow \frac{1}{T}\text{Tr}(\mathbf{K}^{-1}\mathbf{K}')$, which leads to $\gamma_{\Sigma\Sigma'} \rightarrow qdf_{\Sigma\Sigma'}^2$ and $\gamma_{\Sigma\Sigma'\mathbf{K}\mathbf{K}'} \rightarrow qdf_{\Sigma\Sigma'}^2\frac{1}{T}\text{Tr}(\mathbf{K}^{-1}\mathbf{K}')$. Combining these results, we find that

$$\lim_{\lambda \downarrow 0} R_g \simeq \begin{cases} \frac{q^{-1}df_{\mathbf{K}\mathbf{K}'}^2\frac{1}{N}\text{Tr}(\Sigma^{-1}\Sigma')}{1-q^{-1}df_2}\sigma_\epsilon^2 & q < 1, \\ \kappa^2\bar{\mathbf{w}}(\Sigma + \kappa)^{-1}\Sigma'(\Sigma + \kappa)^{-1}\bar{\mathbf{w}} + \kappa^2\frac{qdf_{\Sigma\Sigma'}^2}{1-qdf_2}\bar{\mathbf{w}}\Sigma(\Sigma + \kappa)^{-2}\bar{\mathbf{w}} & \\ + \frac{qdf_{\Sigma\Sigma'}^2\frac{1}{T}\text{Tr}(\mathbf{K}^{-1}\mathbf{K}')}{1-qdf_2}\sigma_\epsilon^2 & q > 1. \end{cases} \quad (\text{F12})$$

The effect of mismatch is easiest to understand in the overparameterized regime, where the factor $\frac{1}{T}\text{Tr}(\mathbf{K}^{-1}\mathbf{K}')$ multiplies the same expression for the noise term that appears for uncorrelated datapoints. In the special case $\mathbf{K}' = \mathbf{I}_T$, we can use Jensen's inequality

to bound

$$\frac{1}{T} \text{Tr}(\mathbf{K}^{-1}) \geq \frac{1}{\text{Tr}(\mathbf{K})/T} = 1 \quad (\text{F13})$$

with equality iff $\mathbf{K} = \mathbf{I}_T$, hence in this special case mismatch generically increases the error.

F.4. Further comments on the effect of mismatched correlations

Here, we briefly comment further on the case in which the noise is uncorrelated ($\mathbf{K}' = \mathbf{I}_T$). In greatest generality, we seek to bound

$$\text{df}_{\mathbf{K}, \mathbf{K}'=\mathbf{I}_T}^2(\tilde{\kappa}) = \frac{1}{T} \text{Tr}[\mathbf{K}(\mathbf{K} + \tilde{\kappa})^{-2}] \quad (\text{F14})$$

in terms of

$$\tilde{\text{df}}_2(\tilde{\kappa}) = \frac{1}{T} \text{Tr}[\mathbf{K}^2(\mathbf{K} + \tilde{\kappa})^{-2}]. \quad (\text{F15})$$

We prove a simple upper bound, which follows from a negative association argument:

Proposition F.6. *For any $\tilde{\kappa} \geq 0$ and $\mathbf{K}$ invertible, we have*

$$\text{df}_{\mathbf{K}, \mathbf{K}'=\mathbf{I}_T}^2(\tilde{\kappa}) \leq \left(\frac{1}{T} \text{Tr}(\mathbf{K}^{-1}) \right) \tilde{\text{df}}_2(\tilde{\kappa}), \quad (\text{F16})$$

with equality when $\tilde{\kappa} = 0$.

Proof. If $\tilde{\kappa} = 0$, the claim obviously holds with equality so long as $\mathbf{K}$ is invertible. Then, for any fixed $\tilde{\kappa} > 0$, define

$$f(\rho) = \frac{\rho^2}{(\rho + \tilde{\kappa})^2} \quad (\text{F17})$$

and

$$g(\rho) = \frac{1}{\rho}, \quad (\text{F18})$$

such that

$$\text{df}_{\mathbf{K}, \mathbf{K}'=\mathbf{I}_T}^2(\tilde{\kappa}) = \mathbb{E}[f(\rho)g(\rho)] \quad (\text{F19})$$

and

$$\tilde{\text{df}}_2(\tilde{\kappa}) = \mathbb{E}[f(\rho)], \quad (\text{F20})$$

where expectation is taken with respect to the distribution of eigenvalues of $\mathbf{K}$. Observe that $f(\rho)$ is for any $\tilde{\kappa} > 0$ a monotone increasing function on $(0, \infty)$, while $g(\rho)$ is a monotone decreasing function on $(0, \infty)$. Then, for any $\rho_1, \rho_2 \in (0, \infty)$, we have

$$[f(\rho_1) - f(\rho_2)][g(\rho_1) - g(\rho_2)] \leq 0, \quad (\text{F21})$$

so upon taking expectations for ρ_1, ρ_2 independently drawn from the eigenvalue distribution of $\mathbf{K}$ we have

$$0 \geq \mathbb{E} \left[[f(\rho_1) - f(\rho_2)][g(\rho_1) - g(\rho_2)] \right] = 2\mathbb{E}[f(\rho)g(\rho)] - 2\mathbb{E}[f(\rho)]\mathbb{E}[g(\rho)] \quad (\text{F22})$$

hence

$$\mathbb{E}[f(\rho)g(\rho)] \leq \mathbb{E}[f(\rho)]\mathbb{E}[g(\rho)]. \quad (\text{F23})$$

