

Capturing the learning curves of generic features maps for realistic data sets with a teacher-student model

Bruno Loureiro^{*1}, Cédric Gerbelot^{†2}, Hugo Cui³, Sebastian Goldt⁴,
Florent Krzakala¹, Marc Mézard², and Lenka Zdeborová³

¹IdePHICS laboratory, École Fédérale Polytechnique de Lausanne (EPFL), Switzerland

²Laboratoire de Physique de l'École Normale Supérieure, Université PSL, CNRS, Sorbonne Université, Université Paris-Diderot, Sorbonne Paris Cité, Paris, France

³SPOC laboratory, Ecole Fédérale Polytechnique de Lausanne (EPFL), Switzerland

⁴International School of Advanced Studies (SISSA), Trieste, Italy

Abstract

Teacher-student models provide a powerful framework in which the typical case performance of high-dimensional supervised learning tasks can be studied in closed form. In this setting, labels are assigned to data - often taken to be Gaussian i.i.d. - by a *teacher model*, and the goal is to characterise the typical performance of the *student model* in recovering the parameters that generated the labels. In this manuscript we discuss a generalisation of this setting where the teacher and student can act on different spaces, generated with fixed, but generic feature maps. This is achieved via the rigorous study of a high-dimensional Gaussian covariate model. Our contribution is two-fold: First, we prove a rigorous formula for the asymptotic training loss and generalisation error achieved by empirical risk minimization for this model. Second, we present a number of situations where the learning curve of the model captures the one of a *realistic data set* learned with kernel regression and classification, with out-of-the-box feature maps such as random projections or scattering transforms, or with pre-learned ones - such as the features learned by training multi-layer neural networks. We discuss both the power and the limitations of the Gaussian teacher-student framework as a typical case analysis capturing learning curves as encountered in practice on real data sets.

^{*}bruno.loureiro@epfl.ch

[†]cedric.gerbelot@ens.fr

Contents

1	Introduction	3
2	Main technical results	4
3	Applications of the Gaussian model	7
3.1	Random kitchen sinks with Gaussian data	8
3.2	Kernel methods with Gaussian data	9
3.3	GAN-generated data and learned teachers	10
3.4	Learning from real data sets	11
A	Main result from the replica method	14
A.1	Replica computation of the free energy	15
A.2	Ridge regression and fixed weights	17
A.3	Taking the $\beta \rightarrow \infty$ limit	18
A.4	Saddle-point equations	18
A.5	Examples	19
A.6	Relation to previous models	20
B	Power law scaling for kernel regression	23
B.1	Self-consistent equations for the generalization error	23
B.2	Limit of infinite student dimension	24
B.3	Infinite sample limit and the scaling of the generalisation error	24
C	Rigorous proof of the Main result	26
C.1	Necessary assumptions	26
C.2	Theoretical toolbox	27
C.3	Determining a candidate primary problem, auxiliary problem and its solution.	31
C.4	Study of the scalar equivalent problem : geometry and asymptotics.	40
C.5	Back to the original problem : proof of Theorem 1 and 2	45
C.6	Relaxing assumption (B1)	48
C.7	The 'vanilla' teacher-student scenario	48
D	Equivalence replica-Gordon	50
D.1	Solution for separable loss and ridge regularization	50
D.2	Matching with Replica equations	51
E	Details on the simulations	53
E.1	Ridge regression on real data	53
E.2	Binary classification on GAN generated data	56
F	Ridge regression with linear teachers	58

1 Introduction

Teacher-student models are a popular framework to study the high-dimensional asymptotic performance of machine learning problems with synthetic data, and have been the subject of intense investigations spanning three decades [1, 2, 3, 4, 5, 6, 7]. In the wake of understanding the limitations of classical statistical learning approaches [8, 9, 10] this direction is witnessing a renewal of interest [11, 12, 10, 13, 10, 14, 15]. In this paper, we extend this line of work, and consider models taking into account, to some extent, the structure and correlations in the data and labels.

We discuss a theory capturing the cases of teachers and students defined by elaborate, deterministic, random, or even learned features maps for the data.

This generic setting is relevant in many situations of practical interest in machine learning, such as kernel regression, the analysis of feature maps - random projections [16], neural tangent kernels [17], scattering transforms [18] - as well as the analysis of transfer learning performance on data generated by generative adversarial networks [19]. The broad spectrum of our analysis is possible because, for all these scenarios, the important correlations between the teacher and the student can be asymptotically captured by a (*jointly*)-*Gaussian covariate model*, as we will discuss in detail below. We derive a rigorous solution of this covariate model, and present some of its applications to realistic (and in some cases real) data sets. We also discuss its limit of applicability, by showing how the Gaussian assumption breaks down in some concrete situations.

Model definition and basic intuition — Our main technical result concerns the following Gaussian covariate model for two vectors $\mathbf{u} \in \mathbb{R}^p$ and $\mathbf{v} \in \mathbb{R}^d$, with correlation matrices $\Psi \in \mathbb{R}^{p \times p}$, $\Omega \in \mathbb{R}^{d \times d}$ and $\Phi \in \mathbb{R}^{p \times d}$, from which we draw n independent samples:

$$\begin{bmatrix} \mathbf{u} \\ \mathbf{v} \end{bmatrix} \in \mathbb{R}^{p+d} \sim \mathcal{N} \left(0, \begin{bmatrix} \Psi & \Phi \\ \Phi^\top & \Omega \end{bmatrix} \right). \quad (1.1)$$

Denoting $U \in \mathbb{R}^{n \times p}$ & $V \in \mathbb{R}^{n \times d}$ the matrices obtained from the concatenation of the n samples, we define the teacher and the student as follows:

- The *labels* \mathbf{y} are generated by a *teacher* function that is only using the vectors \mathbf{u} :

$$\mathbf{y} = \mathbf{f}_0 \left(\frac{1}{\sqrt{p}} U \boldsymbol{\theta}_0 \right), \quad (1.2)$$

where $\mathbf{f}_0 : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is a function that may include randomness such as, for instance, an additive Gaussian noise, and $\boldsymbol{\theta}_0 \in \mathbb{R}^p$ is a vector of teacher-weights which can be either random or deterministic. We will focus on a factorised \mathbf{f}_0 (i.e. the same scalar function f_0 is applied to each sample).

- *Learning* is performed by the *student* via empirical risk minimization that has access only to the vectors \mathbf{v}^{μ} :

$$\hat{\mathbf{w}} = \arg \min_{\mathbf{w} \in \mathbb{R}^d} \left[g \left(\frac{1}{\sqrt{d}} V \mathbf{w}, \mathbf{y} \right) + r(\mathbf{w}) \right], \quad (1.3)$$

where \mathbf{w} are the student-weights, and r & g are proper, convex functions of $\mathbf{w} \in \mathbb{R}^d$ (e.g. g can be a logistic or a square loss and r a ℓ_p ($p=1, 2$) penalty).

Main contributions —

- Our main technical contribution is a rigorous closed-form characterisation of the properties of the estimator $\hat{\mathbf{w}}$, and its corresponding training and generalisation error in the high-dimensional limit. This is achieved using Gaussian comparison inequalities [20].
- We also derive the expression using the replica method from statistical physics [21]. We believe this relationship to be of additional interest given the wide range of applications of the replica approach in machine learning and computer science [22].

- We apply the formula to a range of situations of interests with both synthetic and realistic, correlated data. We show that the theoretical prediction captures the learning curves in non-trivial cases, e.g. when input data are generated using a trained generative adversarial network, while extracting the feature maps from a neural network trained on real data.
- We empirically show that, for ridge regression, the asymptotic formula can be applied *directly* to real data sets, even though the Gaussian hypothesis is not satisfied. This illustrates the power of typical case analysis for concrete practical situations.

Related works — The Gaussian covariate model (1.1) contains the vanilla teacher-student model as a special case where one takes \mathbf{u} and \mathbf{v} *identical*, with unique covariance matrix Ω . This special case has been extensively studied in the statistical physics community using the heuristic replica method [23, 24, 1, 2, 3]. Many recent rigorous results for such models can be rederived as a special case of our formula, e.g. [11, 12, 25, 10, 13, 10, 26, 27, 14, 15, 28]. Numerous of these works are based on the same proof technique as we employed here: the Gordon’s Gaussian min-max inequalities [20, 29, 30]. Closer to our work is the recent work of [31] on the random feature model. For ridge regression, there are also precise predictions thanks to random matrix theory [32, 12, 33, 34, 35]. A related set of results were obtained in [36] for orthogonal random matrix models. The main technical novelty of our proof is the handling of a generic loss and regularisation, not only ridge, representing convex empirical risk minimization, for both classification and regression, with the generic correlation structure of the model (1.1).

The main interest of model (1.1) is that it goes beyond the vanilla case and spans a variety of new interesting scenarios. As we shall see, the covariance matrices Ψ , Φ , and Ω can represent different aspects of the data-generative process and learning model. They can account for learned generative models of the input data, such as generative neural networks, but they can also model the process generating the labels, such as a neural network learned from data, or a given kernel or feature map (detailed mappings are given later). The student (1.3) then corresponds to the last layer of the learning model. A number of recent results conjectured via the heuristic replica method then follow from our proof. For instance, the asymptotic analysis of kernel ridge regression [37], of margin-based classification [38] and of a Gaussian covariate models used to study generative distributions [39], which is a particular cases of (1.1) and (1.2) when $\Psi = \mathbf{I}_d$ and $\boldsymbol{\theta}_0 \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_p)$. See Appendix A.6 for the details on these connections. Other examples include models of the double descent phenomenon [40].

2 Main technical results

Exact asymptotics — The results that follow are valid in the *high-dimensional* limit, characterised by the dimensions of the problem n, p, d going to infinity with fixed ratios:

$$\alpha \equiv \frac{n}{d}, \quad \text{and} \quad \gamma \equiv \frac{p}{d}. \quad (2.1)$$

We define the averaged training loss for a given estimator \mathbf{w} as the expectation of the risk:

$$\mathcal{E}_{\text{train}}(\mathbf{w}) \equiv \frac{1}{n} \mathbb{E} \left[g \left(\frac{1}{\sqrt{d}} V \mathbf{w}, \mathbf{y} \right) + r(\mathbf{w}) \right]. \quad (2.2)$$

where the argument \mathbf{y} in the loss function $g(\cdot, \mathbf{y})$ is understood as a parameter, not a variable. To evaluate the performance, we consider a prediction function $f : \mathbb{R} \rightarrow \mathbb{R}$ used to generate new labels from the learned model, and a performance measure $\hat{g} : \mathbb{R} \rightarrow \mathbb{R}$. A fresh sample then takes the form of a feature vector \mathbf{v}_{new} and the associated label y_{new} originating from a feature \mathbf{u}_{new} . The averaged generalisation error is then defined as:

$$\mathcal{E}_{\text{gen}}(\mathbf{w}) \equiv \mathbb{E} \left[\hat{g}(\hat{f}(\mathbf{v}_{\text{new}}^\top \mathbf{w}), y_{\text{new}}) \right] \equiv \mathbb{E} \left[\hat{g}(\hat{f}(\mathbf{v}_{\text{new}}^\top \mathbf{w}), \mathbf{f}_0(\mathbf{u}_{\text{new}}^\top \boldsymbol{\theta}_0)) \right]. \quad (2.3)$$

From these definitions, we shall deduce asymptotics for the variables $\mathbf{u}_{\text{new}}^\top \boldsymbol{\theta}_0$ and $\mathbf{v}_{\text{new}}^\top \mathbf{w}$. Given an instance of $\boldsymbol{\theta}_0$ and \mathbf{w} , the tuple $\left(\frac{1}{\sqrt{p}} \mathbf{u}_{\text{new}}^\top \boldsymbol{\theta}_0, \frac{1}{\sqrt{d}} \mathbf{v}_{\text{new}}^\top \mathbf{w}\right)$ is a bivariate Gaussian with covariance:

$$\begin{bmatrix} \frac{1}{p} \boldsymbol{\theta}_0^\top \Psi \boldsymbol{\theta}_0 & \frac{1}{\sqrt{dp}} (\Phi^\top \boldsymbol{\theta}_0)^\top \mathbf{w} \\ \frac{1}{\sqrt{dp}} (\Phi^\top \boldsymbol{\theta}_0)^\top \mathbf{w} & \frac{1}{d} \mathbf{w}^\top \Omega \mathbf{w} \end{bmatrix}. \quad (2.4)$$

We thus define the following overlaps, that will play a fundamental role in the analysis:

$$\rho = \frac{1}{p} \boldsymbol{\theta}_0^\top \Psi \boldsymbol{\theta}_0, \quad m = \frac{1}{\sqrt{dp}} (\Phi^\top \boldsymbol{\theta}_0)^\top \mathbf{w}, \quad q = \frac{1}{d} \mathbf{w}^\top \Omega \mathbf{w}, \quad \chi = \frac{1}{d} \boldsymbol{\theta}_0^\top \Phi \Omega^{-1} \Phi^\top \boldsymbol{\theta}_0. \quad (2.5)$$

Our main result is that the overlaps, the training loss and the generalisation error can be computed explicitly in the high-dimensional limit defined above. We derive this result in two different ways: a) heuristically, using the the replica method from statistical physics in Appendix A, and b) rigorously, via the Gordon minimax approach detailed in Appendix C.

First, let's define a scalar optimization problem that will be used to state the asymptotic behaviour of (1.2-1.3):

Definition 1. (*Scalar potentials/replica free energy*) Define the following functions of the scalar variables $\tau_1 > 0, \tau_2 > 0, \kappa \geq 0, \eta \geq 0, \nu, m$:

$$\begin{aligned} \mathcal{L}_g(\tau_1, \kappa, m, \eta) &= \frac{1}{n} \mathbb{E} \left[\mathcal{M}_{\frac{\tau_1}{\kappa} g(\cdot, \mathbf{y})} \left(\frac{m}{\sqrt{\rho}} \mathbf{s} + \eta \mathbf{h} \right) \right], \\ \mathcal{L}_r(\tau_2, \eta, \nu, \kappa) &= \frac{1}{d} \mathbb{E} \left[\mathcal{M}_{\frac{\eta}{\tau_2} r(\Omega^{-1/2} \cdot)} \left(\frac{\eta}{\tau_2} (\nu \mathbf{t} + \kappa \mathbf{g}) \right) \right], \end{aligned} \quad (2.6)$$

where $\mathbf{s}, \mathbf{h} \sim \mathcal{N}(0, I_n)$ and $\mathbf{g} \sim \mathcal{N}(0, I_d)$ are random vectors independent of the other quantities, $\mathbf{t} = \Omega^{-1/2} \Phi^\top \boldsymbol{\theta}_0$, $\mathbf{y} = \mathbf{f}_0(\sqrt{\rho} \mathbf{s})$, and \mathcal{M} denotes the Moreau envelope of a target function.

From these quantities define the following potential:

$$\begin{aligned} \mathcal{E}(\tau_1, \tau_2, \kappa, \eta, \nu, m) &= \frac{\kappa \tau_1}{2} - \frac{\eta \tau_2}{2} + m \nu \sqrt{\gamma} - \frac{\tau_2}{2\eta} \frac{m^2}{\rho} \\ &\quad - \frac{\eta}{2\tau_2} (\nu^2 \chi + \kappa^2) + \alpha \mathcal{L}_g(\tau_1, \kappa, m, \eta) + \mathcal{L}_r(\tau_2, \eta, \nu, \kappa). \end{aligned} \quad (2.7)$$

Under Assumption (C.1), the previously defined quantities all admit finite limits when $n, p, d \rightarrow \infty$.

Proof: This follows directly from Lemma 5.

The next lemma characterizes important properties of the "potential" function $\mathcal{E}(\tau_1, \tau_2, \kappa, \eta, \nu, m)$:

Lemma 1. (*Geometry and minimizers of \mathcal{E}*) The function $\mathcal{E}(\tau_1, \tau_2, \kappa, \eta, \nu, m)$ is jointly convex in (m, η, τ_1) and jointly concave in (ν, κ, τ_2) , and the optimization problem

$$\min_{m, \eta, \tau_1} \max_{\kappa, \nu, \tau_2} \mathcal{E}(\tau_1, \tau_2, \kappa, \eta, \nu, m) \quad (2.8)$$

has a unique solution $(\tau_1^*, \tau_2^*, \kappa^*, \eta^*, \nu^*, m^*)$ on $\text{dom}(\mathcal{E})$.