Using the definitions of f and g , this concludes the proof. $\square$

Appendix G. Asymptotics of previously-proposed extensions of the GCV

Here, we compute the high-dimensional asymptotics of previously-proposed extensions to the GCV in the presence of correlations. We write each estimator in the form

$$\hat{E}_{\text{predicted}} = G\hat{R}_{\text{in}}. \quad (\text{G1})$$

Moreover, we introduce the notation

$$\mathbf{H} = (\hat{\mathbf{K}} + \lambda)^{-1} \hat{\mathbf{K}} \quad (\text{G2})$$

for the smoothing matrix, in terms of which the predictions on the training set are given as

$$\hat{\mathbf{y}} = \mathbf{H}\mathbf{y}, \quad (\text{G3})$$

as these estimators make use of this matrix. For comparison, we recall from the main text that the asymptotically precise CorrGCV estimator is

$$G_{\text{CorrGCV}} = S(\text{df}_1) \frac{\tilde{\text{df}}_1}{\tilde{\text{df}}_1 - \tilde{\text{df}}_2}. \quad (\text{G4})$$

G.1. The GCV

The standard GCV can be written as

$$G_{\text{GCV}} = \frac{1}{[1 - \frac{1}{T}\text{Tr}(\mathbf{H})]^2}, \quad (\text{G5})$$

hence we have immediately that

$$G_{\text{GCV}} \simeq \frac{1}{[1 - \tilde{\text{df}}_1(\tilde{\kappa})]^2}. \quad (\text{G6})$$

We denote this estimator by GCV_1 in figure 5 and in all plots.

G.2. The Altman estimator

Starting from the GCV, Altman [26] and later Opsomer *et al* [27] consider regression with correlated errors

$$\text{cov}(\epsilon) = \sigma_\epsilon^2 \mathbf{K}_\epsilon \quad (\text{G7})$$

and consider the estimator

$$G_{\text{Altman}} = \frac{1}{[1 - \frac{1}{T} \text{Tr}(\mathbf{H} \mathbf{K}_\epsilon)]^2}. \quad (\text{G8})$$

Asymptotically, we have immediately that

$$G_{\text{Altman}} \simeq \frac{1}{\{1 - \frac{1}{T} \text{Tr}[\mathbf{K}_\epsilon \mathbf{K} (\mathbf{K} + \tilde{\kappa})^{-1}]\}^2}. \quad (\text{G9})$$

Putting $\mathbf{K}_\epsilon = \mathbf{K}$, we can write

$$\frac{1}{T} \text{Tr}[\mathbf{K}^2 (\mathbf{K} + \tilde{\kappa})^{-1}] = \frac{1}{T} \text{Tr}[\mathbf{K} (\mathbf{I} - \tilde{\kappa} (\mathbf{K} + \tilde{\kappa})^{-1})] = \frac{1}{T} \text{Tr}(\mathbf{K}) - \tilde{\kappa} \tilde{\text{df}}_1, \quad (\text{G10})$$

hence, as $\frac{1}{T} \text{Tr}(\mathbf{K}) = 1$ by normalization, we have

$$G_{\text{Altman}} \simeq \frac{1}{\tilde{\kappa}^2 \tilde{\text{df}}_1^2}. \quad (\text{G11})$$

Using the duality relation

$$\frac{\kappa \tilde{\kappa}}{\lambda} = \frac{1}{\tilde{\text{df}}_1}, \quad (\text{G12})$$

this reduces to

$$G_{\text{Altman}} \simeq \frac{\kappa^2}{\lambda^2} = S^2. \quad (\text{G13})$$

We denote this estimator by GCV_2 in figure 5 and in all plots.

G.3. The GCCV estimator of Carmack *et al*

In our notation, Carmack *et al* [28] again assume that the label noise has

$$\text{cov}(\epsilon) = \sigma_\epsilon^2 \mathbf{K}_\epsilon \quad (\text{G14})$$

and consider the estimator

$$G_{\text{Carmack}} = \frac{1}{[1 - \frac{1}{T} \text{Tr}(2\mathbf{H} \mathbf{K}_\epsilon - \mathbf{H} \mathbf{K}_\epsilon \mathbf{H}^\top)]^2}. \quad (\text{G15})$$

The first term is identical to Altman's [26] estimator, while the second is an additional correction. We have

$$\frac{1}{T} \text{Tr}(\mathbf{H}\mathbf{K}_\epsilon) \simeq \frac{1}{T} \text{Tr}[\mathbf{K}_\epsilon \mathbf{K}(\mathbf{K} + \tilde{\kappa})^{-1}] \quad (\text{G16})$$

and

$$\frac{1}{T} \text{Tr}(\mathbf{H}\mathbf{K}_\epsilon \mathbf{H}^\top) = \frac{1}{T} \text{Tr}[\mathbf{K}_\epsilon \hat{\mathbf{K}}(\hat{\mathbf{K}} + \lambda)^{-1}] - \lambda \frac{1}{T} \text{Tr}[\mathbf{K}_\epsilon \hat{\mathbf{K}}(\hat{\mathbf{K}} + \lambda)^{-2}] \quad (\text{G17})$$

$$= \frac{1}{T} \text{Tr}[\mathbf{K}_\epsilon \hat{\mathbf{K}}(\hat{\mathbf{K}} + \lambda)^{-1}] + \lambda \partial_\lambda \frac{1}{T} \text{Tr}[\mathbf{K}_\epsilon \hat{\mathbf{K}}(\hat{\mathbf{K}} + \lambda)^{-1}] \quad (\text{G18})$$