Proof: see Appendix C.4. The optimality condition of problem (2.8) yields the set of self-consistent fixed point equations given in Lemma 12 of Appendix C. Finally, define the following variables:

$$\mathbf{w}^* = \Omega^{-1/2} \text{prox}_{\frac{\eta^*}{\tau_2^*} r(\Omega^{-1/2} \cdot)} \left(\frac{\eta^*}{\tau_2^*} (\nu^* \mathbf{t} + \kappa^* \mathbf{g}) \right), \quad \mathbf{z}^* = \text{prox}_{\frac{\tau_1^*}{\kappa^*} g(\cdot, \mathbf{y})} \left(\frac{m^*}{\sqrt{\rho}} \mathbf{s} + \eta^* \mathbf{h} \right). \quad (2.9)$$

where prox denotes the proximal operator. With these definitions, we can now state our main result:

Theorem 1. (Training loss and generalisation error) Under Assumption (C.1), there exist constants $C, c, c' > 0$ such that, for any optimal solution $\hat{\mathbf{w}}$ to (1.3), the training loss and generalisation error respectively defined by equations (2.2) and (2.3) verify, for any $0 < \epsilon < c'$:

$$\mathbb{P}(|\mathcal{E}_{\text{train}}(\hat{\mathbf{w}}) - \mathcal{E}_{\text{train}}^*| \geq \epsilon) \leq \frac{C}{\epsilon} e^{-c\epsilon^2}, \quad (2.10)$$

$$\mathbb{P}\left(\left|\mathcal{E}_{\text{gen}}(\hat{\mathbf{w}}) - \mathbb{E}_{\omega, \xi} \left[\hat{g}(f_0(\omega), \hat{f}(\xi)) \right] \right| \geq \epsilon\right) \leq \frac{C}{\epsilon} e^{-c\epsilon^2},$$

where $\mathcal{E}_{\text{train}}^*$ is defined as follows:

$$\mathcal{E}_{\text{train}}^* = \frac{1}{n} \mathbb{E} [g(\mathbf{z}^*, \mathbf{y})] + \frac{1}{\alpha d} \mathbb{E} [r(\mathbf{w}^*)], \quad (2.11)$$

and the random variables (ω, ξ) are jointly Gaussian with covariance

$$(\omega, \xi) \sim \mathcal{N}\left(0, \begin{bmatrix} \rho & m^* \\ m^* & q^* \end{bmatrix}\right), \quad q^* = (\eta^*)^2 + \frac{(m^*)^2}{\rho}. \quad (2.12)$$

Proof: see Appendix C.5. Note that the regularisation may be removed to evaluate the training loss. A more generic result, aiming directly at the estimator $\hat{\mathbf{w}}$, can also be stated:

Theorem 2. Under Assumption (C.1), for any optimal solution $\hat{\mathbf{w}}$ to (1.3), denote $\hat{\mathbf{z}} = \frac{1}{\sqrt{d}} V \hat{\mathbf{w}}$. Then, there exist constants $C, c, c' > 0$ such that, for any Lipschitz function $\phi_1 : \mathbb{R}^d \rightarrow \mathbb{R}$, and separable, pseudo-Lipschitz function $\phi_2 : \mathbb{R}^n \rightarrow \mathbb{R}$ and any $0 < \epsilon < c'$:

$$\mathbb{P}\left(\left|\phi_1\left(\frac{\hat{\mathbf{w}}}{\sqrt{d}}\right) - \mathbb{E}\left[\phi_1\left(\frac{\mathbf{w}^*}{\sqrt{d}}\right)\right]\right| \geq \epsilon\right) \leq \frac{C}{\epsilon} e^{-c\epsilon^2}, \quad (2.13)$$

$$\mathbb{P}\left(\left|\phi_2\left(\frac{\hat{\mathbf{z}}}{\sqrt{n}}\right) - \mathbb{E}\left[\phi_2\left(\frac{\mathbf{z}^*}{\sqrt{n}}\right)\right]\right| \geq \epsilon\right) \leq \frac{C}{\epsilon} e^{-c\epsilon^2}. \quad (2.14)$$

Proof: see Appendix C.5. Concentration still holds for a larger class of functions $\phi_{1,2}$, but exponential rates are lost. This is discussed in Appendix C.1.

Factorised loss and ℓ_2 regularisation — A more explicit form of the scalar optimisation problem (2.7) arises when the loss g is factorisable & the regularisation is the ℓ_2 penalty, as shall be the case for the remainder of this paper:

Corollary 1. (Factorised loss and ℓ_2 regularisation) For a factorised g and an ℓ_2 regularisation with parameter $\lambda \geq 0$, letting $G = (\frac{\eta}{\tau_2} \Omega + \lambda \mathbb{I}_d)^{-1}$, the function $\mathcal{E}(\tau_1, \tau_2, \kappa, \eta, \nu, m)$ takes the form

$$\begin{aligned} \mathcal{E}(\tau_1, \tau_2, \kappa, \eta, \nu, m) &= \frac{\kappa \tau_1}{2} - \frac{\eta \tau_2}{2} + m \nu \sqrt{\gamma} - \frac{\tau_2}{2\eta} m^2 + \alpha \mathbb{E}[\mathcal{M}_{\frac{\tau_1}{\kappa} g(\cdot, \mathbf{y})}(ms + \eta h)] - \frac{1}{2d} \frac{\nu^2}{\rho} \boldsymbol{\theta}_0^\top \Phi G \Phi^\top \boldsymbol{\theta}_0 \\ &\quad - \frac{1}{2d} \kappa^2 \text{tr}(G \Omega), \end{aligned} \quad (2.15)$$

where s, h are independent standard normal random variables. Additionally, the optimality condition prescribing the unique solution $(\tau_1^*, \tau_2^*, \kappa^*, \eta^*, \nu^*, m^*)$ to the convex-concave min-max problem:

$$\mathcal{E}^* = \min_{m, \eta, \tau_1} \max_{\kappa, \nu, \tau_2} \mathcal{E}(\tau_1, \tau_2, \kappa, \eta, \nu, m) \quad (2.16)$$

matches the one obtained with the heuristic replica method.

Proof: see Appendix D.

Computing the errors — As shown in Theorem 1, the asymptotic training loss and generalisation error are fully characterised by a convex low-dimensional minimisation problem given by eq. (2.8). Since the potential \mathcal{E} has a unique minimiser, we can simply look for critical points. Focusing on the factorised case of Corollary 1, the zero-gradient condition yields a set of self-consistent equations. In the notations of Definition (1) and Corollary (1):

$$\begin{cases} \tau_1 = \frac{\kappa}{d} \text{tr}(G\Omega) \\ m = \frac{1}{d} \frac{\nu}{\sqrt{\gamma}} \boldsymbol{\theta}_0 \Phi G \Phi^\top \boldsymbol{\theta}_0 \\ \eta^2 + \frac{m^2}{\rho} = \frac{1}{d} (\nu^2 \mathbf{t}^\top \Omega G^2 \Omega \mathbf{t} + \kappa^2 \text{tr}(G^2 \Omega^2)) \end{cases}, \quad \begin{cases} \nu \sqrt{\gamma} = \alpha \frac{\kappa}{\sqrt{\rho \tau_1}} (\mathbb{E}[sz(\frac{\tau_1}{\kappa}, m, h)] - \frac{m}{\sqrt{\rho}} \mathbb{E}[z'(\frac{\tau_1}{\kappa}, m, h)]) \\ \frac{\tau_2}{\eta} = \alpha \frac{\kappa}{\tau_1} (1 - \mathbb{E}[z'(\frac{\tau_1}{\kappa}, m, h)]) \\ \tau_1^2 = \alpha \mathbb{E}[(\frac{m}{\sqrt{\rho}} s + \eta h - z(\frac{\tau_1}{\kappa}, m, \eta))^2] \end{cases} \quad (2.17)$$

where we defined the scalar random functions $z(\frac{\tau_1}{\kappa}, m, h) = \text{prox}_{\frac{\kappa}{\tau_1} g(\cdot, f_0(\sqrt{\rho}s))}(\frac{m}{\sqrt{\rho}} s + \eta h)$ and $z'(\frac{\tau_1}{\kappa}, m, h) = \text{prox}'_{\frac{\kappa}{\tau_1} g(\cdot, f_0(\sqrt{\rho}s))}(\frac{m}{\sqrt{\rho}} s + \eta h)$ as the first derivative of the proximal operator. For a given problem defined by the parameters $(f_0, \hat{f}, g, r, \boldsymbol{\theta}_0, \Omega, \Phi, \Psi)$, one can easily iterate the above equations to find the fixed point $(\kappa^*, \nu^*, \tau_1^*, \tau_2^*, m^*, \eta^*)$ which appears in the expression for the training and generalisation errors $(\mathcal{E}_{\text{train}}^*, \mathcal{E}_{\text{gen}}^*)$, see eq. (2.10). Two cases are of particular relevance for the experiments that follow. The first is the case of *ridge regression*, in which $f_0(x) = \hat{f}(x)$ and both the loss g and the performance measure \hat{g} are taken to be the *mean-squared error* $\text{mse}(y, \hat{y}) = \frac{1}{2}(y - \hat{y})^2$, and the asymptotic errors are given by the simple closed-form expression:

$$\mathcal{E}_{\text{gen}}^* = \rho + q^* - 2m^*, \quad \mathcal{E}_{\text{train}}^{*, \text{unreg}} = \frac{\mathcal{E}_{\text{gen}}^*}{(1 + \tau_1^*/\kappa^*)^2}, \quad (2.18)$$

where $q^* = (\eta^*)^2 + (m^*)^2/\rho$ as defined in eq. (2.12). The second case of interest is the one of a binary classification task, for which $f_0(x) = \hat{f}(x) = \text{sign}(x)$, and we choose the performance measure to be the *classification error* $\hat{g}(y, \hat{y}) = \mathbb{P}(y \neq \hat{y})$. In the same notation as before, the asymptotic generalisation error in this case reads:

$$\mathcal{E}_{\text{gen}}^* = \frac{1}{\pi} \cos^{-1} \left(\frac{m^*}{\sqrt{\rho q^*}} \right), \quad (2.19)$$

while the training loss $\mathcal{E}_{\text{train}}^*$ depends on the choice of g - which we will take to be the logistic loss $g(y, x) = \log(1 + e^{-xy})$ in all of the binary classification experiments. An out-of-the-box iterator for all the cases studied hereafter is provided in the GitHub repository for this manuscript [41].

3 Applications of the Gaussian model

Learning with features maps — We now discuss how the theorems above can be applied to characterise the learning curves, i.e. the dependence of the training and test errors on the number of samples, for a range of concrete problems of interest. Let $\{\mathbf{x}^\mu\}_{\mu=1}^n$ denote a data set with n independent samples on $\mathcal{X} \subset \mathbb{R}^D$. We consider here the cases in which the variables \mathbf{u}, \mathbf{v} are given by (potentially) elaborated features of \mathbf{x} , i.e. $\mathbf{u} = \boldsymbol{\varphi}_t(\mathbf{x})$ and $\mathbf{v} = \boldsymbol{\varphi}_s(\mathbf{x})$ for given feature maps $\boldsymbol{\varphi}_t : \mathcal{X} \rightarrow \mathbb{R}^p$ and $\boldsymbol{\varphi}_s : \mathcal{X} \rightarrow \mathbb{R}^d$, see Fig. 1. Our strategy is then to use the following conjecture:

Conjecture 1. (*Gaussian equivalent model*) For a wide class of data distributions $\{\mathbf{x}^\mu\}_{\mu=1}^n$, and features maps $\mathbf{u} = \boldsymbol{\varphi}_t(\mathbf{x}), \mathbf{v} = \boldsymbol{\varphi}_s(\mathbf{x})$, the generalisation and training errors of (1.3) are asymptotically captured by the equivalent Gaussian model (1.1), where $[\mathbf{u}, \mathbf{v}]$ are jointly Gaussian variables, and thus by the closed-form expressions of Theorem 1.

Conjecture 1 holds naturally in the case in which \mathbf{x} are Gaussian variables and the feature maps $(\boldsymbol{\varphi}_t, \boldsymbol{\varphi}_s)$ preserve the Gaussianity, for example linear features. In particular, this is the case for $\mathbf{u} = \mathbf{v} = \mathbf{x}$, which is the widely-studied vanilla teacher-student model. Beyond this situation, a similar Gaussian conjecture has been discussed in a series of recent works, and some proved partial results in this direction [12, 11, 27, 42, 43, 39, 31,

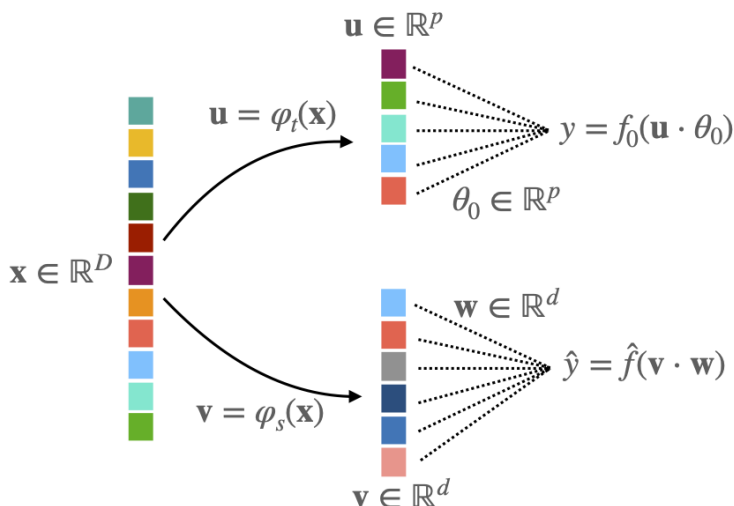


Figure 1: The approach used for capturing realistic learning curves: Given a data set $\{\mathbf{x}^\mu\}_{\mu=1}^n$, teacher $\mathbf{u} = \varphi_t(\mathbf{x})$ and student maps $\mathbf{v} = \varphi_s(\mathbf{x})$, we assume $[\mathbf{u}, \mathbf{v}]$ to be jointly Gaussian random variables and apply the results of the Gaussian covariate model (1.1).

44]. In particular, [39] studied a special case of the Gaussian model (corresponding to $\varphi_t = \text{id}$ here), and proved a Gaussian equivalence theorem (GET) for feature maps φ_s given by single-layer neural networks with fixed weights. They showed that for Gaussian data $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_D)$, feature maps of the form $\mathbf{v} = \sigma(\mathbf{W}\mathbf{x})$ (with some technical restriction on the weights) led to the jointly-Gaussian property for the two scalars $(\mathbf{v} \cdot \mathbf{w}, \mathbf{u} \cdot \boldsymbol{\theta}_0)$ for *almost* any vector \mathbf{w} . Remarkably, [44] showed that the joint Gaussianity is actually satisfied for $\hat{\mathbf{w}}$, establishing the universality of the Gaussian covariate model (1.1) for random feature maps.

We expect Conjecture 1 to hold in a much broader context than what is discussed in the above papers. Authors of [39] provided strong empirical evidence supporting the conjecture for data coming from pre-trained generative networks. However, their stringent assumptions on random teacher weights limited the scope of applications to unrealistic label models. A related line of work discussed similar universality through the lens of random matrix theory [45, 46, 47]. In particular [48] showed that, in our notations, vectors $[\mathbf{u}, \mathbf{v}]$ obtained from Gaussian inputs $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_D)$ with Lipschitz feature maps satisfy, in the authors terminology, a concentrated property. In this case, again, one can expect the two scalars $(\mathbf{v} \cdot \mathbf{w}, \mathbf{u} \cdot \boldsymbol{\theta}_0)$ to be jointly Gaussian.

The rest of this manuscript is devoted to a discussion of cases - some rather surprising - for which Conjecture 1 seems valid. Our aim is to bring forward these cases and to call for a more formal analysis of the range of validity of Conjecture 1, which as we show in this manuscript is clearly larger than what can be established from the existing literature.

3.1 Random kitchen sinks with Gaussian data

A first application is to choose random feature maps, as in random kitchen sinks [16], where $\varphi_s(\mathbf{x}) = \sigma(\mathbf{F}\mathbf{x})$ for a random matrix \mathbf{F} and a chosen scalar function σ acting component-wise. This model has seen a surge of interest recently, and a sharp asymptotic analysis was provided in the particular case of uncorrelated Gaussian data $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_D)$ and $\varphi_t = \text{id}$ in [11, 12] for ridge regression and generalised by [42, 44] for generic convex losses. Both results can be framed as a Gaussian covariate model with:

$$\Psi = \mathbf{I}_p, \quad \Phi = \kappa_1 \mathbf{F}^\top, \quad \Omega = \kappa_0^2 \mathbf{1}_d \mathbf{1}_d^\top + \kappa_1^2 \frac{\mathbf{F}\mathbf{F}^\top}{d} + \kappa_\star^2 \mathbf{I}_d, \quad (3.1)$$

where $\mathbf{1}_d \in \mathbb{R}^d$ is the all-ones vector and the constants $(\kappa_0, \kappa_1, \kappa_\star)$ are related to the non-linearity σ as:

$$\kappa_0 = \mathbb{E}_{z \sim \mathcal{N}(0,1)} [\sigma(z)], \quad \kappa_1 = \mathbb{E}_{z \sim \mathcal{N}(0,1)} [z\sigma(z)], \quad \kappa_\star = \sqrt{\mathbb{E}_{z \sim \mathcal{N}(0,1)} [\sigma(z)^2] - \kappa_0^2 - \kappa_1^2}. \quad (3.2)$$




Figure 2: Learning in kernel space: Teacher & student live in the same (Hilbert) feature space $\mathbf{v} = \mathbf{u} \in$ with $d \gg n$, and the performance only depends on the relative decay between the student spectrum $\omega_i = d i^{-2}$ (the capacity) and the teacher weights in feature space $\theta_{0i}^2 \omega_i = d i^{-a}$ (the source). Left: a task with sign teacher (in kernel space), fitted with a max-margin support vector machine (logistic regression with vanishing regularisation [54]). Right: a task with linear teacher (in kernel space) fitted via kernel ridge regression with vanishing regularisation. Points are simulation that matches the theory (lines).

Note, however, that our present framework can accommodate more involved random sinks models, such as when the teacher features are also a random feature model or multi-layer random architectures.

3.2 Kernel methods with Gaussian data

Another direct application of our formalism is to kernel methods. Kernel methods admit a dual representation in terms of optimization over feature space [49]. The connection is given by Mercer's theorem, which provides an eigen-decomposition of the kernel and of the target function in the feature basis, effectively mapping kernel regression to a teacher-student problem on feature space. The classical way of studying the performance of kernel methods [50, 51] is then to directly analyse the performance of convex learning in this space. In our notation, the teacher and student feature maps are equal, and we thus set

$$p = d, \quad \Psi = \Phi = \Omega = \text{diag}(\omega_i) \quad (3.3)$$

where ω_i are the eigenvalues of the kernel and we take the teacher weights θ_0 to be the decomposition of the target function in the kernel feature basis.

There are many deep results in classical learning theories on this problem for the case of ridge regression (where the teacher is usually called "the source" and the eigenvalues of the kernel matrix the "capacity", see e.g. [50, 52]). However, these are worst case approaches, where no assumption is made on the true distribution of the data. Most of these results are also restricted to ridge regression. In contrast, here we follow a *typical case* analysis, assuming Gaussianity in feature space. Through Theorem 2, this allow us to go beyond the restriction of the ridge loss. An example of application to this scenario is shown in Fig. 2.

For the particular case of kernel ridge regression, Theorem 2 provides a rigorous proof of the formula conjectured in [37]. In Appendix A.6 we present an explicit mapping to their results, as well as a discussion of the asymptotic performance of kernel ridge regression as discussed in [53, 37]. Our theorem puts these works on rigorous grounds, and extends them to kernel tasks with more general losses.




Figure 3: generalisation classification error, eq. (2.19) (top) and (unregularised) training loss, eq. (2.2) (bottom) as a function of the sample complexity $\alpha = n/d$ for logistic regression on a learned feature map trained on dcGAN-generated CIFAR10-like images labelled by a teacher fully-connected neural network (see Appendix E.1 for architecture details), with vanishing ℓ_2 regularisation. The different curves compare featured maps at different epochs of training. The theoretical predictions based on the Gaussian covariate model (full lines) are in very good agreement with the actual performance (points).

3.3 GAN-generated data and learned teachers

To approach more realistic data sets, we now consider the case in which the input data $\mathbf{x} \in \mathcal{X}$ is given by a generative neural network $\mathbf{x} = \mathcal{G}(\mathbf{z})$, where \mathbf{z} is a Gaussian i.i.d. latent vector. Therefore, the covariates $[\mathbf{u}, \mathbf{v}]$ are the result of the following Markov chain:

$$\mathbf{z} \xrightarrow{\mathcal{G}} \mathbf{x} \in \mathcal{X} \xrightarrow{\varphi_t} \mathbf{u} \in \mathbb{R}^p, \quad \mathbf{z} \xrightarrow{\mathcal{G}} \mathbf{x} \in \mathcal{X} \xrightarrow{\varphi_s} \mathbf{v} \in \mathbb{R}^d. \quad (3.4)$$

With a model for the covariates, the missing ingredient is the teacher weights $\boldsymbol{\theta}_0 \in \mathbb{R}^p$, which determine the label assignment: $y = f_0(\mathbf{u}^\top \boldsymbol{\theta}_0)$. In the experiments that follows, we fit the teacher weights *from the original data set in which the generative model \mathcal{G} was trained*. Different choices for the fitting yield different teacher weights, and the quality of label assignment can be accessed by the performance of the fit on the test set. The set $(\varphi_t, \varphi_s, \mathcal{G}, \boldsymbol{\theta}_0)$ fully defines the data generative process. For predicting the learning curves from the iterative eqs. (2.17) we also need to estimate the *population* covariances (Ψ, Φ, Ω) . This is done from the generative process in eq. (3.4) with a Monte Carlo sampling algorithm. This pipeline is explained in full detail in Appendix E. An open source implementation of the algorithms used in the experiments is available in [41].

Fig. 3 shows an example of the learning curves resulting from the pipeline discussed above in a logistic regression task on data generated by a GAN trained on CIFAR10 images. More concretely, we used a pre-trained five-layers Deep convolutional GAN (dcGAN) from [55], which maps 100 dimensional i.i.d. Gaussian noise into $k = 32 \times 32 \times 3$ realistic looking CIFAR10-like images: $\mathcal{G} : \mathbf{z} \in \mathbb{R}^{100} \mapsto \mathbf{x} \in \mathbb{R}^{32 \times 32 \times 3}$. To generate labels, we trained a simple fully-connected four-layer neural network on the *real* CIFAR10 data set, on a odd ($y = +1$) vs. even ($y = -1$) task, achieving $\sim 75\%$ classification accuracy on the test set. The teacher weights $\boldsymbol{\theta}_0 \in \mathbb{R}^p$ were taken from the last layer of the network, and the teacher feature map φ_t from the three previous layers. For the student model, we trained a completely independent fully connected 3-layer neural network on the dcGAN-generated CIFAR10-like images and took snapshots of the feature maps φ_s^i induced by the 2-first layers during the first $i \in \{0, 5, 50, 200\}$ epochs of training. Finally, once $(\mathcal{G}, \varphi_t, \varphi_s^i, \boldsymbol{\theta}_0)$ have been fixed, we

estimated the covariances (Ψ, Φ, Ω) with a Monte Carlo algorithm. Full detail of the architectures used and of the training procedure can be found in Appendix E.1.

Fig. 3 depicts the resulting learning curves obtained by training the last layer of the student. Interestingly, the performance of the feature map at epoch 0 (random initialisation) beats the performance of the learned features during early phases of training in this experiment. Another interesting behaviour is given by the separability threshold of the learned features, i.e. the number of samples for which the training loss becomes larger than 0 in logistic regression. At epoch 50 the learned features are separable at lower sample complexity $\alpha = n/d$ than at epoch 200 - even though in the later the training and generalisation performances are better.

3.4 Learning from real data sets

Applying teacher/students to a real data set — Given that the learning curves of realistic-looking inputs can be captured by the Gaussian covariate model, it is fair to ask whether the same might true for *real data sets*. To test this idea, we first need to cast the real data set into the teacher-student formalism, and then compute the covariances matrices Ω, Ψ, Φ and teacher vector θ_0 required by model (1.1).

Let $\{\mathbf{x}^\mu, y^\mu\}_{\mu=1}^{n_{\text{tot}}}$ denote a real data set, e.g. MNIST or Fashion-MNIST for concreteness, where $n_{\text{tot}} = 7 \times 10^4$, $\mathbf{x}^\mu \in \mathbb{R}^D$ with $D = 784$. Without loss of generality, we can assume the data is centred. To generate the teacher, let $\mathbf{u}^\mu = \varphi_t(\mathbf{x}^\mu) \in \mathbb{R}^p$ be a feature map such data is invertible in feature space, i.e. that $y^\mu = \theta_0^\top \mathbf{u}^\mu$ for some teacher weights $\theta_0 \in \mathbb{R}^p$, which should be computed from the samples. Similarly, let $\mathbf{v}^\mu = \varphi_s(\mathbf{x}^\mu) \in \mathbb{R}^d$ be a feature map we are interested in studying (e.g. random features or a pre-trained neural network). Then, we can estimate the population covariances (Ψ, Φ, Ω) empirically from the *entire* data set as:

$$\Psi = \sum_{\mu=1}^{n_{\text{tot}}} \frac{\mathbf{u}^\mu \mathbf{u}^{\mu\top}}{n_{\text{tot}}}, \quad \Phi = \sum_{\mu=1}^{n_{\text{tot}}} \frac{\mathbf{u}^\mu \mathbf{v}^{\mu\top}}{n_{\text{tot}}}, \quad \Omega = \sum_{\mu=1}^{n_{\text{tot}}} \frac{\mathbf{v}^\mu \mathbf{v}^{\mu\top}}{n_{\text{tot}}}. \quad (3.5)$$

At this point, we have all we need to run the self-consistent equations (2.17). The issue with this approach is that there is not a unique teacher map φ_t and teacher vector θ_0 that fits the true labels: this requirement seems ill-defined, as it depends on this rather arbitrary choice. However, we can show that *all interpolating linear teachers are equivalent*:

Theorem 3. (*Universality of linear teachers*) *For any teacher feature map φ_t , and for any θ_0 that interpolates the real data set so that $y^\mu = \theta_0^\top \mathbf{u}^\mu \forall \mu$, the asymptotic predictions of model (1.1) are equivalent.*

Proof. It follows from the fact that the teacher weights and covariances only appear in eq. (2.17) through $\rho = \frac{1}{p} \theta_0^\top \Psi \theta_0$ and the projection $\Phi \theta_0$. Using the estimation (3.5) and the assumption that it exists $y^\mu = \theta_0^\top \mathbf{u}^\mu$, one can write these quantities directly from the labels y^μ :

$$\rho = \frac{1}{n_{\text{tot}}} \sum_{\mu=1}^{n_{\text{tot}}} (y^\mu)^2, \quad \Phi \theta_0 = \frac{1}{n_{\text{tot}}} \sum_{\mu=1}^{n_{\text{tot}}} y^\mu \mathbf{v}^\mu. \quad (3.6)$$

For any linear interpolating teacher, the theory is thus independent of the choice of the teacher feature map. \square

Ridge Regression with linear teachers — Using these ideas, we now test the prediction of model (1.1) on real data sets, and show that it is surprisingly effective in predicting the learning curves, at least for the ridge regression task. We have trained a 3-layer fully connected neural network with ReLU activations on the full Fashion-MNIST data set to distinguish clothing used above vs. below the waist [57]. The student feature map $\varphi_s : \mathbb{R}^{784} \rightarrow \mathbb{R}^d$ is obtained by removing the last layer, see Appendix E.1 for a detailed description. In Fig. 4 we show the test and training errors of the ridge estimator on a sub-sample of $n < n_{\text{tot}}$ on the Fashion-MNIST images. We observe remarkable agreement between the learning curve obtained from simulations and the theoretical prediction by the matching Gaussian covariate model.

Fig. 4 shows a similar experiment on the MNIST data set, but for different out-of-the-box feature maps, such as random features and the scattering transform [58]. Note the characteristic double-descent behaviour [9], and

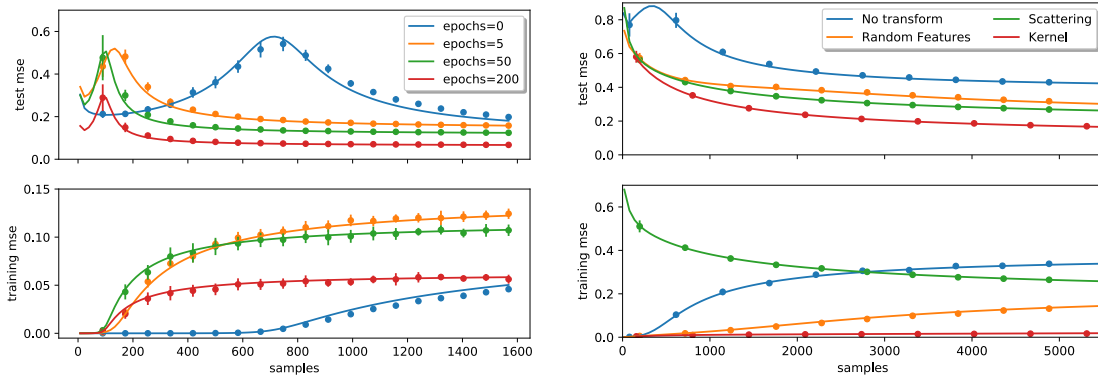


Figure 4: (Left) Test and training mean-squared errors eqs. (2.18) as a function of the number of samples n for ridge regression on the Fashion-MNIST data set, with vanishing regularisation $\lambda = 10^{-5}$. In this plot, the student feature map φ_s is a 3-layer fully-connected neural network with $d = 2352$ hidden neurons trained on the full data set with the square loss. Different curves correspond to the feature map obtained at different stages of training. (Right) Test and training mean-squared errors eq. (2.18) as a function of the number of samples n for ridge regression on a MNIST even vs odd digit task, with regularisation $\lambda = 10^{-2}$. The different curves compare the performance of no feature map (blue), random feature map with $\sigma = \text{erf}$ and Gaussian projection (orange), the scattering transform with parameters $J = 3, L = 8$ [18] (green) and of the limiting kernel of the RF map above [56] (red). Note we chose the number of random features $d = 1953$ to match the number of features from the scattering transform. Full details on the simulations are described in Appendix E.1

the accurate prediction of the peak where the interpolation transition occurs. We note in Appendix E.1 that for both Fig. 4 and 4, for a number of samples n closer to n_{tot} we start to see deviations between the real learning curve and the theory: this is to be expected since in that case the teacher can be learned with zero test error, while the test error on the real data set saturates.

Why is the Gaussian model so effective for describing learning with data that are *not* Gaussian? The point is that ridge regression is sensitive only to second order statistics, and not to the full distribution of the data. It is a classical property (see Appendix F) that the training and generalisation errors are only a function of the spectrum of the *empirical* and *population* covariances, and of their products. Random matrix theory teach us that such quantities are very robust, and their asymptotic behaviour is universal for a broad class of distributions of $[\mathbf{u}, \mathbf{v}]$ [59, 60, 61, 47]. The asymptotic behavior of kernel matrices has indeed been the subject of intense scrutiny [45, 62, 46, 11, 63, 48]. We thus expect the validity of model (1.1) for ridge regression, with a linear teacher, to go way beyond the Gaussian assumption.

Beyond ridge regression — The same strategy fails beyond ridge regression and mean-squared test error. This suggests a limit in the application of model (1.1) to real (non-Gaussian) data to the universal linear teacher. To illustrate this, consider the setting of Figs. 4 and 4, and compare the model predictions for the binary classification error instead of the ℓ_2 one. There is a clear mismatch between the simulated performance and prediction given by the theory (see Appendix E.1) due to the fact that the classification error does not depends only on the two first moments.

We present an additional experiment in Fig. 5. We compare the learning curves of logistic regression on a classification task on the *real* CIFAR10 images with the real labels versus the one on dcGAN-generated CIFAR10-like images and teacher generated labels from Sec. 3.3. While the Gaussian theory captures well the behaviour of the later, it fails on the former. A histogram of the distribution of the product $\mathbf{u}^\top \hat{\mathbf{w}}$ for a fixed number of samples illustrates well the deviation from the prediction of the theory with the real case, in particular on the tails of the distribution. The difference between GAN generated data (that fits the Gaussian theory) and real data is striking. Given that for classification problems there exists a number of choices of "sign" teachers and feature maps that give the exact same labels as in the data set, an interesting question is: *is there a teacher*




Figure 5: (Left) Test classification error, eq. (2.19) (top) and (unregularised) training loss, eq. (2.2) (bottom) for logistic regression as a function of the number of samples n for an animal vs not-animal binary classification task with ℓ_2 penalty $\lambda = 10^{-2}$, comparing real CIFAR10 grey-scale images (blue) with dcGAN-generated CIFAR10-like gray-scale images (red). The real-data learning curve was estimated, just as in Figs. 4 and 4 by estimating the population covariances on the full data set. (Right) Histograms of the variable $\frac{1}{\sqrt{d}} \mathbf{v}^\top \hat{\mathbf{w}}$ for a fixed number of samples $n = 2d = 2048$ and the respective p.d.f. given by the theoretical prediction (solid line) from eq. (2.4).

that allows to reproduce the learning curves more accurately. This question is left for future work.

Acknowledgements

We acknowledge funding from the ERC under the European Union’s Horizon 2020 Research and Innovation Programme Grant Agreement 714608-SMiLe, from “Chaire de recherche sur les modèles et sciences des données”, Fondation CFM pour la Recherche-ENS, and from the French National Research Agency grants ANR-17-CE23-0023-01 PAIL and ANR-19-P3IA-0001 PRAIRIE.

A Main result from the replica method

In this appendix we derive the formula for the performance of the Gaussian covariate model from a heuristic replica analysis. The computation closely follows the recent developments in [42, 14]. We refer to [21, 3, 22] for an introduction to this remarkable heuristic (but seemingly never failing) approach.