$$\simeq \frac{1}{T} \text{Tr}[\mathbf{K}_\epsilon \mathbf{K}(\mathbf{K} + \tilde{\kappa})^{-1}] - \lambda \frac{\partial \tilde{\kappa}}{\partial \lambda} \frac{1}{T} \text{Tr}[\mathbf{K}_\epsilon \mathbf{K}(\mathbf{K} + \tilde{\kappa})^{-2}], \quad (\text{G19})$$

so

$$\frac{1}{T} \text{Tr}(2\mathbf{H}\mathbf{K}_\epsilon - \mathbf{H}\mathbf{K}_\epsilon \mathbf{H}^\top) \simeq \frac{1}{T} \text{Tr}[\mathbf{K}_\epsilon \mathbf{K}(\mathbf{K} + \tilde{\kappa})^{-1}] + \lambda \frac{\partial \tilde{\kappa}}{\partial \lambda} \frac{1}{T} \text{Tr}[\mathbf{K}_\epsilon \mathbf{K}(\mathbf{K} + \tilde{\kappa})^{-2}]. \quad (\text{G20})$$

Setting $\mathbf{K}_\epsilon = \mathbf{K}$, we have

$$1 - \frac{1}{T} \text{Tr}(2\mathbf{H}\mathbf{K} - \mathbf{H}\mathbf{K}\mathbf{H}^\top) \simeq \tilde{\kappa} \tilde{\text{df}}_1 - \lambda \frac{\partial \tilde{\kappa}}{\partial \lambda} \tilde{\text{df}}_2 \quad (\text{G21})$$

as $\frac{1}{T} \text{Tr}(\mathbf{K}) = 1$. Now, using (C48) and the defining equation $\lambda \tilde{S} = \tilde{\kappa}$, we have

$$1 - \frac{1}{T} \text{Tr}(2\mathbf{H}\mathbf{K} - \mathbf{H}\mathbf{K}\mathbf{H}^\top) \simeq \tilde{\kappa} \tilde{\text{df}}_1 - \tilde{\kappa} \frac{1 - \frac{\text{df}_2}{\text{df}_1}}{1 - \gamma} \tilde{\text{df}}_2 \quad (\text{G22})$$

with

$$\gamma = \frac{\text{df}_2}{\text{df}_1} \frac{\tilde{\text{df}}_2}{\tilde{\text{df}}_1}. \quad (\text{G23})$$

Then, we have

$$G_{\text{Carmack}} \simeq \left[\tilde{\kappa} \tilde{\text{df}}_1 - \tilde{\kappa} \frac{1 - \frac{\text{df}_2}{\text{df}_1}}{1 - \gamma} \tilde{\text{df}}_2 \right]^{-2}. \quad (\text{G24})$$

We now write

$$\frac{\text{df}_2}{\text{df}_1} = \gamma \frac{\tilde{\text{df}}_1}{\tilde{\text{df}}_2} \quad (\text{G25})$$

and use the duality relation to expand $\tilde{\kappa} \tilde{\text{df}}_1 = \lambda / \kappa = 1/S$, which upon combining terms gives

$$G_{\text{Carmack}} \simeq \left[\frac{1}{S} \frac{1}{1 - \gamma} \frac{\tilde{\text{df}}_1 - \tilde{\text{df}}_2}{\tilde{\text{df}}_1} \right]^{-2}. \quad (\text{G26})$$

Risk and cross validation in ridge regression with correlated samples




Figure 11. Plot of the Carmack *et al* [28] estimator relative to other GCV estimators defined earlier. We note that similar to the naive GCVs 1 and 2, the Carmack estimator fails to correctly predict R . (a) Unstructured covariates. (b) Structured covariates with source and capacity exponents as labeled.

We recognize the term in brackets as the noise term when $\mathbf{K} = \mathbf{K}'$ in (C59). In general this implies that

$$G_{\text{Carmack}} = (1 - \gamma)^2 S^2 \underbrace{\left(\frac{\tilde{\text{df}}_1}{\tilde{\text{df}}_1 - \tilde{\text{df}}_2} \right)^2}_{\text{CorrGCV}^2} = S^2 \left(\frac{1 - \frac{\text{df}_2 \tilde{\text{df}}_2}{\text{df}_1 \tilde{\text{df}}_1}}{1 - \frac{\tilde{\text{df}}_2}{\tilde{\text{df}}_1}} \right)^2. \quad (\text{G27})$$

When $\mathbf{K} = \mathbf{I}_T$, $\tilde{\text{df}}_1/(\tilde{\text{df}}_1 - \tilde{\text{df}}_2) = 1/(1 - 1/(\tilde{\kappa} + 1))$

We plot an example of this estimator for correlated data in figure 11; note the substantial deviation. This is to be expected, as there was no claim in Carmack *et al* [28] that this estimator should work for $\mathbf{X}$ correlations.

Appendix H. Weighted risks and the MMSE estimator

In this appendix, we briefly comment on related risks and the MMSE estimator. As discussed in the main text, one might consider a weighted loss

$$L_M(\mathbf{w}) = \frac{1}{T} (\mathbf{X}\mathbf{w} - \mathbf{y})^\top \mathbf{M} (\mathbf{X}\mathbf{w} - \mathbf{y}) + \lambda \|\mathbf{w}\|^2 \quad (\text{H1})$$

for $\mathbf{M}$ a general positive-definite symmetric weighting matrix. Under our Gaussian statistical assumptions, this is equivalent to using an unweighted loss with $\mathbf{M} \leftarrow \mathbf{I}_T$, $\mathbf{K} \leftarrow \mathbf{M}^{1/2} \mathbf{K} \mathbf{M}^{1/2}$, and $\mathbf{K}' \leftarrow \mathbf{M}^{1/2} \mathbf{K}' \mathbf{M}^{1/2}$.