The data: First, let's recall the definition of our model. Consider synthetic labelled data $(\mathbf{v}, y) \in \mathbb{R}^d \times \mathbb{R}$ drawn independently from a joint distribution with density:

$$p_{\theta_0}(\mathbf{v}, y) = \int_{\mathbb{R}^p} d\mathbf{u} P_0(y|\mathbf{u}^\top \boldsymbol{\theta}_0) \mathcal{N}(\mathbf{u}, \mathbf{v}; \mathbf{0}, \Sigma) \quad (\text{A.1})$$

where P_0 is a given likelihood on \mathbb{R} , $\boldsymbol{\theta}_0 \in \mathbb{R}^p$ is a fixed vector of parameters and Σ is a correlation matrix given by:

$$\Sigma = \begin{bmatrix} \Psi & \Phi \\ \Phi^\top & \Omega \end{bmatrix} \in \mathbb{R}^{(p+d) \times (p+d)} \quad (\text{A.2})$$

for symmetric positive semi-definite matrices Ψ and Ω and $\Phi \in \mathbb{R}^{p \times d}$. In its simplest form, which we will mostly be using in the applications, we take the likelihood $P_0(y|x) = \delta(y - f_0(x))$ to be a deterministic function with $f_0 : \mathbb{R} \rightarrow \mathbb{R}$ a non-linearity, e.g. $f_0(x) = \text{sign}(x)$ to generate binary labels.

The task: In our analysis, we are interested in the training and generalisation performance of a linear classifier $\hat{y} = f_{\mathbf{w}}(\mathbf{v}) = \hat{f}(\mathbf{w}^\top \mathbf{v})$ trained on n independent samples $\mathcal{D} = \{(\mathbf{v}^\mu, y^\mu)\}_{\mu=1}^n$ from p_{θ_0} by minimising the regularised empirical risk:

$$\hat{\mathbf{w}} = \arg \min_{\mathbf{w} \in \mathbb{R}^d} \left[\sum_{\mu=1}^n g(y^\mu, \mathbf{w}^\top \mathbf{v}^\mu) + \frac{\lambda}{2} \|\mathbf{w}\|_2^2 \right], \quad (\text{A.3})$$

where $\lambda > 0$ is the regularisation strength. We define the *sample complexity* $\alpha = n/d$ and the *aspect ratio* $\gamma = p/d$.

Gibbs minimisation: As it was proven in Theorem 1 of the main manuscript, the asymptotic performance of the estimator in eq. (A.3) is fully characterised by the following scalar parameters:

$$\rho = \frac{1}{p} \boldsymbol{\theta}_0^\top \Psi \boldsymbol{\theta}_0, \quad m^* = \frac{1}{\sqrt{pd}} \boldsymbol{\theta}_0^\top \Phi \hat{\mathbf{w}}, \quad q^* = \frac{1}{d} \hat{\mathbf{w}}^\top \Omega \hat{\mathbf{w}} \quad (\text{A.4})$$

The replica method is precisely a heuristic tool allowing us to circumvent the high-dimensional estimation problem defined in eq. (A.3) and giving us direct access to (m^*, q^*) .

The starting point is to define the following Gibbs measure over weights $\mathbf{w} \in \mathbb{R}^d$:

$$\mu_\beta(d\mathbf{w}) = \frac{1}{\mathcal{Z}_\beta} e^{-\beta \left[\sum_{\mu=1}^n g(y^\mu, \mathbf{w}^\top \mathbf{v}^\mu) + \frac{\lambda}{2} \sum_{i=1}^d w_i^2 \right]} d\mathbf{w} = \frac{1}{\mathcal{Z}_\beta} \underbrace{\prod_{\mu=1}^n e^{-\beta g(y^\mu, \mathbf{w}^\top \mathbf{v}^\mu)}}_{P_g} \underbrace{\prod_{i=1}^d e^{-\frac{\beta \lambda}{2} w_i^2}}_{P_w} d\mathbf{w}_i \quad (\text{A.5})$$

Note that P_g and P_w can be interpreted as a (unnormalised) likelihood and prior distribution respectively. In the limit $\beta \rightarrow \infty$, the measure μ_β concentrates around solutions of the minimisation in eq. (A.3). The aim in the replica method is to compute the free energy density, defined as:

$$\beta f_\beta = - \lim_{d \rightarrow \infty} \frac{1}{d} \mathbb{E}_{\mathcal{D}} \log \mathcal{Z}_\beta. \quad (\text{A.6})$$

A.1 Replica computation of the free energy

The average in eq. (A.6) is not straightforward due to the logarithm term. The replica method consists of computing it using the following trick to get rid of the logarithm:

$$\log \mathcal{Z}_\beta = \lim_{r \rightarrow 0^+} \frac{1}{r} \partial_r \mathcal{Z}_\beta^r \quad (\text{A.7})$$

Averaging

Applying the trick above, the computation of the free energy density boils down to the evaluation of the averaged replicated partition function:

$$\begin{aligned} \mathbb{E}_{\mathcal{D}} \mathcal{Z}_\beta^r &= \prod_{\mu=1}^n \mathbb{E}_{(\mathbf{v}^\mu, y^\mu)} \prod_{a=1}^r \int_{\mathbb{R}^d} P_w(d\mathbf{w}^a) P_g \left(y^\mu \middle| \frac{\mathbf{v}^\mu \cdot \mathbf{w}^a}{\sqrt{d}} \right) \\ &= \prod_{\mu=1}^n \int_{\mathbb{R}} dy^\mu \int_{\mathbb{R}^p} P_{\theta_0}(d\theta_0) \int_{\mathbb{R}^{d \times r}} \left(\prod_{a=1}^r P_w(d\mathbf{w}^a) \right) \underbrace{\mathbb{E}_{\mathbf{u}^\mu, \mathbf{v}^\mu} \left[P_0 \left(y^\mu \middle| \frac{\mathbf{u}^\mu \cdot \boldsymbol{\theta}_0}{\sqrt{p}} \right) \prod_{a=1}^r P_g \left(y^\mu \middle| \frac{\mathbf{v}^\mu \cdot \mathbf{w}^a}{\sqrt{d}} \right) \right]}_{(*)} \end{aligned} \quad (\text{A.8})$$

Note that in the above we included an average over the parameters $\boldsymbol{\theta}_0 \in \mathbb{R}^p$, which the deterministic case in which this is fixed given by a point mass $P_{\boldsymbol{\theta}_0} = \delta_{\boldsymbol{\theta}_0}$. Focusing on the average term in brackets:

$$\begin{aligned} (*) &= \mathbb{E}_{(\mathbf{u}, \mathbf{v})} \left[P_0 \left(y \middle| \frac{\mathbf{u} \cdot \boldsymbol{\theta}_0}{\sqrt{p}} \right) \prod_{a=1}^r P_g \left(y \middle| \frac{\mathbf{v} \cdot \mathbf{w}^a}{\sqrt{d}} \right) \right] \\ &= \int_{\mathbb{R}} d\nu_\mu P_0(y|\nu_\mu) \int_{\mathbb{R}^r} \left(\prod_{a=1}^r d\lambda_\mu^a P_g(y^\mu|\lambda_\mu^a) \right) \underbrace{\mathbb{E}_{(\mathbf{u}^\mu, \mathbf{v}^\mu)} \left[\delta \left(\nu_\mu - \frac{\mathbf{u}^\mu \cdot \boldsymbol{\theta}_0}{\sqrt{p}} \right) \prod_{a=1}^r \delta \left(\lambda_\mu^a - \frac{\mathbf{v}^\mu \cdot \mathbf{w}^a}{\sqrt{d}} \right) \right]}_{P(\nu, \lambda)} \end{aligned}$$

Note that the term in brackets defines the joint density over (ν_μ, λ_μ^a) . It is easy to check that these are Gaussian random variables with zero mean and covariance matrix given by:

$$\Sigma^{ab} = \begin{pmatrix} \rho & m^a \\ m^a & Q^{ab} \end{pmatrix}. \quad (\text{A.9})$$

where the so-called overlap parameters (ρ, m^a, Q^{ab}) are related to the weights $\boldsymbol{\theta}_0, \mathbf{w}$:

$$\rho \equiv \mathbb{E} [\nu_\mu^2] = \frac{1}{p} \boldsymbol{\theta}_0^\top \Psi \boldsymbol{\theta}_0, \quad m^a \equiv \mathbb{E} [\lambda_\mu^a \nu_\mu] = \frac{1}{\sqrt{pd}} \boldsymbol{\theta}_0^\top \Phi \mathbf{w}^a, \quad Q^{ab} \equiv \mathbb{E} [\lambda_\mu^a \lambda_\mu^b] = \frac{1}{d} \mathbf{w}^{a\top} \Omega \mathbf{w}^b$$

We can therefore write the averaged replicated partition function as:

$$\mathbb{E}_{\mathcal{D}} \mathcal{Z}_\beta^r = \prod_{\mu=1}^n \int dy^\mu \int_{\mathbb{R}^p} P_{\theta_0}(d\boldsymbol{\theta}_0) \int_{\mathbb{R}^{d \times r}} \left(\prod_{a=1}^r P_w(d\mathbf{w}^a) \right) \mathcal{N}(\nu_\mu, \lambda_\mu^a; \mathbf{0}, \Sigma^{ab}) \quad (\text{A.10})$$

Rewriting as a saddle-point problem

The next step is to free the overlap parameters by introducing delta functions:

$$\begin{aligned}
1 &\propto \int_{\mathbb{R}} d\rho \delta(p\rho - \boldsymbol{\theta}_0^\top \Psi \boldsymbol{\theta}_0) \int_{\mathbb{R}^r} \prod_{a=1}^r dm^a \delta(\sqrt{p}dm^a - \boldsymbol{\theta}_0^\top \Phi \mathbf{w}^a) \times \int_{\mathbb{R}^{r \times r}} \prod_{1 \leq a \leq b \leq r} dQ^{ab} \delta(dQ^{ab} - \mathbf{w}^{a \top} \Omega \mathbf{w}^b) \\
&= \int_{\mathbb{R}} \frac{d\rho d\hat{\rho}}{2\pi} e^{i\hat{\rho}(p\rho - \boldsymbol{\theta}_0^\top \Psi \boldsymbol{\theta}_0)} \int_{\mathbb{R}^r} \prod_{a=1}^r \frac{dm^a d\hat{m}^a}{2\pi} e^{i \sum_{a=1}^r \hat{m}^a (\sqrt{p}dm^a - \boldsymbol{\theta}_0^\top \Phi \mathbf{w}^a)} \times \\
&\quad \times \int_{\mathbb{R}^{r \times r}} \prod_{1 \leq a \leq b \leq r} \frac{dQ^{ab} d\hat{Q}^{ab}}{2\pi} e^{i \sum_{1 \leq a \leq b \leq r} \hat{Q}^{ab} (dQ^{ab} - \mathbf{w}^{a \top} \Omega \mathbf{w}^b)} \tag{A.11}
\end{aligned}$$

Inserting this in eq. (A.10) allow us to rewrite:

$$\mathbb{E}_{\mathcal{D}} \mathcal{Z}_{\beta}^r = \int_{\mathbb{R}} \frac{d\rho d\hat{\rho}}{2\pi} \int_{\mathbb{R}^r} \prod_{a=1}^r \frac{dm^a d\hat{m}^a}{2\pi} \int_{\mathbb{R}^{r \times r}} \prod_{1 \leq a \leq b \leq r} \frac{dQ^{ab} d\hat{Q}^{ab}}{2\pi} e^{d\Phi^{(r)}} \tag{A.12}$$

where we have absorbed a $-i$ factor in the integrals (this won't matter since we will look to the saddle-point) and defined the potential:

$$\Phi^{(r)} = -\gamma\rho\hat{\rho} - \sqrt{\gamma} \sum_{a=1}^r m^a \hat{m}^a - \sum_{1 \leq a \leq b \leq r} Q^{ab} \hat{Q}^{ab} + \alpha \Psi_y^{(r)}(\rho, m^a, Q^{ab}) + \Psi_w^{(r)}(\hat{\rho}, \hat{m}^a, \hat{Q}^{ab}) \tag{A.13}$$

where we recall that $\alpha = n/d$, $\gamma = p/d$ and:

$$\Psi_w^{(r)} = \frac{1}{d} \log \int_{\mathbb{R}^p} P_{\boldsymbol{\theta}_0} (d\boldsymbol{\theta}_0) \int_{\mathbb{R}^{d \times r}} \prod_{a=1}^r P_w (d\mathbf{w}^a) e^{\hat{\rho} \boldsymbol{\theta}_0^\top \Psi \boldsymbol{\theta}_0 + \sum_{a=1}^r \hat{m}^a \boldsymbol{\theta}_0^\top \Phi \mathbf{w}^a + \sum_{1 \leq a \leq b \leq r} \hat{Q}^{ab} \mathbf{w}^{a \top} \Omega \mathbf{w}^b} \tag{A.14}$$

$$\Psi_y^{(r)} = \log \int_{\mathbb{R}} dy \int_{\mathbb{R}} d\nu P_0(y|\nu) \int \prod_{a=1}^r d\lambda^a P_g(y|\lambda^a) \mathcal{N}(\nu, \lambda^a; \mathbf{0}, \Sigma^{ab}) \tag{A.15}$$

In the high-dimensional limit where $d \rightarrow \infty$ while $\alpha = n/d$ and $\gamma = p/d$ stay finite, the integral in eq. (A.12) concentrate around the values of the overlaps that extremise $\Phi^{(r)}$, and therefore we can write:

$$\beta f_{\beta} = - \lim_{r \rightarrow 0^+} \frac{1}{r} \mathbf{extr} \Phi^{(r)} \left(\hat{\rho}, \hat{m}^a, \hat{Q}^{ab}; \rho, m^a, Q^{ab} \right) \tag{A.16}$$

Replica symmetric ansatz

In order to proceed with the $r \rightarrow 0^+$ limit, we restrict the extremisation above to the following replica symmetric ansatz:

$$\begin{aligned}
m^a &= m, & \hat{m}^a &= \hat{m}, & \text{for } a &= 1, \dots, r \\
Q^{aa} &= r, & \hat{Q}^{aa} &= -\frac{1}{2} \hat{r}, & \text{for } a &= 1, \dots, r \\
Q^{ab} &= q, & \hat{Q}^{ab} &= \hat{q}, & \text{for } 1 \leq a < b \leq r &
\end{aligned} \tag{A.17}$$

Inserting this ansatz in eq. (A.13) allow us to explicitly take the $r \rightarrow 0^+$ limit for each term. The first three terms are straightforward to obtain. The limit of $\Psi_y^{(r)}$ is cumbersome, but it common to many replica computations for the generalised linear likelihood P_g . We refer the curious reader to Appendix C of [42] or to Appendix IV of [14] for details, and write the final result here:

$$\Psi_y \equiv \lim_{r \rightarrow 0^+} \frac{1}{r} \Psi_w^{(r)} = \mathbb{E}_{\xi \sim \mathcal{N}(0,1)} \left[\int_{\mathbb{R}} dy \mathcal{Z}_0 \left(y, \frac{m}{\sqrt{q}} \xi, \rho - \frac{m^2}{q} \right) \log \mathcal{Z}_g(y, \sqrt{q} \xi, V) \right] \tag{A.18}$$

where we have defined $V = r - q$ and:

$$\mathcal{Z}_{g/0}(y, \omega, V) = \mathbb{E}_{x \sim \mathcal{N}(\omega, V)} [P_{g/0}(y|x)]. \quad (\text{A.19})$$

Note that as in [42], the consistency condition of the zeroth order term in the free energy fix $\rho = \mathbb{E}_{\theta_0} \left[\frac{1}{p} \theta_0^\top \Psi \theta_0 \right]$ and $\hat{\rho} = 0$. On the other hand, the limit of the prior term here is exactly as the one discussed in Appendix C of [39], and is given by:

$$\Psi_w \equiv \lim_{r \rightarrow 0^+} \frac{1}{r} \Psi_w^{(r)} = \frac{1}{d} \mathbb{E}_{\xi, \theta_0} \log \int_{\mathbb{R}^d} P_w(\mathbf{d}\mathbf{w}) e^{-\frac{\hat{V}}{2} \mathbf{w}^\top \Omega \mathbf{w} + \mathbf{w}^\top (\hat{m} \Phi^\top \theta_0 + \hat{q} \Omega^{1/2} \xi)}. \quad (\text{A.20})$$

Summary

The replica symmetric free energy density is simply given by:

$$f_\beta = \mathbf{extr}_{q, m, \hat{q}, \hat{m}} \left\{ -\frac{1}{2} r \hat{r} - \frac{1}{2} q \hat{q} + \sqrt{\gamma} m \hat{m} - \alpha \Psi_y(r, m, q) - \Psi_w(\hat{r}, \hat{m}, \hat{q}) \right\} \quad (\text{A.21})$$

where

$$\begin{aligned} \Psi_w &= \lim_{d \rightarrow \infty} \frac{1}{d} \mathbb{E}_{\xi, \theta_0} \log \int_{\mathbb{R}^d} P_w(\mathbf{d}\mathbf{w}) e^{-\frac{\hat{V}}{2} \mathbf{w}^\top \Omega \mathbf{w} + \mathbf{w}^\top (\hat{m} \Phi^\top \theta_0 + \hat{q} \Omega^{1/2} \xi)} \\ \Psi_y &= \mathbb{E}_{\xi \sim \mathcal{N}(0,1)} \left[\int_{\mathbb{R}} dy \mathcal{Z}_0 \left(y, \frac{m}{\sqrt{q}} \xi, \rho - \frac{m^2}{q} \right) \log \mathcal{Z}_g(y, \sqrt{q} \xi, V) \right] \end{aligned} \quad (\text{A.22})$$