We now observe that the Bayesian MMSE estimator is equivalent to minimizing a weighted loss with the particular choice $\mathbf{M} = (\mathbf{K}')^{-1}$. Under our statistical assumptions, the MMSE estimator is simply given by the posterior mean:

$$\hat{\mathbf{w}}_{\text{MMSE}} = \int d\mathbf{w} \mathbf{w} p(\mathbf{w} | \mathbf{X}, \mathbf{y}). \quad (\text{H2})$$

Using our Gaussian assumptions on the data and a Gaussian prior of variance ρ^2 , the posterior is

$$p(\mathbf{w} | \mathbf{X}, \mathbf{y}) \propto p(\mathbf{X}, \mathbf{y} | \mathbf{w})p(\mathbf{w}) \quad (\text{H3})$$

$$= p(\mathbf{y} | \mathbf{X}, \mathbf{w})p(\mathbf{X})p(\mathbf{w}) \quad (\text{H4})$$

$$\propto \exp \left(-\frac{1}{2\sigma_\epsilon^2} (\mathbf{y} - \mathbf{X}\mathbf{w})^\top (\mathbf{K}')^{-1} (\mathbf{y} - \mathbf{X}\mathbf{w}) - \frac{1}{2} \text{Tr}(\mathbf{K}^{-1} \mathbf{X} \Sigma^{-1} \mathbf{X}^\top) - \frac{1}{2\rho^2} \|\mathbf{w}\|^2 \right). \quad (\text{H5})$$

Discarding $\mathbf{w}$ -independent terms and completing the square, this means that the posterior over $\mathbf{w}$ is Gaussian with mean

$$\frac{1}{T} \left(\frac{1}{T} \mathbf{X}^\top (\mathbf{K}')^{-1} \mathbf{X} + \lambda \mathbf{I} \right)^{-1} \mathbf{X}^\top (\mathbf{K}')^{-1} \mathbf{y} \quad (\text{H6})$$

and covariance

$$\frac{\sigma_\epsilon^2}{T} \left(\frac{1}{T} \mathbf{X}^\top (\mathbf{K}')^{-1} \mathbf{X} + \lambda \mathbf{I} \right)^{-1}, \quad (\text{H7})$$

where we take $\lambda = \frac{\sigma_\epsilon^2}{T\rho^2}$. Therefore, the MMSE estimator corresponds to minimizing a weighted loss with

$$\mathbf{M} = (\mathbf{K}')^{-1}. \quad (\text{H8})$$

If $\mathbf{K}' = \mathbf{K}$, this then corresponds to minimizing an unweighted loss for uncorrelated datapoints.

However, we now observe that this procedure is only possible if one has omniscient knowledge of $\mathbf{K}'$, as reliably estimating $(\mathbf{K}')^{-1}$ from samples is challenging.

Appendix I. S-transforms for certain Toeplitz covariance matrices

In this appendix, we record formulas for the spectral statistics of a few tractable and practically-relevant classes of Toeplitz covariance matrices. We direct the interested reader to work by Kühn and Sollich [70] or Basak *et al* [71] for studies of the limiting spectral properties of *empirical* autocovariance matrices.

I.1. Nearest-neighbor correlations

We begin with the simplest non-trivial example: nearest-neighbor correlations. In this case we give a self-contained analysis, which hints at some of the approaches that can be used for more general classes of Toeplitz matrices. Suppose that

$$K_{ts} = \begin{cases} 1 & t = s \\ b/2 & t = s \pm 1 \\ 0 & \text{otherwise} \end{cases} \quad (\text{I1})$$

is a symmetric tridiagonal Toeplitz matrix with diagonal elements 1 and off-diagonal elements $b/2$. We have set the diagonal elements equal to 1 without loss of generality as this is equivalent to choosing an overall scale. Tridiagonal Toeplitz matrices are exactly diagonalized by Fourier modes [72], and their eigenvalues are known to be

$$\lambda_t = 1 + b \cos \frac{\pi t}{T+1} \quad (\text{I2})$$

where $t = 1, \dots, T$. Clearly, for the matrix to be positive-definite we must have $|b| < 1$; we will assume $b > 0$ without loss of generality as the cosine term is symmetric. For a test function ϕ , we therefore have

$$\frac{1}{T} \sum_{t=1}^T \phi(\lambda_t) = \frac{1}{T} \sum_{t=1}^T \phi \left(1 + b \cos \frac{\pi t}{T+1} \right). \quad (\text{I3})$$

Recognizing this as a Riemann sum for an integral with respect to $x = t/(T+1)$, it is easy to see that

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \phi(\lambda_t) = \int_0^1 dx \phi(1 + b \cos \pi x). \quad (\text{I4})$$

Putting $\lambda = 1 + b \cos(\pi x)$, we have $\lambda = 1 + b$ when $x = 0$ and $\lambda = 1 - b$ when $x = \pi$. In this range, we can invert the relationship to find that $x = \arccos[(\lambda - 1)/b]/\pi$. Differentiating, we have $dx = -\frac{1}{\pi \sqrt{(\lambda - 1 + b)(1 + b - \lambda)}} d\lambda$. Thus, the integral becomes

$$\int_{1-b}^{1+b} d\lambda \frac{1}{\pi \sqrt{(\lambda - 1 + b)(1 + b - \lambda)}} \phi(\lambda). \quad (\text{I5})$$

This shows that for any $b > 0$ the limiting density of eigenvalues is

$$\frac{1}{\pi \sqrt{(\lambda - 1 + b)(1 + b - \lambda)}} 1_{\lambda \in [1-b, 1+b]}, \quad (\text{I6})$$

which is an arcsine distribution centered at 1 of width b . In other words, the limiting distribution is simply the pushforward of the uniform measure on $[0, 1]$ by $x \mapsto 1 + b \cos(\pi x)$, much as the distribution at finite size is the pushforward of the uniform measure on $\{1/(T+1), \dots, T/(T+1)\}$ by the same function. This is a (very) special case of more general limit theorems for the spectra of symmetric Toeplitz matrices [68].