A.2 Ridge regression and fixed weights

For an ℓ_2 -regularisation term, we have:

$$P_w(\mathbf{d}\mathbf{w}) = \frac{1}{(2\pi)^{d/2}} e^{-\frac{\beta\lambda}{2} \|\mathbf{w}\|_2^2} \mathbf{d}\mathbf{w} \quad (\text{A.23})$$

where we have included a convenient constant, and therefore:

$$\begin{aligned} \int_{\mathbb{R}^d} P_w(\mathbf{d}\mathbf{w}) e^{-\frac{\hat{V}}{2} \mathbf{w}^\top \Omega \mathbf{w} + \mathbf{w}^\top (\hat{m} \Phi^\top \theta_0 + \sqrt{\hat{q}} \Omega^{1/2} \xi)} &= \int_{\mathbb{R}^d} \frac{\mathbf{d}\mathbf{w}}{(2\pi)^{p/2}} e^{-\frac{1}{2} \mathbf{w}^\top (\beta\lambda \mathbf{I}_d + \hat{V} \Omega) \mathbf{w} + \mathbf{w}^\top (\hat{m} \Phi^\top \theta_0 + \sqrt{\hat{q}} \Omega^{1/2} \xi)} \\ &= \frac{\exp \left(\frac{1}{2} (\hat{m} \Phi^\top \theta_0 + \sqrt{\hat{q}} \Omega^{1/2} \xi)^\top (\beta\lambda \mathbf{I}_d + \hat{V} \Omega)^{-1} (\hat{m} \Phi^\top \theta_0 + \sqrt{\hat{q}} \Omega^{1/2} \xi) \right)}{\sqrt{\det (\beta\lambda \mathbf{I}_d + \hat{V} \Omega)}} \end{aligned} \quad (\text{A.24})$$

taking the log and using $\log \det = \text{tr} \log$, up to the limit:

$$\begin{aligned} \Psi_w &= \frac{1}{2d} \mathbb{E}_{\xi, \theta_0} \left[(\hat{m} \Phi^\top \theta_0 + \sqrt{\hat{q}} \Omega^{1/2} \xi)^\top (\beta\lambda \mathbf{I}_d + \hat{V} \Omega)^{-1} (\hat{m} \Phi^\top \theta_0 + \sqrt{\hat{q}} \Omega^{1/2} \xi) \right] \\ &\quad - \frac{1}{2d} \text{tr} \log (\beta\lambda \mathbf{I}_d + \hat{V} \Omega) \end{aligned} \quad (\text{A.25})$$

Defining the shorthand $\mathbf{A} = (\beta\lambda \mathbf{I}_d + \hat{V} \Omega)^{-1}$, we can now take the averages over ξ explicitly:

$$\mathbb{E}_\xi \left[(\hat{m} \Phi^\top \theta_0 + \sqrt{\hat{q}} \Omega^{1/2} \xi)^\top \mathbf{A} (\hat{m} \Phi^\top \theta_0 + \sqrt{\hat{q}} \Omega^{1/2} \xi) \right] = \hat{m}^2 \theta_0^\top \Phi \mathbf{A} \Phi^\top \theta_0 + \hat{q} \text{tr} \Omega^{1/2} \mathbf{A} \Omega^{1/2} \quad (\text{A.26})$$

Putting together, up to the limit:

$$\Psi_w = -\frac{1}{2d} \text{tr} \log (\beta\lambda \mathbf{I}_d + \hat{V} \Omega) + \frac{1}{2d} \text{tr} \left(\hat{m}^2 \Phi^\top \theta_0 \theta_0^\top \Phi + \hat{q} \Omega \right) (\beta\lambda \mathbf{I}_d + \hat{V} \Omega)^{-1} \quad (\text{A.27})$$

A.3 Taking the $\beta \rightarrow \infty$ limit

Finally, in order to take the $\beta \rightarrow \infty$ limit explicitly, we note that under the rescaling

$$\begin{aligned} V &\rightarrow \beta^{-1}V & q &\rightarrow q & m &\rightarrow m \\ \hat{V} &\rightarrow \beta\hat{V} & \hat{q} &\rightarrow \beta^2\hat{q} & \hat{m} &\rightarrow \beta\hat{m}. \end{aligned} \quad (\text{A.28})$$

The potential Ψ_w has a trivial limit:

$$\lim_{\beta \rightarrow \infty} \frac{1}{\beta} \Psi_w = -\frac{1}{2d} \text{tr} \log \left(\lambda \mathbf{I}_d + \hat{V} \Omega \right) + \frac{1}{2d} \text{tr} \left(\hat{m}^2 \Phi^\top \boldsymbol{\theta}_0 \boldsymbol{\theta}_0^\top \Phi + \hat{q} \Omega \right) \left(\lambda \mathbf{I}_d + \hat{V} \Omega \right)^{-1} \quad (\text{A.29})$$

while Ψ_y requires more attention. Since \mathcal{Z}_0 only depends on (q, m) , it is invariant under the rescaling. On the other hand, we have that:

$$\mathcal{Z}_g(y, \sqrt{q}\xi, V) = \sqrt{\beta} \int \frac{dx}{\sqrt{2\pi V}} e^{-\beta \left[\frac{(x - \sqrt{q}\xi)^2}{2V} + g(y, x) \right]} \underset{\beta \rightarrow \infty}{=} e^{-\beta \mathcal{M}_{Vg(y, \cdot)}(\sqrt{q}\xi)} \quad (\text{A.30})$$

where \mathcal{M} is the Moreau envelope associated to the loss g :

$$\mathcal{M}_{\tau g(y, \cdot)}(x) = \inf_{z \in \mathbb{R}} \left[\frac{(z - x)^2}{2\tau} + g(y, z) \right] \quad (\text{A.31})$$

and therefore:

$$\lim_{\beta \rightarrow \infty} \frac{1}{\beta} \Psi_y = -\mathbb{E}_{\xi \sim \mathcal{N}(0,1)} \left[\int dy \mathcal{Z}_0 \left(y, \frac{m}{\sqrt{q}} \xi, \rho - \frac{m^2}{q} \right) \mathcal{M}_{Vg(y, \cdot)}(\sqrt{q}\xi) \right] \quad (\text{A.32})$$

The zero temperature therefore is simply given by:

$$\begin{aligned} \lim_{\beta \rightarrow \infty} f_\beta = \mathbf{extr}_{V, q, m, \hat{V}, \hat{q}, \hat{m}} & \left\{ -\frac{1}{2} \hat{V} V - \frac{1}{2} (q \hat{V} - \hat{q} V) + \sqrt{\gamma} m \hat{m} + \alpha \mathbb{E}_{\xi \sim \mathcal{N}(0,1)} \left[\int dy \mathcal{Z}_0 \mathcal{M}_{Vg(y, \cdot)} \right] \right. \\ & \left. + \frac{1}{2d} \text{tr} \log \left(\lambda \mathbf{I}_d + \hat{V} \Omega \right) - \frac{1}{2d} \text{tr} \left(\hat{m}^2 \Phi^\top \boldsymbol{\theta}_0 \boldsymbol{\theta}_0^\top \Phi + \hat{q} \Omega \right) \left(\lambda \mathbf{I}_d + \hat{V} \Omega \right)^{-1} \right\} \end{aligned} \quad (\text{A.33})$$

A.4 Saddle-point equations

To solve the extremisation problem defined by eq. (A.33), we search for vanishing gradient points of the potential. This lead to a set of self-consistent *saddle-point* equations:

$$\begin{cases} \hat{V} = -\alpha \mathbb{E}_\xi \left[\int_{\mathbb{R}} dy \mathcal{Z}_0 \partial_\omega f_g \right] \\ \hat{q} = \alpha \mathbb{E}_\xi \left[\int_{\mathbb{R}} dy \mathcal{Z}_0 f_g^2 \right] \\ \hat{m} = \frac{\alpha}{\sqrt{\gamma}} \mathbb{E}_\xi \left[\int_{\mathbb{R}} dy \partial_\omega \mathcal{Z}_0 f_g \right] \end{cases} \begin{cases} V = \frac{1}{d} \text{tr} \left(\lambda \mathbf{I}_d + \hat{V} \Omega \right)^{-1} \Omega \\ q = \frac{1}{d} \text{tr} \left[(\hat{q} \Omega + \hat{m}^2 \Phi^\top \boldsymbol{\theta}_0 \boldsymbol{\theta}_0^\top \Phi) \Omega \left(\lambda \mathbf{I}_d + \hat{V} \Omega \right)^{-2} \right] \\ m = \frac{1}{\sqrt{\gamma}} \frac{\hat{m}}{p} \text{tr} \Phi^\top \boldsymbol{\theta}_0 \boldsymbol{\theta}_0^\top \Phi \left(\lambda \mathbf{I}_d + \hat{V} \Omega \right)^{-1} \end{cases} \quad (\text{A.34})$$

where $f_g(y, \omega, V) = -\partial_\omega \mathcal{M}_{Vg(y, \cdot)}(\omega)$, which can also be obtained from the proximal operator

$$\text{prox}_{Vg(y, \cdot)}(\omega) = \arg \min_{z \in \mathbb{R}} \left[\frac{(z - \omega)^2}{2V} + g(y, z) \right] \quad (\text{A.35})$$

using the envelope theorem $\mathcal{M}'_{\tau f}(x) = \tau^{-1} (x - \text{prox}_{\tau f}(x))$. A python implementation of the saddle-point equations for the losses discussed below is available in [41].

A.5 Examples

We now discuss a couple of examples in which the equations above simplify.

Ridge regression: Consider a ridge regression task with $f_0(x) = \hat{f}(x) = x$, loss $g(y, x) = \frac{1}{2}(y - x)^2$ and choose $\hat{g}(y, x) = \frac{1}{2}(y - x)^2$. In this case, our model is closely related to the mismatched models in [12] and [25]. In the first, labels are generated in a higher-dimensional space which contains the features as a subspace, and can be mapped to our model in the case $p > d$ by defining the projection of the teacher weights in the student space $\Phi^\top \theta_0 \in \mathbb{R}^d$ and its orthogonal complement $(\Phi^\top \theta_0)^\perp \in \mathbb{R}^{p-d}$. In the second, the teacher acts on an orthogonal subset of the features, and can be mapped with a similar construction to our model in the case $p < d$. These two cases were studied for specific linear tasks, such as ridge and random features regression, with the covariances modelling structure in the data. Conceptually, our model differs slightly in the sense that any additional fixed feature layer, e.g. random projections or a pre-trained feature map, is also contained in the covariances.

For the linear task, the asymptotic training and generalisation errors read:

$$\mathcal{E}_{\text{train.}}^* = \frac{\rho + q^* - 2m^*}{(1 + V^*)^2}, \quad \mathcal{E}_{\text{gen.}}^* = \rho + q^* - 2m^* \quad (\text{A.36})$$

where $\rho = \frac{1}{p} \theta_0^\top \Psi \theta_0$ and (V^*, q^*, m^*) are the fixed point of the following set of self-consistent equations:

$$\begin{cases} \hat{V} = \frac{\alpha}{1+V} \\ \hat{q} = \alpha \frac{\rho + q - 2m}{(1+V)^2} \\ \hat{m} = \frac{1}{\sqrt{\gamma}} \frac{\alpha}{1+V} \end{cases}, \quad \begin{cases} V = \frac{1}{d} \text{tr} \left(\lambda \mathbf{I}_d + \hat{V} \Omega \right)^{-1} \Omega \\ q = \frac{1}{d} \text{tr} \left[\left(\hat{q} \Omega + \hat{m}^2 \Phi^\top \theta_0 \theta_0^\top \Phi \right) \Omega \left(\lambda \mathbf{I}_d + \hat{V} \Omega \right)^{-2} \right] \\ m = \frac{1}{\sqrt{\gamma}} \frac{\hat{m}}{d} \text{tr} \Phi^\top \theta_0 \theta_0^\top \Phi \left(\lambda \mathbf{I}_d + \hat{V} \Omega \right)^{-1} \end{cases}. \quad (\text{A.37})$$

Note that quite interestingly we have the following relationship between the training and generalisation error:

$$\mathcal{E}_{\text{train.}}^* = \frac{\mathcal{E}_{\text{gen.}}^*}{(1 + V^*)^2}. \quad (\text{A.38})$$

This give us an interesting interpretation of V^* as parametrising the variance gap between the generalisation and training error¹. In particular, note that V^* only depends on the spectrum of the population covariance, since it is the solution of:

$$V = \int \nu_\Omega(d\omega) \frac{\omega}{\lambda + \frac{\alpha\omega}{1+V}} \quad (\text{A.39})$$

where ν_Ω is the spectral density of Ω .

Binary classification

For a binary classification task, we tak $f_0(x) = \hat{f}(x) = \text{sign}(x) \in \{-1, 1\}$. Our equations generalise the ones derived [38] in the specific case of $d = p$ and $\Psi = \mathbf{I}_d$, $\theta_0 \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$. For binary classification, the asymptotic classification error $\mathcal{E}_{\text{gen.}}(\hat{w}) = \mathbb{P}(y \neq \text{sign}(\hat{w}^\top \mathbf{u}))$ can be explicitly written is terms of the overlaps as:

$$\mathcal{E}_{\text{gen.}}^* = \frac{1}{\pi} \cos^{-1} \left(\frac{m^*}{\sqrt{\rho q^*}} \right). \quad (\text{A.40})$$

where again (q^*, m^*) are solutions of the self-consistent saddle-point equations. The teacher measure is given by:

$$\mathcal{Z}_0(y, \omega, V) = \frac{\delta_{y,1} + \delta_{y,-1}}{2} \left(1 + \text{erf} \left(\frac{y\omega}{\sqrt{2V}} \right) \right) \quad (\text{A.41})$$

The explicit form of the equation depends on the choice of the loss function, three of which are of particular interest:

¹We thank Stéphane d'Ascoli for bringing this relation to our attention.

Square-loss: As in the ridge case, for $g(y, x) = \frac{1}{2}(y - x)^2$ the saddle-point equations simplify considerably:

$$\begin{cases} \hat{V} = \frac{\alpha}{1+V} \\ \hat{q} = \alpha \frac{1+q-2m\sqrt{\frac{2}{\pi\rho}}}{(1+V)^2} \\ \hat{m} = \sqrt{\frac{2}{\pi\rho}} \frac{\alpha}{1+V} \end{cases}, \quad \begin{cases} V = \frac{1}{d} \text{tr} \left(\lambda \mathbf{I}_d + \hat{V} \Omega \right)^{-1} \Omega \\ q = \frac{1}{d} \text{tr} \left[\left(\hat{q} \Omega + \hat{m}^2 \Phi^\top \boldsymbol{\theta}_0 \boldsymbol{\theta}_0^\top \Phi \right) \Omega \left(\lambda \mathbf{I}_d + \hat{V} \Omega \right)^{-2} \right] \\ m = \frac{1}{\sqrt{\gamma}} \frac{\hat{m} \gamma}{d} \text{tr} \Phi^\top \boldsymbol{\theta}_0 \boldsymbol{\theta}_0^\top \Phi \left(\lambda \mathbf{I}_d + \hat{V} \Omega \right)^{-1} \end{cases}. \quad (\text{A.42})$$

Similarly, the asymptotic training error also admits a simple expression:

$$\mathcal{E}_{\text{train}}^* = \frac{1}{4} \frac{1 + q^* - 2m^* \sqrt{\frac{2}{\pi\rho}}}{(1 + V^*)^2} \quad (\text{A.43})$$

Logistic regression: Different from the previous cases, for logistic loss $g(y, x) = \log(1 + e^{-yx})$ the equations for $(\hat{V}, \hat{q}, \hat{m})$ cannot be integrated explicitly, since the proximal operator doesn't admit a closed form solution. Instead, f_g can be found by solving the following self-consistent equation:

$$f_g = \frac{y}{1 + e^{-y(Vf_g + \omega)}}. \quad (\text{A.44})$$

Note that this also allow us to express $\partial_\omega f_g = (2 \cosh(\frac{1}{2}y(Vf_g + \omega)))^{-2}$.