Applying the result of this digression, we have

$$\text{df}_K^1(z) \rightarrow \int_{1-b}^{1+b} d\lambda \frac{1}{\pi \sqrt{(\lambda - 1 + b)(1 + b - \lambda)}} \frac{\lambda}{\lambda + z} \quad (\text{I7})$$

$$= 1 - \frac{z}{\sqrt{(1+z)^2 - b^2}}. \quad (\text{I8})$$

From this, we can obtain with a bit of algebra the corresponding S -transform

$$S_K(t) = \frac{\sqrt{1 - b^2 t(2 - t)} - (1 - t)}{(1 - b^2)t}. \quad (\text{I9})$$

We observe that the limit as $b \downarrow 0$ of this result gives

$$\lim_{b \downarrow 0} \text{df}_K^1(z) = \frac{1}{1 + z} \quad (\text{I10})$$

$$\lim_{b \downarrow 0} S_K(t) = 1, \quad (\text{I11})$$

which recover the expected results for the identity matrix.

I.2. Exponential correlations

Now we consider the case of exponential correlations

$$K_{ts} = e^{-|t-s|/\xi} \quad (\text{I12})$$

for some correlation length ξ , which is treated in the textbook of Potters and Bouchaud [32]. Though the eigenvectors of this matrix are not precisely Fourier modes at finite size, one can argue that the approximation error in the resolvent resulting from treating $\mathbf{K}$ as circulant becomes negligible, and from that determine the limit. In the end, one finds that

$$S_K(t) = \frac{bt + \sqrt{1 + (b^2 - 1)t^2}}{t + 1}, \quad b = \coth 1/\xi. \quad (\text{I13})$$

In this case, we should recover the identity matrix on taking $\xi \downarrow 0$, in which case we have $b \downarrow 1$. This yields $\lim_{b \downarrow 1} S_K(t) = 1$, as expected.

Appendix J. Further experiments

J.1. Exponential correlations

In this section we study the predictions of the theory for various relevant quantities in the case of both strong and weak exponential correlations. Plots of the error curves in this setting are numerous in the main text. We test this across both isotropic and structured data.

In figure 12, we plot several different curves for $\kappa, \tilde{\kappa}$ as T varies from small to large across differently structured datasets. We then plot the corresponding degrees of freedom in figure 13. Finally, we verify the duality relation holds empirically in figure 14.

J.2. Nearest-neighbor correlations

The S -transform for a correlation matrix $\mathbf{K}$ where only nearest neighbor data points are correlated can be straightforwardly obtained. We do so in section I.1. In general, the effect of correlations is quite weak unless the off-diagonal correlation element $b/2$ has b

Risk and cross validation in ridge regression with correlated samples




Figure 12. Plots of the resolutions $\kappa, \tilde{\kappa}$ under exponential correlations. Top: Isotropic data. Bottom: Anisotropic power law data with the eigenvalues of Σ having power law decay $k^{-\alpha}$, $\alpha = 1.8$. Overlaid in dashed is the expected scaling in the overparameterized regime. Left: $\xi = 10^{-2}$, essentially uncorrelated. Right: $\xi = 10^2$, strongly correlated. We fix $N = 100$ and vary T . Overlaid in dashed is the scaling expected in the overparameterized regime when $\alpha > 1$.

close to 1. In figure 15, we plot the risk curves and sources of variance in both the setting of isotropic and anisotropic data. In general, all estimators tend to agree better for the choices of b listed compared to the strongly exponentially correlated data. However, only the CorrGCV consistently correctly estimates the out-of-sample risk across T , b , λ , and σ_ϵ . Especially at small value of T , we see that only the CorrGCV estimator is the most reliable.

We further plot the resolutions in figure 16 and degrees of freedom in figure 17 across two different values of $b = 0.5, 0.95$ in both the isotropic and anisotropic case.

J.3. Power-law correlations

For power law correlations, we do not have an explicit analytic formula for S_K . We instead estimate it by interpolating df_K as a function of λ and estimating the functional inverse df_K^{-1} , as discussed in section A.2. We ensure our interpolator is compatible with autograd so that we can run algorithm 7 to get estimators for $\text{df}_2, \tilde{\text{df}}_2$ and thus for the CorrGCV.

Risk and cross validation in ridge regression with correlated samples




Figure 13. Plots of the degrees of freedom $df_1, \tilde{df}_1, df_2, \tilde{df}_2$ under exponential correlations as in figure 12. Solid lines are theory, dots with error bars are empirics. The lack of substantial error bars stems from the fact that all quantities concentrate. We find that although df_2 can be numerically sensitive, $\tilde{df}_2$ remains numerically precise even in the presence of strong correlations.