Soft-margin regression: Another useful case in which the proximal operator has a closed form solution is for the hinge loss $g(y, x) = \max(0, 1 - yx)$. In this case:

$$f_g(y, \omega, V) = \begin{cases} y & \text{if } \omega y < 1 - V \\ \frac{y - \omega}{V} & \text{if } 1 - V < \omega y < 1 \\ 0 & \text{otherwise} \end{cases}, \quad \partial_\omega f_g(y, \omega, V) = \begin{cases} -\frac{1}{V} & \text{if } 1 - V < \omega y < 1 \\ 0 & \text{otherwise} \end{cases} \quad (\text{A.45})$$

Again, the equations cannot be integrated explicitly. Note that in the limit $\lambda \rightarrow 0$, both the logistic and soft-margin solutions converge to the max-margin estimator

A.6 Relation to previous models

Random features: The feature map for random features learning can be written as:

$$\Phi_{\text{F}} : \mathbf{u} \in \mathbb{R}^p \mapsto \mathbf{v} = \sigma \left(\frac{1}{\sqrt{k}} \mathbf{F} \mathbf{u} \right) \in \mathbb{R}^d \quad (\text{A.46})$$

where $\mathbf{u} \in \mathbb{R}^p$ is the original data, $\mathbf{F} \in \mathbb{R}^{d \times k}$ is a chosen random projection matrix and $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ is a chosen non-linearity acting component-wise in \mathbb{R}^d , see [16]. Random features learning has attracted a lot of interest recently, and has been studied in [11, 42, 38, 31] in the case of Gaussian data $\mathbf{u} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_p)$. Our model encompasses all of these works, and in the case of Gaussian data the covariance (Ψ, Ω, Φ) can be explicitly related to the projection matrix \mathbf{F} :

$$\Psi = \mathbf{I}_p, \quad \Phi = \kappa_1 \mathbf{F}, \quad \Omega = \kappa_0^2 \mathbf{1}_d \mathbf{1}_d^\top + \kappa_1^2 \frac{\mathbf{F} \mathbf{F}^\top}{d} + \kappa_\star^2 \mathbf{I}_d \quad (\text{A.47})$$

where $\mathbf{1}_d \in \mathbb{R}^d$ is the all-ones vector and the constants $(\kappa_0, \kappa_1, \kappa_\star)$ are related to σ as:

$$\kappa_0 = \mathbb{E}_{z \sim \mathcal{N}(0,1)} [\sigma(z)], \quad \kappa_1 = \mathbb{E}_{z \sim \mathcal{N}(0,1)} [z \sigma(z)], \quad \kappa_\star = \sqrt{\mathbb{E}_{z \sim \mathcal{N}(0,1)} [\sigma(z)^2] - \kappa_0^2 - \kappa_1^2} \quad (\text{A.48})$$

These relations hold asymptotically, and rely on the *Gaussian equivalence theorem* (GET), see [39] for a proof.

Generative models: In [39], a similar Gaussian covariate model was used to study the performance of random feature regression on data generated from pre-trained generative models:

$$\mathbf{v} = \mathcal{G}(\mathbf{u}) \in \mathbb{R}^d, \quad \mathbf{u} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_p). \quad (\text{A.49})$$

where $\mathcal{G} : \mathbb{R}^p \rightarrow \mathbb{R}^d$ is a generative network mapping the latent space \mathbb{R}^p to the input space \mathbb{R}^d (e.g. a pre-trained GAN). Labels were generated directly in the latent space \mathbb{R}^p using a generalised linear model on random weights: $y = f_0(\boldsymbol{\theta}_0^\top \mathbf{u})$ with $\mathbf{u} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_p)$. A *Gaussian Equivalence Principle* (GEP) stating that the asymptotic generalisation and training performances of this model are fully captured by second order statistics was conjectured and shown to hold numerically for different choices of generative models \mathcal{G} . Indeed, this model is a particular case of ours when $\Psi = \mathbf{I}_p$ and $\boldsymbol{\theta}_0 \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_p)$. Assuming that the GEP holds, our model therefore can be seen as a generalisation of [39] to structured teachers. For instance, in Section 3.3 of the main we show several cases in which the teacher $\mathbf{u} = \tilde{\mathcal{G}}(\mathbf{c})$ for a latent vector $\mathbf{c} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_k)$ and a pre-trained map $\tilde{\mathcal{G}}$ that can include a generative model and a fixed feature map (e.g. random features, scattering transform, pre-learned neural network, etc.). Also, it is important to stress that our model also account for the case in which the teacher weights $\boldsymbol{\theta}_0 \in \mathbb{R}^p$ are fixed, and therefore can be also learned.

Kernel methods: Let \mathcal{H} be a Kernel Reproducing Hilbert space (RKHS) associated to a given kernel K and $\mathcal{D} = \{\mathbf{x}^\mu, y^\mu\}_{\mu=1}^n$ be a labelled data set with $\mathbf{x} \sim p_x$ independently, and set $\mathcal{X} = \text{supp}(p_x)$. In Kernel regression, the aim is to solve:

$$\min_{f \in \mathcal{H}} \left[\frac{1}{2} \sum_{\mu=1}^n (y^\mu - f(\mathbf{x}^\mu))^2 + \frac{\lambda}{2} \|f\|_{\mathcal{H}}^2 \right] \quad (\text{A.50})$$

where $\|\cdot\|_{\mathcal{H}}$ is the norm induced by the scalar product in \mathcal{H} . An alternative representation of this problem is given by the feature decomposition of the kernel given by Mercer's theorem:

$$K(\mathbf{x}, \mathbf{x}') = \sum_{i=1}^{\infty} \omega_i e_i(\mathbf{x}') e_i(\mathbf{x}) \quad (\text{A.51})$$

where ω_i and $e_i(\mathbf{x})$ are the eigenvalues and eigenvectors associated with the kernel:

$$\int_{\mathbb{R}^k} p_x(d\mathbf{x}') K(\mathbf{x}, \mathbf{x}') e_i(\mathbf{x}') = \omega_i e_i(\mathbf{x}) \quad (\text{A.52})$$

Note that $\{e_i(\mathbf{x})\}_i^\infty$ form an orthonormal basis of the space of square-integrable functions $L^2(\mathcal{X})$ (with respect to the standard scalar product of L^2). It is also convenient to define the feature map $\varphi_i(\mathbf{x}) = \sqrt{\omega_i} e_i(\mathbf{x})$, which is an orthonormal basis of $\mathcal{H} \subset L^2(\mathcal{X})$ (with respect to the scalar product induced by K). Therefore, if we assume that the labels $y^\mu = f_0(\mathbf{x}^\mu)$ are generated from a ground truth target function (not necessarily part of \mathcal{H}), we can expand both f and f_0 the feature basis:

$$f(\mathbf{x}) = \sum_{i=1}^{\infty} w_i \varphi_i(\mathbf{x}), \quad f_0(\mathbf{x}) = \sum_{i=1}^{\infty} \theta_{0i} \varphi_i(\mathbf{x}) \quad (\text{A.53})$$

Note that $f \in \mathcal{H}$ implies that for this sum to make sense w_i needs to decay fast enough with respect to $\sqrt{\omega_i}$, but in general we can have $f_0 \notin \mathcal{H}$ meaning that θ_{0i} decays slower than $\sqrt{\omega_i}$ but still fast enough such that $f_0 \in L^2(\mathcal{X})$. If the number of features is finite ($\omega_i = 0$ for $i \geq d$) or if we introduce a cut-off $d \gg n, |\mathcal{X}|$, the representation in the feature basis in eq. (A.53) allow us to rewrite Kernel regression problem in eq. (A.50) simply as ridge regression in feature space:

$$\min_{\mathbf{w} \in \mathbb{R}^d} \left[\frac{1}{2} \sum_{\mu=1}^n \left(\boldsymbol{\theta}_0^\top \boldsymbol{\varphi}(\mathbf{x}^\mu) - \mathbf{w}^\top \boldsymbol{\varphi}(\mathbf{x}^\mu) \right)^2 + \frac{\lambda}{2} \|\mathbf{w}\|_2^2 \right]. \quad (\text{A.54})$$

Letting $\mathbf{v} = \varphi(\mathbf{x}) \in \mathbb{R}^d$, this formulation is equivalent to our model with $p = d$ and covariance matrices given by:

$$\Psi = \Phi = \Omega = \text{diag}(\omega_i). \quad (\text{A.55})$$

Indeed, inserting this expression equation (A.37):

$$\begin{cases} \hat{V} = \hat{m} = \frac{\alpha}{1+\hat{V}} \\ \hat{q} = \alpha \frac{\rho+q-2m}{(1+\hat{V})^2} \end{cases}, \quad \begin{cases} V = \frac{1}{d} \sum_{i=1}^d \frac{\omega_i}{\lambda+\hat{V}\omega_i} \\ q = \frac{1}{d} \sum_{i=1}^d \frac{\hat{q}\omega_i^2 + \theta_{0i}^2 \omega_i^3 \hat{m}^2}{(\lambda+\hat{V}\omega_i)^2} \\ m = \frac{\hat{m}}{d} \sum_{i=1}^d \frac{\omega_i^2 \theta_{0i}^2}{\lambda+\hat{V}\omega_i} \end{cases}. \quad (\text{A.56})$$

and making a change of variables $\hat{q} \leftarrow \hat{q} \frac{d^2}{n}$, $\hat{m} \leftarrow \hat{m} \frac{d}{n}$, $\hat{V} \leftarrow \hat{q} \frac{d}{n}$, $\rho \leftarrow d\rho$, $m \leftarrow dm$, $q \leftarrow dq$, $\lambda \leftarrow d\lambda$ we recover exactly the self-consistent equations of [37] for the performance of kernel ridge regression directly from our equations. In Appendix B we show how to recover the results of [37] and [53] for the power law scaling of the generalisation error in kernel ridge regression. Moreover, our model allow to generalise this discussion to more involved kernel tasks such as kernel logistic regression and support vector machines.

B Power law scaling for kernel regression

In this appendix we derive the scaling results from [53, 37] for the decay of the generalisation error in kernel ridge regression directly from our framework.

As discussed in Appendix A.6, kernel ridge regression can be written as ridge regression in feature space, see eqs. (A.56). Here, we are interested in understanding how the generalisation error depends on the decay of the teacher weights and of the student covariance Ω , in the case considered in [37, 53] where the student dimension is infinity $d = \infty$.

We remind that in the kernel ridge regression case, the saddle point equations read:

$$\begin{cases} \hat{V} = \hat{m} = \frac{\alpha}{1+\hat{V}} \\ \hat{q} = \alpha \frac{\rho+q-2\hat{m}}{(1+\hat{V})^2} \end{cases}, \quad \begin{cases} V = \frac{1}{d} \sum_{i=1}^d \frac{\omega_i}{\lambda + \hat{V}\omega_i} \\ q = \frac{1}{d} \sum_{i=1}^d \frac{\hat{q}\omega_i^2 + \tilde{\theta}_i\omega_i\hat{m}^2}{(\lambda + \hat{V}\omega_i)^2} \\ m = \frac{\hat{m}}{d} \sum_{i=1}^d \frac{\omega_i\tilde{\theta}_i}{\lambda + \hat{V}\omega_i} \end{cases}, \quad (\text{B.1})$$

where we have defined $\tilde{\theta}_i = \theta_{0i}^2\omega_i$ for convenience, with ω_i being the i^{th} eigenvalue of the student covariance Ω . In subsection B.1, we derive a pair of coupled self-consistent equations determining the generalization error ϵ_g , which we simplify in the $d \rightarrow \infty$ limit in subsection B.2 and extract therefrom the asymptotic behaviour for ϵ_g in subsection B.3 as $n \gg 1$.

B.1 Self-consistent equations for the generalization error

Define $z = \alpha \frac{\lambda}{\hat{V}}$. Note that we can find a closed equation for z by combining the equations for (\hat{V}, V) :

$$z \equiv \alpha \frac{\lambda}{\hat{V}} = \alpha \frac{\lambda}{\alpha} (1 + V) = \lambda + \frac{1}{d} \sum_{i=1}^d \frac{\lambda\omega_i}{\lambda + \hat{V}\omega_i}. \quad (\text{B.2})$$

Therefore:

$$z = \lambda + \frac{z}{n} \sum_{i=1}^d \frac{\omega_i}{\frac{z}{\alpha} + \omega_i}. \quad (\text{B.3})$$

The test error is given by $\epsilon_g = \rho + q - 2m$. Using the saddle-point equations (B.1), we can write:

$$\begin{aligned} \epsilon_g &\stackrel{(a)}{=} \frac{1}{d} \sum_{i=1}^d \left[\tilde{\theta}_i + \frac{\hat{q}\omega_i^2 + \tilde{\theta}_i\omega_i\hat{m}^2}{(\lambda + \hat{V}\omega_i)^2} - \frac{2\hat{m}\tilde{\theta}_i\omega_i}{\lambda + \hat{V}\omega_i} \right] \\ &\stackrel{(b)}{=} \frac{1}{d} \sum_{i=1}^d \frac{(\lambda + \hat{V}\omega_i)^2 + \alpha^{-1}\omega_i^2\hat{V}^2\epsilon_g + \hat{V}^2\tilde{\theta}_i\omega_i - 2\tilde{\theta}_i\hat{V}\omega_i(\lambda + \hat{V}\omega_i)}{(\lambda + \hat{V}\omega_i)^2} \end{aligned} \quad (\text{B.4})$$

$$= \frac{1}{d} \sum_{i=1}^d \frac{\alpha^{-1}\omega_i^2\hat{V}^2\epsilon_g + \lambda^2\tilde{\theta}_i}{(\lambda + \hat{V}\omega_i)^2}, \quad (\text{B.5})$$

where in (a) we simply used the equations for q, m and the definition of ρ and in (b) we used $\hat{m} = \hat{V}$ and $\hat{q} = \alpha^{-1}\hat{V}^2\epsilon_g$. Note that the right-hand side now only depends on the overlap \hat{V} . Solving for ϵ_g and writing everything as a function of $z = \alpha\lambda/\hat{V}$:

$$\epsilon_g = \frac{\frac{z^2 d}{n^2} \sum_{i=1}^d \frac{\tilde{\theta}_i}{(z \frac{d}{n} + \omega_i)^2}}{1 - \frac{1}{n} \sum_{i=1}^d \frac{\omega_i^2}{(z \frac{d}{n} + \omega_i)^2}}. \quad (\text{B.6})$$

Equation (B.6) together with equation (B.3) make for a pair of coupled equations whose solution allows to determine ϵ_g . As written, solving for (B.3) and (B.6) is not tractable. We show in subsection B.2 how the expressions (B.3) and (B.6) simplify as the $d \rightarrow \infty$ limit is taken. In subsection B.3, we further show how in the $n \gg 1$ regime the asymptotic behaviour for the generalization error (B.6) can be extracted.

B.2 Limit of infinite student dimension

Now assume the following power law decay for the covariance Ω and the teacher θ_0 :

$$\omega_i = i^{-b}d, \quad \tilde{\theta}_i = i^{-a}d, \quad (\text{B.7})$$

i.e. the eigenvalues of Ω possess a power-law decay with exponent b , while the components of the teacher θ_{0i} decay with exponent $\frac{a-b}{2}$. Note that ansatz (B.7) slightly differs from the one assumed in [37, 53] by a factor d , resulting from differing normalization conventions in the loss considered. Since labels $y^\mu = \frac{1}{\sqrt{d}}\theta_0 \cdot x^\mu$ are defined with a normalization \sqrt{d} in the present work, the factor d is necessary in (B.7) to keep labels of $\mathcal{O}(1)$. Then equation (B.6) can be simplified to

$$\epsilon_g = \frac{\frac{z^2}{n^2} \sum_{i=1}^d \frac{i^{-a}}{(z\frac{1}{n} + i^{-b})^2}}{1 - \frac{1}{n} \sum_{i=1}^d \frac{i^{-2b}}{(z\frac{1}{n} + i^{-b})^2}}, \quad (\text{B.8})$$

which has a meaningful limit as $d \rightarrow \infty$ (with n, λ kept fixed):

$$\epsilon_g = \frac{\sum_{i=1}^{\infty} \frac{i^{-a}}{(1+nz^{-1}i^{-b})^2}}{1 - \frac{n}{z^2} \sum_{i=1}^{\infty} \frac{i^{-2b}}{(1+nz^{-1}i^{-b})^2}}. \quad (\text{B.9})$$

In the same limit, the equation defining z (B.3) is amenable to being rewritten:

$$z = \lambda + \frac{z}{n} \sum_{i=1}^{\infty} \frac{1}{1 + \frac{z}{n} i^b}. \quad (\text{B.10})$$

B.3 Infinite sample limit and the scaling of the generalisation error

Consider now the limit $n \gg 1$ with λ kept fixed. We can use a Riemann approximation to rewrite (B.10) in integral form, for $n \gg 1$:

$$z \approx \lambda + \left(\frac{z}{n}\right)^{1-\frac{1}{b}} \int_{\left(\frac{z}{n}\right)^{1/b}}^{\infty} \frac{dx}{1+x^b}. \quad (\text{B.11})$$

Note that the scalings of z with respect to n differ according to the regularisation λ , depending on which of the two terms on the right hand side of equation (B.11) dominates. If the first λ term dominates, then (B.11) simplifies to $z \approx \lambda$. For this to be self-consistent, we must have $(z/n)^{1-\frac{1}{b}} \approx (\lambda/n)^{1-\frac{1}{b}} \ll \lambda$, i.e. $n \gg \lambda^{-\frac{1}{b-1}}$. In the converse case where the second term in (B.11) dominates, $z \sim n^{1-b}$. For this to consistently hold, one needs $(z/n)^{1-\frac{1}{b}} \approx n^{1-b} \gg \lambda$, i.e. $n \ll \lambda^{-\frac{1}{b-1}}$. Therefore there exist in total two regimes:

- in the $n \ll \lambda^{-\frac{1}{b-1}}$ regime, $z \sim n^{1-b}$,
- in the $n \gg \lambda^{-\frac{1}{b-1}}$ regime, $z \approx \lambda$.

In the first regime, the numerator of ϵ_g reads possesses different scalings according to exponents a, b . if $2b - a > -1$,

$$\sum_{i=1}^{\infty} \frac{i^{-a}}{(1 + nz^{-1}i^{-b})^2} \sim n^{1-a} \sum_{i=1}^{\infty} \frac{\left(\frac{i}{n}\right)^{-a}}{\left(1 + \left(\frac{i}{n}\right)^{-b}\right)^2} \frac{1}{n} \sim n^{1-a} \int_0^{\infty} \frac{x^{2b-a}}{(1+x^b)^2} = \mathcal{O}(n^{1-a}). \quad (\text{B.12})$$

If $2b - a < -1$, it is no longer possible to write the Riemann sum as an integral, and

$$\sum_{i=1}^{\infty} \frac{i^{-a}}{(1 + nz^{-1}i^{-b})^2} = \sum_{i=1}^n \frac{i^{-a}}{(1 + n^b i^{-b})^2} + n^{1-a} \sum_{i=n}^{\infty} \frac{\left(\frac{i}{n}\right)^{-a}}{\left(1 + \left(\frac{i}{n}\right)^{-b}\right)^2} \frac{1}{n} = \mathcal{O}(n^{-2b}). \quad (\text{B.13})$$

It is possible to similarly decompose the sum in the denominator to find

$$\frac{n}{z^2} \sum_{i=1}^{\infty} \frac{i^{-2b}}{(1 + nz^{-1}i^{-b})^2} = \mathcal{O}(1). \quad (\text{B.14})$$

From this, it follows that:

$$\epsilon_g = \mathcal{O}\left(n^{-\min(a-1, 2b)}\right) \quad (\text{B.15})$$

By the same token, in the second $n \gg \lambda^{-\frac{1}{b-1}}$ regime, provided $2b - a > -1$, one can write the numerator as a Riemann sum:

$$\sum_{i=1}^{\infty} \frac{i^{-a}}{(1 + nz^{-1}i^{-b})^2} \sim \left(\frac{\lambda}{n}\right)^{\frac{a-1}{b}} \sum_{i=1}^{\infty} \frac{\left(i \left(\frac{\lambda}{n}\right)^{\frac{1}{b}}\right)^{2b-a}}{\left(\left(i \left(\frac{\lambda}{n}\right)^{\frac{1}{b}}\right)^b + 1\right)^2} \left(\frac{\lambda}{n}\right)^{\frac{1}{b}} \sim \left(\frac{\lambda}{n}\right)^{\frac{a-1}{b}} \int_0^{\infty} \frac{x^{2b-a}}{(1+x^b)^2} = \mathcal{O}(n^{\frac{1-a}{b}}). \quad (\text{B.16})$$

In the $2b - a < -1$ case,

$$\sum_{i=1}^{\infty} \frac{i^{-a}}{(1 + nz^{-1}i^{-b})^2} = \sum_{i=1}^{\frac{n}{\lambda}} \frac{i^{-a}}{\left(1 + \frac{n}{\lambda} i^{-b}\right)^2} + \left(\frac{\lambda}{n}\right)^{\frac{a-1}{b}} \sum_{i=\frac{n}{\lambda}}^{\infty} \frac{\left(i \left(\frac{\lambda}{n}\right)^{\frac{1}{b}}\right)^{2b-a}}{\left(\left(i \left(\frac{\lambda}{n}\right)^{\frac{1}{b}}\right)^b + 1\right)^2} \left(\frac{\lambda}{n}\right)^{\frac{1}{b}}. \quad (\text{B.17})$$

Upper and lower bounds can be straightforwardly found for the first sum and the following equivalence established

$$\sum_{i=1}^{\frac{n}{\lambda}} \frac{i^{-a}}{\left(1 + \frac{n}{\lambda} i^{-b}\right)^2} \sim n^{-2} \sum_{i=1}^{\frac{n}{\lambda}} i^{2b-a} = \mathcal{O}(n^{-2}), \quad (\text{B.18})$$

while the second sum is a Riemann sum of order $\mathcal{O}(n^{\frac{1-a}{b}}) = o(n^{-2})$. Finally, the scaling of the sum in the denominator is found along similar lines to be

$$\frac{n}{z^2} \sum_{i=1}^{\infty} \frac{i^{-2b}}{(1 + nz^{-1}i^{-b})^2} = \mathcal{O}(n^{\frac{1-b}{b}}) = o(1). \quad (\text{B.19})$$

Putting everything together,

$$\epsilon_g = \mathcal{O}\left(n^{-\frac{\min(a-1, 2b)}{b}}\right). \quad (\text{B.20})$$

The scaling exponents (B.15) and (B.20) are in agreement with those derived through other means in [37]. (B.15) recovers the exponents derived using Fourier analysis in [53], with $a = \frac{\alpha_T}{d}$, $b = \frac{\alpha_S}{d}$ in their notations.

C Rigorous proof of the Main result

In this Appendix, we discuss the complete proof of Theorem 1 and 2 using the Gaussian comparison Theorem 4.

C.1 Necessary assumptions

We start with a list of the necessary assumptions for our result to hold. We also briefly discuss how they are relevant in a supervised machine learning context.

- (A1)** The vector θ_0 is pulled from any given distribution $p_{\theta_0} \in \mathbb{R}^p$ with sub-Gaussian one dimensional marginals (this includes deterministic vectors with bounded norm), and is independent of the matrices U and V . Additionally, the signal is non-vanishing and has finite energy, i.e. the following holds almost surely:

$$\lim_{p \rightarrow \infty} 0 < \mathbb{E} \left[\frac{\theta_0^\top \theta_0}{p} \right] < +\infty \quad (\text{C.1})$$

- (A2)** The covariance matrices verify:

$$(\Psi, \Omega) \in \mathbb{S}_p^{++} \times \mathbb{S}_d^{++}, \quad \Omega - \Phi^\top \Psi^{-1} \Phi \succeq 0 \quad (\text{C.2})$$

Additionally, the spectral distributions of the matrices Φ , Ψ and Ω converge to distributions with almost surely bounded eigenvalues.

- (A3)** The functions r and g are proper, lower semi-continuous, convex functions. Additionally, we assume that the cost function $r + g$ is coercive, i.e.:

$$\lim_{\|\mathbf{w}\|_2 \rightarrow +\infty} (r + g)(\mathbf{w}) = +\infty \quad (\text{C.3})$$

and that the following scaling condition holds : for all $n, d \in \mathbb{N}$, $\mathbf{z} \in \mathbb{R}^n$ and any constant $c > 0$, there exist finite, positive constants C_1, C_2, C_3 , such that, for any standard normal random vectors $\mathbf{h} \in \mathbb{R}^d$ and $\mathbf{g} \in \mathbb{R}^n$:

$$\|\mathbf{z}\|_2 \leq c\sqrt{n} \implies \sup_{\mathbf{x} \in \partial g(\mathbf{z})} \|\mathbf{x}\|_2 \leq C_1\sqrt{n}, \quad \frac{1}{d}\mathbb{E}[r(\mathbf{h})] < +\infty, \quad \frac{1}{n}\mathbb{E}[g(\mathbf{g})] < +\infty \quad (\text{C.4})$$

- (A4)** The random elements of the function f_0 are independent of the matrices U and V . Additionally the following limit exists and is finite

$$\lim_{n \rightarrow \infty} \mathbb{E} \left[\frac{1}{n} f_0(U\theta_0)^\top f_0(U\theta_0) \right] < +\infty$$

- (A5)** When we send the dimensions n, p, d to infinity, they grow with finite ratios $\alpha = n/d, \gamma = p/d$.
- (A6) Additional assumptions for linear finite sample size rates :** the functions r, g, ϕ_1, ϕ_2 are pseudo-Lipschitz of finite order. The eigenvalues of the covariance matrices are bounded with probability one.
- (A7) Additional assumptions for exponential finite sample size rates:** the loss function g is separable and pseudo-Lipschitz of order 2, the regularisation is either a ridge or a Lipschitz function, the functions ϕ_1, ϕ_2 are respectively separable, pseudo-Lipschitz of order 2, and a square or Lipschitz function.

The first assumption (A1) simply ensures that the teacher distribution is non-trivial and has a non-vanishing structure. The positive definiteness in (A2) means the covariance matrices of the blocks U and V are well-specified. Note that the cross-correlation matrix Φ can have singular values equal to zero. The assumption about the limiting spectral distribution is a classical one in random matrix theory and ensures that the eigenvalues are almost surely non-diverging. The scaling assumptions from (A3) are natural as they imply that non-diverging

inputs result in non-diverging outputs in the functions f and g , as well as the sub-differentials. Similar scaling assumptions are encountered in proofs such as [26]. They also allow to show Gaussian concentration of Moreau envelopes, as we will see in Lemma 5. The *coercivity* assumption is verified in most common machine learning setups : any convex loss with ridge regularisation, or any convex loss that is bounded below with a coercive regularisation, see Corollary 11.15 from [64]. Most machine learning setups of interest are thus included : ridge regularized classification or regression, LASSO, log-barrier regularisation, etc... An important remark is the case of unregularised minimisation problems which can be non-coercive. A typical example would be the plain logistic regression with no penalty, where a phase transition in the existence of a minimiser can be observed. This is recovered by taking the limit of vanishing regularisation parameter in the expressions obtained here. Assumption (A4) is a classical assumption of teacher-student setups, where any correlation between the teacher and the student is modeled by the covariance matrices and not by the label generating function f_0 . The summability condition ensures generalization error is well-defined for squared performance measures. Finally, (A5) is the typical *high-dimensional* limit used in statistical physics of learning, random matrix theory and a large recent body of work in high-dimensional statistical learning.

C.2 Theoretical toolbox

Here we remind a few known results that are used throughout the proof. We also provide proofs of useful, straightforward consequences of these results that do not appear explicitly in the literature for completeness.

C.2.1 A Gaussian comparison theorem

We start with the Convex Gaussian Min-max Theorem, as presented in [26], which is a tight version of an inequality initially derived in [20].

Theorem 4. (CGMT) *Let $\mathbf{G} \in \mathbb{R}^{m \times n}$ be an i.i.d. standard normal matrix and $\mathbf{g} \in \mathbb{R}^m$, $\mathbf{h} \in \mathbb{R}^n$ two i.i.d. standard normal vectors independent of one another. Let $\mathcal{S}_{\mathbf{w}}, \mathcal{S}_{\mathbf{u}}$ be two compact sets such that $\mathcal{S}_{\mathbf{w}} \subset \mathbb{R}^n$ and $\mathcal{S}_{\mathbf{u}} \subset \mathbb{R}^m$. Consider the two following optimization problems for any continuous ψ on $\mathcal{S}_{\mathbf{w}} \times \mathcal{S}_{\mathbf{u}}$:*

$$\mathbf{C}(\mathbf{G}) := \min_{\mathbf{w} \in \mathcal{S}_{\mathbf{w}}} \max_{\mathbf{u} \in \mathcal{S}_{\mathbf{u}}} \mathbf{u}^\top \mathbf{G} \mathbf{w} + \psi(\mathbf{w}, \mathbf{u}), \quad (\text{C.5})$$

$$\mathcal{C}(\mathbf{g}, \mathbf{h}) := \min_{\mathbf{w} \in \mathcal{S}_{\mathbf{w}}} \max_{\mathbf{u} \in \mathcal{S}_{\mathbf{u}}} \|\mathbf{w}\|_2 \mathbf{g}^\top \mathbf{u} + \|\mathbf{u}\|_2 \mathbf{h}^\top \mathbf{w} + \psi(\mathbf{w}, \mathbf{u}) \quad (\text{C.6})$$

then the following holds:

1. For all $c \in \mathbb{R}$:

$$\mathbb{P}(\mathbf{C}(\mathbf{G}) < c) \leq 2\mathbb{P}(\mathcal{C}(\mathbf{g}, \mathbf{h}) \leq c)$$

2. Further assume that $\mathcal{S}_{\mathbf{w}}, \mathcal{S}_{\mathbf{u}}$ are convex sets and ψ is convex-concave on $\mathcal{S}_{\mathbf{w}} \times \mathcal{S}_{\mathbf{u}}$. Then, for all $c \in \mathbb{R}$,

$$\mathbb{P}(\mathbf{C}(\mathbf{G}) > c) \leq 2\mathbb{P}(\mathcal{C}(\mathbf{g}, \mathbf{h}) \geq c)$$

In particular, for all $\mu \in \mathbb{R}, t > 0, \mathbb{P}(|\mathbf{C}(\mathbf{G}) - \mu| > t) \leq 2\mathbb{P}(|\mathcal{C}(\mathbf{g}, \mathbf{h}) - \mu| \geq t)$.

Following [26], we will say that any reformulation of a target problem matching the form of (C.5) is an acceptable primary optimization problem (PO), and the corresponding form (C.6) is an acceptable auxiliary problem (AO). The main idea of this approach is to study the asymptotic properties of the (PO) by studying the simpler (AO).

C.2.2 Proximal operators and Moreau envelopes : differentials and useful functions

Here we remind the definition and some important properties of Moreau envelopes and proximal operators, key elements of convex analysis. Other properties will be used throughout the proof but at less crucial stages, thus we don't remind them explicitly. Our main reference for these properties will be [64].

Consider a closed, proper function f such that $\text{dom}(f) \subset \mathbb{R}^n$. Its Moreau envelope and proximal operator are respectively defined by :

$$\mathcal{M}_{\tau f}(\mathbf{x}) = \min_{\mathbf{z} \in \text{dom}(f)} \left\{ f(\mathbf{z}) + \frac{1}{2\tau} \|\mathbf{x} - \mathbf{z}\|_2^2 \right\}, \quad \text{prox}_{\tau f}(\mathbf{x}) = \arg \min_{\mathbf{z} \in \text{dom}(f)} \left\{ f(\mathbf{z}) + \frac{1}{2\tau} \|\mathbf{x} - \mathbf{z}\|_2^2 \right\} \quad (\text{C.7})$$

As reminded in [26], the Moreau envelope is jointly convex in (τ, \mathbf{x}) and differentiable almost everywhere, with gradients:

$$\nabla_{\mathbf{x}} \mathcal{M}_{\tau f}(\mathbf{x}) = \frac{1}{\tau} (\mathbf{x} - \text{prox}_{\tau f}(\mathbf{x})) \quad (\text{C.8})$$

$$\frac{\partial}{\partial \tau} \mathcal{M}_{\tau f}(\mathbf{x}) = -\frac{1}{2\tau^2} \|\mathbf{x} - \text{prox}_{\tau f}(\mathbf{x})\|_2^2 \quad (\text{C.9})$$

We remind that $\text{prox}_{\tau f}(\mathbf{x})$ is the unique point which solves the strongly convex optimization problem defining the Moreau envelope, i.e.:

$$\mathcal{M}_{\tau f}(\mathbf{x}) = f(\text{prox}_{\tau f}(\mathbf{x})) + \frac{1}{2\tau} \|\mathbf{x} - \text{prox}_{\tau f}(\mathbf{x})\|_2^2 \quad (\text{C.10})$$

We also remind the definition of order k pseudo-Lipschitz function.

Definition 2. *Pseudo-Lipschitz function* For $k \in \mathbb{N}^*$ and any $n, m \in \mathbb{N}^*$, a function $\phi : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is called a pseudo-Lipschitz of order k if there exists a constant $L(k)$ such that for any $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$,

$$\|\phi(\mathbf{x}) - \phi(\mathbf{y})\|_2 \leq L(k) \left(1 + (\|\mathbf{x}\|_2)^{k-1} + (\|\mathbf{y}\|_2)^{k-1} \right) \|\mathbf{x} - \mathbf{y}\|_2 \quad (\text{C.11})$$

We now give some further properties with straightforward proofs that will be helpful throughout the proof.

Lemma 2. *(Moreau envelope of pseudo-Lipschitz function)* Consider a proper, lower-semicontinuous, convex, pseudo-Lipschitz function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ of order k . Then its Moreau envelope is also pseudo-Lipschitz of order k .