We show two examples with power-law correlations in figure 18. We further plot the resolutions in figure 19 and degrees of freedom in figure 20 across two different values of $\chi = 1.5, 0.1$ in both the isotropic and anisotropic case.

Risk and cross validation in ridge regression with correlated samples




Figure 14. Plots verifying the duality relation $\frac{\kappa\tilde{\kappa}}{\lambda} = \frac{1}{\text{df}_1}$ hold for the settings in figures 12 and 13 under exponential correlations. Dashed black lines are theory, dots with error bars are empirics. The lack of substantial error bars stems from the fact that all quantities concentrate.




Figure 15. Risks, sources of variance, and risk estimators, similar to figure 1, but with nearest neighbor correlations. (a) Isotropic covariates. (b) Structured covariates under source and capacity conditions.

Risk and cross validation in ridge regression with correlated samples




Figure 16. Plots of the resolution $\kappa, \tilde{\kappa}$ under nearest neighbor correlations. Top: Isotropic data. Bottom: Anisotropic power law data with the eigenvalues of Σ having power law decay $k^{-\alpha}$, $\alpha = 1.2$. Left: $b = 0.5$, weakly correlated. Right: $b = 0.95$, strongly correlated. We fix $N = 100$ and vary T . Solid lines are theory, dots with error bars are empirics. Overlaid in dashed is the expected scaling in the overparameterized regime when $\alpha > 1$.

Risk and cross validation in ridge regression with correlated samples




Figure 17. Plots of the the degrees of freedom $df_1, \tilde{df}_1, df_2, \tilde{df}_2$ under nearest-neighbor correlations as in figure 16. Solid lines are theory, dots with error bars are empirics. The lack of substantial error bars stems from the fact that all quantities concentrate. We find that although df_2 can be numerically sensitive, $\tilde{df}_2$ remains numerically precise even in the presence of strong correlations.




Figure 18. Risks, sources of variance, and risk estimators, similar to figure 1, but with power law correlations $\mathbb{E}[\mathbf{x}_t \cdot \mathbf{x}_s] \propto (1 + |t - s|)^{-\chi}$ with exponent $\chi = 0.7$. (a) Isotropic covariates. (b) Structured covariates under source and capacity conditions.

Risk and cross validation in ridge regression with correlated samples




Figure 19. Plots of the resolution $\kappa, \tilde{\kappa}$ under power law correlations. Top: Isotropic data. Bottom: Anisotropic power law data with the eigenvalues of Σ having power law decay $k^{-\alpha}$, $\alpha = 1.2$. Overlaid in dashed is the expected scaling in the overparameterized regime. Left: $\chi = 1.5$, weakly correlated. Right: $\chi = 0.1$, strongly correlated. We fix $N = 100$ and vary T . For strongly correlated data, we see a slight deviation between κ in theory and empirics. This is due to the low-degree polynomial interpolation estimator of $S_K(qdf_1)$ performing more poorly near $qdf_1 = 1$. This can be resolved with a more flexible interpolant.

Risk and cross validation in ridge regression with correlated samples




Figure 20. Plots of the the degrees of freedom $df_1, \tilde{df}_1, df_2, \tilde{df}_2$ under power law correlations as in figure 19. Solid lines are theory, dots with error bars are empirics. The lack of substantial error bars stems from the fact that all quantities concentrate. We find that although df_2 can be numerically sensitive, $\tilde{df}_2$ remains numerically precise even in the presence of strong correlations.

Appendix K. Experimental details

Error bars are always reported over ten runs of the same regression over different datasets. We ensembled over these datasets to calculate the $\text{Var}_{\mathbf{X}}$ term empirically. We also held out a target free of label noise to calculate the $\text{Var}_{\mathbf{X}\epsilon}$ component of the variance empirically.

We used JAX [46] to perform all the linear algebraic manipulations. All experiments were done primarily on a CPU with very little compute required. We also tested a speedup of running our code on a GPU, amounting to less than 1 GPU-day of compute usage. As mentioned in the main text, all code is available at: <https://github.com/Pehlevan-Group/S.transform>.

ORCID iDs

Alexander Atanasov  0000-0002-3338-0324Jacob A Zavatone-Veth  0000-0002-4060-1738Cengiz Pehlevan  0000-0001-9767-6063

References

- [1] Bouchaud J-P and Potters M 2003 *Theory of Financial Risk and Derivative Pricing: From Statistical Physics to Risk Management* 2nd edn (Cambridge University Press)
- [2] Mudelsee M 2014 Climate time series analysis: classical statistical and bootstrap methods *Atmospheric and Oceanographic Sciences Library* (Springer)
- [3] Williams A H and Linderman S W 2021 Statistical neuroscience in the single trial limit *Curr. Opin. Neurobiol.* **70** 193
- [4] Golub G H, Heath M and Wahba G 1979 Generalized cross-validation as a method for choosing a good ridge parameter *Technometrics* **21** 215
- [5] Craven P and Wahba G 1978 Smoothing noisy data with spline functions: estimating the correct degree of smoothing by the method of generalized cross-validation *Numer. Math.* **31** 377
- [6] Bates S, Hastie T and Tibshirani R 2024 Cross-validation: what does it estimate and how well does it do it? *J. Am. Stat. Assoc.* **119** 1434
- [7] Xu G and Huang J Z 2012 Asymptotic optimality and efficient computation of the leave-subject-out cross-validation *Ann. Stat.* **40** 3003
- [8] Gu C and Ma P 2005 Optimal smoothing in nonparametric mixed-effect models *Ann. Stat.* **33** 1357
- [9] Xu G, Shang Z and Cheng G 2018 Optimal tuning for divide-and-conquer kernel ridge regression with massive data *Proc. 35th Int. Conf. on Machine Learning, (Series Proc. Machine Learning Research)* vol, ed J Dy and A Krause (PMLR) pp 5483–91
- [10] Xu G, Shang Z and Cheng G 2019 Distributed generalized cross-validation for divide-and-conquer kernel ridge regression and its asymptotic optimality *J. Comput. Graph. Stat.* **28** 891
- [11] Hastie T, Montanari A, Rosset S and Tibshirani R J 2022 Surprises in high-dimensional ridgeless least squares interpolation *Ann. Stat.* **50** 949
- [12] Dobriban E and Wager S 2018 High-dimensional asymptotics of prediction: ridge regression and classification *Ann. Stat.* **46** 247
- [13] Advani M and Ganguli S 2016 Statistical mechanics of optimal convex inference in high dimensions *Phys. Rev. X* **6** 031034
- [14] Canatar A, Bordelon B and Pehlevan C 2021 Spectral bias and task-model alignment explain generalization in kernel regression and infinitely wide neural networks *Nat. Commun.* **12** 2914
- [15] Loureiro B, Gerbelot C, Cui H, Goldt S, Krzakala F, Mezard M and Zdeborová L 2021 Learning curves of generic features maps for realistic datasets with a teacher-student model *Advances in Neural Information Processing Systems* p 18137
- [16] Atanasov A, Zavatone-Veth J A and Pehlevan C 2026 Scaling and renormalization in high-dimensional regression *J. Stat. Mech.* **2026** 043404
- [17] Mel G and Pennington J 2021 Anisotropic random feature regression in high dimensions *Int. Conf. on Learning Representations*
- [18] Mel G and Ganguli S 2021 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 *Proc. 38th Int. Conf. on Machine Learning*, ed M Meila and T Zhang (PMLR) pp 7578–87
- [19] Bordelon B, Canatar A and Pehlevan C 2020 Spectrum dependent learning curves in kernel regression and wide neural networks *Int. Conf. on Machine Learning* (PMLR) pp 1024–34
- [20] Jacot A, Simsek B, Spadaro F, Hongler C and Gabriel F 2020 Kernel alignment risk estimator: risk prediction from training data *Advances in Neural Information Processing Systems* p 15568
- [21] Gerace F, Loureiro B, Krzakala F, Mézard M and Zdeborová L 2020 Generalisation error in learning with random features and the hidden manifold model *Int. Conf. on Machine Learning* (PMLR) pp 3452–62
- [22] Hastie T, Tibshirani R, Friedman J H and Friedman J H 2009 *The Elements of Statistical Learning: Data Mining, Inference and Prediction* vol 2 (Springer)
- [23] Montanari A and Saeed B N 2022 Universality of empirical risk minimization *Proc. Thirty 5th Conf. on Learning Theory, (Series Proc. Machine Learning Research)* vol **178**, ed P-L Loh and M Raginsky (PMLR) pp 4310–2

- [24] Hu H and Lu Y M 2023 Universality laws for high-dimensional learning with random features *IEEE Trans. Inf. Theory* **69** 1932
- [25] Bigot J, Dabo I-M and Male C 2024 High-dimensional analysis of ridge regression for non-identically distributed data with a variance profile (arXiv:2403.20200 [math.ST])
- [26] Altman N S 1990 Kernel smoothing of data with correlated errors *J. Am. Stat. Assoc.* **85** 749
- [27] Opsomer J, Wang Y and Yang Y 2001 Nonparametric regression with correlated errors *Stat. Sci.* **16** 134
- [28] Carmack P S, Spence J S and Schucany W R 2012 Generalised correlated cross-validation *J. Nonparametric Stat.* **24** 269
- [29] Rabinowicz A and Rosset S 2022 Cross-validation for correlated data *J. Am. Stat. Assoc.* **117** 718
- [30] Lukas M A 2010 Robust GCV choice of the regularization parameter for correlated data *J. Integral Equ. Appl.* **22** 519
- [31] Wang Y 1998 Smoothing spline models with correlated random errors *J. Am. Stat. Assoc.* **93** 341
- [32] Potters M and Bouchaud J-P 2020 *A First Course in Random Matrix Theory: For Physicists, Engineers and Data Scientists* (Cambridge University Press)
- [33] Burda Z, Jurkiewicz J and Waclaw B 2005 Spectral moments of correlated Wishart matrices *Phys. Rev. E* **71** 026111
- [34] Burda Z, Jurkiewicz J and Waclaw B 2005 Eigenvalue density of empirical covariance matrix for correlated samples *Acta Phys. Pol. B* **36** 2641–52
- [35] Burda Z and Jarosz A 2022 Cleaning large-dimensional covariance matrices for correlated samples *Phys. Rev. E* **105** 034136
- [36] Bach F 2024 High-dimensional analysis of double descent for linear regression with random projections *SIAM J. Math. Data Sci.* **6** 26–50
- [37] Bun J, Allez R, Bouchaud J-P and Potters M 2016 Rotational invariant estimator for general noisy matrices *IEEE Trans. Inf. Theory* **62** 7475
- [38] Patil P and LeJeune D 2024 Asymptotically free sketched ridge ensembles: risks, cross-validation and tuning *12th Int. Conf. on Learning Representations*
- [39] Caponnetto A and De Vito E 2007 Optimal rates for the regularized least-squares algorithm *Found. Comput. Math.* **7** 331
- [40] Spigler S, Geiger M and Wyart M 2020 Asymptotic learning curves of kernel methods: empirical data versus teacher–student paradigm *J. Stat. Mech.* **2020** 124001
- [41] Cui H, Loureiro B, Krzakala F and Zdeborová L 2021 Generalization error rates in kernel regression: the crossover from the noiseless to noisy regime *Advances in Neural Information Processing Systems* vol **34** p 10131
- [42] Patil P, Du J-H and Tibshirani R J 2024 Optimal ridge regularization for out-of-distribution prediction (arXiv:2404.01233 [math.ST])
- [43] Canatar A, Bordelon B and Pehlevan C 2021 Out-of-distribution generalization in kernel regression *Advances in Neural Information Processing Systems* vol **34**, ed M Ranzato, A Beygelzimer, Y Dauphin, P Liang and J W Vaughan (Curran Associates Inc) pp 12600–12
- [44] Tripuraneni N, Adlam B and Pennington J 2021 Covariate shift in high-dimensional random feature regression (arXiv:2111.08234)
- [45] Advani M S, Saxe A M and Sompolinsky H 2020 High-dimensional dynamics of generalization error in neural networks *Neural Netw.* **132** 428
- [46] Bradbury J *et al* 2018 JAX: composable transformations of Python+NumPy programs
- [47] Misiakiewicz T and Saeed B 2024 A non-asymptotic theory of kernel ridge regression: deterministic equivalents, test error, and GCV estimator (arXiv:2403.08938)
- [48] Luo K, Li Y and Sur P 2024 Roti-gcv: Generalized cross-validation for right-rotationally invariant data (arXiv:2406.11666 [math.ST])
- [49] Lu Y M, Letey M, Zavatone-Veth J A, Maiti A and Pehlevan C 2025 Asymptotic theory of in-context learning by linear attention *Proc. Natl Acad. Sci.* **122** e2502599122
- [50] Mingo J A and Speicher R 2017 *Free Probability and Random Matrices* vol 35 (Springer)
- [51] Voiculescu D V 1997 *Free Probability Theory* vol 12 (American Mathematical Society)
- [52] Nica A and Speicher R 2006 *Lectures on the Combinatorics of Free Probability* vol 13 (Cambridge University Press)
- [53] Voiculescu D V, Dykema K J and Nica A 1992 *Free Random Variables* (American Mathematical Society)
- [54] Burda Z, Janik R and Nowak M 2011 Multiplication law and S transform for non-Hermitian random matrices *Phys. Rev. E* **84** 061125
- [55] Hooft G t 1973 A planar diagram theory for strong interactions *Nucl. Phys. B* **72** 461
- [56] Weingarten D 1978 Asymptotic behavior of group integrals in the limit of infinite rank *J. Math. Phys.* **19** 999

- [57] Atanasov A, Bordelon B, Zavatone-Veth J A, Paquette C and Pehlevan C 2026 Two-point deterministic equivalence for stochastic gradient dynamics in linear models *Adv. Theor. Math. Phys.* **30** 3
- [58] Adlam B and Pennington J 2020 Understanding double descent requires a fine-grained bias-variance decomposition *Advances in Neural Information Processing Systems* vol **33** pp 11022
- [59] Lin L and Dobriban E 2021 What causes the test error? Going beyond bias-variance via anova *J. Mach. Learn. Res.* **22** 1-82
- [60] Caponnetto A and Vito E D 2005 Fast rates for regularized least-squares algorithm *Tech. Rep.* (Institution Massachusetts Institute of Technology Computer Science and Artificial Intelligence Laboratory)
- [61] Jacot A, Gabriel F and Hongler C 2018 Neural tangent kernel: convergence and generalization in neural networks *Advances in Neural Information Processing Systems* vol **31**
- [62] Kaplan J, McCandlish S, Henighan T, Brown T B, Chess B, Child R, Gray S, Radford A, Wu J and Amodei D 2020 Scaling laws for neural language models (arXiv:2001.08361)
- [63] Bahri Y, Dyer E, Kaplan J, Lee J and Sharma U 2024 Explaining neural scaling laws *Proc. Natl Acad. Sci.* **121** e2311878121
- [64] Chizat L, Oyallon E and Bach F 2019 On lazy training in differentiable programming *Advances in Neural Information Processing Systems* vol **32**
- [65] Maloney A, Roberts D A and Sully J 2022 A solvable model of neural scaling laws (arXiv:2210.16859)
- [66] Bordelon B, Atanasov A and Pehlevan C 2024 A dynamical model of neural scaling laws (arXiv:2402.01092)
- [67] Defilippis L, Loureiro B and Misiakiewicz T 2024 Dimension-free deterministic equivalents for random feature regression (arXiv:2405.15699)
- [68] Böttcher A, Grudsky S, Maksimenko E A and Unterberger J 2009 The first order asymptotics of the extreme eigenvectors of certain Hermitian Toeplitz matrices *Integral Equ. Oper. Theory* **63** 165
- [69] Zavatone-Veth J A and Pehlevan C 2023 Learning curves for deep structured Gaussian feature models *Advances in Neural Information Processing Systems*
- [70] Kühn R and Sollich P 2012 Spectra of empirical auto-covariance matrices *Europhys. Lett.* **99** 20008
- [71] Basak A, Bose A and Sen S 2014 Limiting spectral distribution of sample autocovariance matrices *Bernoulli* **20** 1234
- [72] Noschese S, Pasquini L and Reichel L 2013 Tridiagonal Toeplitz matrices: properties and novel applications *Numer. Linear Algebra Appl.* **20** 302