Proof of Lemma 2: For any \mathbf{x}, \mathbf{y} in $\text{dom}(f)$, we have, using the pseudo-Lipschitz property:

$$|f(\text{prox}_{\tau f}(\mathbf{x})) - f(\text{prox}_{\tau f}(\mathbf{y}))| \leq L(k) \left(1 + \left(\|\text{prox}_{\tau f}(\mathbf{x})\|_2 \right)^{k-1} + \left(\|\text{prox}_{\tau f}(\mathbf{y})\|_2 \right)^{k-1} \right) \|\text{prox}_{\tau f}(\mathbf{x}) - \text{prox}_{\tau f}(\mathbf{y})\|_2 \quad (\text{C.12})$$

$$\leq L(k) \left(1 + (\|\mathbf{x}\|_2)^{k-1} + (\|\mathbf{y}\|_2)^{k-1} \right) \|\mathbf{x} - \mathbf{y}\|_2 \quad (\text{C.13})$$

where the second line follows immediately with the same constant $L(k)$ owing to the firm-nonexpansiveness of the proximal operator. Furthermore

$$\begin{aligned} & \|\mathbf{x} - \text{prox}_{\tau f}(\mathbf{x})\|_2^2 - \|\mathbf{y} - \text{prox}_{\tau f}(\mathbf{y})\|_2^2 = \\ & \tau \left| \partial f(\text{prox}_{\tau f}(\mathbf{x})) + \partial f(\text{prox}_{\tau f}(\mathbf{y})) \right| \left\| (\mathbf{x} - \text{prox}_{\tau f}(\mathbf{x}) - \mathbf{y} + \text{prox}_{\tau f}(\mathbf{y})) \right\| \\ & \leq \tau \left\| \partial f(\text{prox}_{\tau f}(\mathbf{x})) + \partial f(\text{prox}_{\tau f}(\mathbf{y})) \right\|_2 \left\| (\mathbf{x} - \text{prox}_{\tau f}(\mathbf{x}) - \mathbf{y} + \text{prox}_{\tau f}(\mathbf{y})) \right\|_2 \end{aligned} \quad (\text{C.14})$$

due to the pseudo-Lipschitz property, one has

$$\partial f(\text{prox}_{\tau f}(\mathbf{x})) \leq L(k) \left(1 + 2 \|\text{prox}_{\tau f}(\mathbf{x})\|_2^{k-1} \right) \quad (\text{C.15})$$

This, along with the firm-nonexpansiveness of $\text{Id} - \text{prox}$, concludes the proof. \square

Lemma 3. (Useful functions) For any $\mathbf{x} \in \mathbb{R}^n$, $\tau > 0$, $\theta \in \mathbb{R}$ and any proper, convex lower semi-continuous function f , define the following functions:

$$\begin{aligned} h_1 : \mathbb{R} &\rightarrow \mathbb{R} \\ \theta &\mapsto \mathbf{x}^T \text{prox}_{\tau f(\cdot)}(\theta \mathbf{x}) \end{aligned} \quad (\text{C.16})$$

$$\begin{aligned} h_2 : \mathbb{R} &\rightarrow \mathbb{R} \\ \tau &\mapsto \frac{1}{2\tau^2} \left\| \mathbf{x} - \text{prox}_{\tau f(\cdot)}(\mathbf{x}) \right\|_2^2 \end{aligned} \quad (\text{C.17})$$

$$\begin{aligned} h_3 : \mathbb{R} &\rightarrow \mathbb{R} \\ \tau &\mapsto \left\| \text{prox}_{\frac{f}{\tau}(\cdot)}\left(\frac{\mathbf{x}}{\tau}\right) \right\|_2^2 \end{aligned} \quad (\text{C.18})$$

$$\begin{aligned} h_4 : \mathbb{R} &\rightarrow \mathbb{R} \\ \tau &\mapsto \left\| \mathbf{x} - \text{prox}_{\tau f}(\mathbf{x}) \right\|_2^2 \end{aligned} \quad (\text{C.19})$$

h_1 is nondecreasing, and h_2, h_3, h_4 are nonincreasing.

Proof of Lemma 3: For any $\theta, \tilde{\theta} \in \mathbb{R}$:

$$\begin{aligned} (\theta - \tilde{\theta})(h_1(\theta) - h_1(\tilde{\theta})) &= (\theta \mathbf{x} - \tilde{\theta} \mathbf{x})^\top \left(\text{prox}_{\tau f(\cdot)}(\theta \mathbf{x}) - \text{prox}_{\tau f(\cdot)}(\tilde{\theta} \mathbf{x}) \right) \\ &\geq \left\| \text{prox}_{\tau f(\cdot)}(\theta \mathbf{x}) - \text{prox}_{\tau f(\cdot)}(\tilde{\theta} \mathbf{x}) \right\|_2^2 \\ &\geq 0 \end{aligned} \quad (\text{C.20})$$

where the inequality comes from the firm non-expansiveness of the proximal operator. Thus h_1 is nondecreasing. Since the Moreau envelope $\mathcal{M}_{\tau f}(\mathbf{x})$ is convex in τ , we have, for any $\tau, \tilde{\tau}$ in \mathbb{R}_{++}

$$(\tau - \tilde{\tau}) \left(\frac{\partial}{\partial \tau} \mathcal{M}_{\tau f}(\mathbf{x}) - \frac{\partial}{\partial \tilde{\tau}} \mathcal{M}_{\tilde{\tau} f}(\mathbf{x}) \right) \geq 0, \quad \iff \quad (\tau - \tilde{\tau})(h_2(\tilde{\tau}) - h_2(\tau)) \geq 0 \quad (\text{C.21})$$

which implies that h_2 is non-increasing.

Using the Moreau decomposition, see e.g. [64], we have:

$$h_2(\tau) = \frac{1}{2\tau^2} \left\| \mathbf{x} - \left(\mathbf{x} - \tau \text{prox}_{\frac{f^*}{\tau}}\left(\frac{\mathbf{x}}{\tau}\right) \right) \right\|_2^2 = \left\| \text{prox}_{\frac{f^*}{\tau}}\left(\frac{\mathbf{x}}{\tau}\right) \right\|_2^2 \quad (\text{C.22})$$

which is a nonincreasing function of τ . Since f is convex, we can restart this short process with the conjugate of f to obtain the desired result. Thus h_3 is nonincreasing and $(\tau - \tilde{\tau})(h_3(\tau) - h_3(\tilde{\tau})) \leq 0$.

Moving to h_4 , proving that it is nonincreasing is equivalent to proving that the following function is increasing

$$h_5(\tau) = \text{prox}_{\tau f}(\mathbf{x})^\top (2\mathbf{x} - \text{prox}_{\tau f}(\mathbf{x})) \quad (\text{C.23})$$

using the Moreau decomposition again

$$h_5(\tau) = \left(\mathbf{x} - \tau \text{prox}_{\frac{f^*}{\tau}}\left(\frac{\mathbf{x}}{\tau}\right) \right)^\top \left(\mathbf{x} + \tau \text{prox}_{\frac{f^*}{\tau}}\left(\frac{\mathbf{x}}{\tau}\right) \right) \quad (\text{C.24})$$

then, for any $\tau, \tilde{\tau}$ in \mathbb{R}_{++} :

$$(\tau - \tilde{\tau})(h_5(\tau) - h_5(\tilde{\tau})) = (\tau - \tilde{\tau}) \left(\tilde{\tau}^2 \left\| \text{prox}_{\frac{f^*}{\tilde{\tau}}}\left(\frac{\mathbf{x}}{\tilde{\tau}}\right) \right\|_2^2 - \tau^2 \left\| \text{prox}_{\frac{f^*}{\tau}}\left(\frac{\mathbf{x}}{\tau}\right) \right\|_2^2 \right) \quad (\text{C.25})$$

separating the cases $\tau \leq \tilde{\tau}$ and $\tau \geq \tilde{\tau}$, and using the result on h_3 then gives the desired result. \square

The following inequality is similar to one that appeared in one-dimensional form in [26].

Lemma 4. (A useful inequality) For any proper, lower semi-continuous convex function f , any $\mathbf{x}, \tilde{\mathbf{x}}$ in $\text{dom}(f)$, and any $\gamma, \tilde{\gamma} \in \mathbb{R}_{++}$, the following holds:

$$\begin{aligned} & (\text{prox}_{\tilde{\gamma}f}(\tilde{\mathbf{x}}) - \text{prox}_{\gamma f}(\mathbf{x}))^\top \left(\frac{\tilde{\mathbf{x}} - \mathbf{x}}{\tilde{\gamma}} - \frac{\mathbf{x}}{\gamma} - \frac{1}{2} \left(\frac{1}{\tilde{\gamma}} - \frac{1}{\gamma} \right) (\text{prox}_{\tilde{\gamma}f}(\tilde{\mathbf{x}}) + \text{prox}_{\gamma f}(\mathbf{x})) \right) \\ & \geq \left(\frac{1}{2\tilde{\gamma}} + \frac{1}{2\gamma} \right) \left\| (\text{prox}_{\tilde{\gamma}f}(\tilde{\mathbf{x}}) - \text{prox}_{\gamma f}(\mathbf{x})) \right\|_2^2 \end{aligned} \quad (\text{C.26})$$

Proof of Lemma 4 : the subdifferential of a proper convex function is a monotone operator, thus:

$$(\text{prox}_{\tilde{\gamma}f}(\tilde{\mathbf{x}}) - \text{prox}_{\gamma f}(\mathbf{x}))^\top (\partial f(\text{prox}_{\tilde{\gamma}f}(\tilde{\mathbf{x}})) - \partial f(\text{prox}_{\gamma f}(\mathbf{x}))) \geq 0 \quad (\text{C.27})$$

additionally, $\text{prox}_{\gamma f}(\mathbf{x}) = (\text{Id} + \gamma \partial f)^{-1}(\mathbf{x})$, hence:

$$\begin{aligned} \partial f(\text{prox}_{\tilde{\gamma}f}(\tilde{\mathbf{x}})) - \partial f(\text{prox}_{\gamma f}(\mathbf{x})) &= \left(\frac{\tilde{\mathbf{x}} - \mathbf{x}}{\tilde{\gamma}} - \frac{\mathbf{x}}{\gamma} - \frac{1}{\tilde{\gamma}} \text{prox}_{\tilde{\gamma}f}(\tilde{\mathbf{x}}) + \frac{1}{\gamma} \text{prox}_{\gamma f}(\mathbf{x}) \right) \\ &= \frac{\tilde{\mathbf{x}} - \mathbf{x}}{\tilde{\gamma}} - \frac{\mathbf{x}}{\gamma} - \frac{1}{\tilde{\gamma}} \text{prox}_{\tilde{\gamma}f}(\tilde{\mathbf{x}}) + \frac{1}{\gamma} \text{prox}_{\gamma f}(\mathbf{x}) - \frac{1}{2} \left(\frac{1}{\tilde{\gamma}} - \frac{1}{\gamma} \right) (\text{prox}_{\tilde{\gamma}f}(\tilde{\mathbf{x}}) + \text{prox}_{\gamma f}(\mathbf{x})) \\ &\quad + \frac{1}{2} \left(\frac{1}{\tilde{\gamma}} - \frac{1}{\gamma} \right) (\text{prox}_{\tilde{\gamma}f}(\tilde{\mathbf{x}}) + \text{prox}_{\gamma f}(\mathbf{x})) \\ &= \left(\frac{\tilde{\mathbf{x}} - \mathbf{x}}{\tilde{\gamma}} - \frac{\mathbf{x}}{\gamma} - \frac{1}{2} \left(\frac{1}{\tilde{\gamma}} - \frac{1}{\gamma} \right) (\text{prox}_{\tilde{\gamma}f}(\tilde{\mathbf{x}}) + \text{prox}_{\gamma f}(\mathbf{x})) \right) - \left(\frac{1}{2\tilde{\gamma}} + \frac{1}{2\gamma} \right) (\text{prox}_{\tilde{\gamma}f}(\tilde{\mathbf{x}}) - \text{prox}_{\gamma f}(\mathbf{x})) \end{aligned} \quad (\text{C.28})$$

which gives the desired inequality. \square

C.2.3 Useful concentration of measure elements

We begin by reminding the Gaussian-Poincaré inequality, see e.g. [65].

Proposition 1. (Gaussian Poincaré inequality)

Let $\mathbf{g} \in \mathbb{R}^n$ be a $\mathcal{N}(0, I_n)$ random vector. Then for any continuous, weakly differentiable φ , there exists a constant c such that:

$$\text{Var}[\varphi(\mathbf{g})] \leq c \mathbb{E} \left[\|\nabla \varphi(\mathbf{g})\|_2^2 \right] \quad (\text{C.29})$$

We now use this previous result to show Gaussian concentration of Moreau envelopes of appropriately scaled convex functions.

Lemma 5. (Gaussian concentration of Moreau envelopes)

Consider a proper, convex function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ verifying the scaling conditions of Assumptions C.1 and let $\mathbf{g} \in \mathbb{R}^n$ be a standard normal random vector. Then, for any parameter $\tau > 0$ and any $\epsilon > 0$, there exists a constant c such that the following holds:

$$\mathbb{P} \left(\left| \frac{1}{n} \mathcal{M}_{\tau f(\cdot)}(\mathbf{g}) - \mathbb{E} \left[\frac{1}{n} \mathcal{M}_{\tau f(\cdot)}(\mathbf{g}) \right] \right| \geq \epsilon \right) \leq \frac{c}{n\tau^2\epsilon^2} \quad (\text{C.30})$$

Proof of Lemma 5:

We start by showing that the Moreau envelope of a proper, convex function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ verifying the scaling conditions of Assumptions C.1 is integrable with respect to the Gaussian measure. Using the convexity of the optimization problem defining the Moreau envelope, and the fact that f is proper, there exists $\mathbf{z}_0 \in \mathbb{R}^n$ and a finite constant \mathcal{K} such that :

$$\begin{aligned} \frac{1}{n} \mathcal{M}_{\tau f(\cdot)}(\mathbf{g}) &\leq \frac{1}{n} f(\mathbf{z}_0) + \frac{1}{2n\tau} \|\mathbf{z}_0 - \mathbf{g}\|_2^2 \\ &\leq \mathcal{K} + \frac{1}{2n\tau} \|\mathbf{z}_0 - \mathbf{g}\|_2^2 \end{aligned} \quad (\text{C.31})$$

where the second line is clearly intergable under a multivariate Gaussian measure. Then, using Proposition 1, we get:

$$\text{Var} \left[\frac{1}{n} \mathcal{M}_{\tau f(\cdot)}(\mathbf{g}) \right] \leq \frac{c}{n^2} \mathbb{E} \left[\left\| \nabla_{\mathbf{z}} \mathcal{M}_{\tau f(\cdot)}(\mathbf{g}) \right\|_2^2 \right] \quad (\text{C.32})$$

$$= \frac{c}{n^2} \mathbb{E} \left[\left\| \frac{1}{\tau} (\mathbf{z} - \text{prox}_{\tau f}(\mathbf{g})) \right\|_2^2 \right] \quad (\text{C.33})$$

Using Proposition 12.27 and Corollary 4.3 from [64], $\mathbf{g} \rightarrow \mathbf{z} - \text{prox}_{\tau f}(\mathbf{g})$ is firmly non-expansive and:

$$\left\| \mathbf{g} - \text{prox}_{\tau f}(\mathbf{g}) \right\|_2^2 \leq \langle \mathbf{g} | \mathbf{g} - \text{prox}_{\tau f}(\mathbf{g}) \rangle \quad \text{which implies} \quad (\text{C.34})$$

$$\left\| \mathbf{g} - \text{prox}_{\tau f}(\mathbf{g}) \right\|_2^2 \leq \left\| \mathbf{g} \right\|_2^2 \quad \text{using the Cauchy-Schwarz inequality} \quad (\text{C.35})$$

then

$$\text{Var} \left[\frac{1}{n} \mathcal{M}_{\tau f(\cdot)}(\mathbf{g}) \right] \leq \frac{c}{n^2 \tau^2} \mathbb{E} \left[\left\| \mathbf{g} \right\|_2^2 \right] = \frac{c}{n \tau^2} \quad (\text{C.36})$$

Chebyshev's inequality then gives, for any $\epsilon > 0$:

$$\mathbb{P} \left(\left| \frac{1}{n} \mathcal{M}_{\tau f(\cdot)}(\mathbf{g}) - \mathbb{E} \left[\frac{1}{n} \mathcal{M}_{\tau f(\cdot)}(\mathbf{g}) \right] \right| \geq \epsilon \right) \leq \frac{c}{n \tau^2 \epsilon^2} \quad (\text{C.37})$$

□

Gaussian concentration of pseudo-Lipschitz functions of finite order can also be proven using the Gaussian Poincaré inequality to yield a bound similar to the one obtained for Moreau envelopes. We thus give the result without proof:

Lemma 6. (Concentration of pseudo-Lipschitz functions) Consider a pseudo-Lipschitz function of finite order k , $f : \mathbb{R}^n \rightarrow \mathbb{R}$. Then for any vector $\mathbf{g} \sim \mathcal{N}(0, I_n)$ and any $\epsilon > 0$, there exists a constant $C(k) > 0$ such that

$$\mathbb{P} \left(\left| \frac{1}{n} f(\mathbf{g}) - \mathbb{E} \left[\frac{1}{n} f(\mathbf{g}) \right] \right| \geq \epsilon \right) \leq \frac{L(k)C(k)}{n \epsilon^2} \quad (\text{C.38})$$

We now cite an exponential concentration lemma for separable, pseudo-Lipschitz functions of order 2, taken from [66].

Lemma 7. (Lemma B.5 from [66]) Consider a separable, pseudo-Lipschitz function of order 2, $f : \mathbb{R}^n \rightarrow \mathbb{R}$. Then for any vector $\mathbf{g} \sim \mathcal{N}(0, I_n)$ and any $\epsilon > 0$, there exists constants $C, c, c' > 0$ such that

$$\mathbb{P} \left(\left| \frac{1}{n} f(\mathbf{g}) - \mathbb{E} \left[\frac{1}{n} f(\mathbf{g}) \right] \right| \geq c' \epsilon \right) \leq C e^{-c n \epsilon^2} \quad (\text{C.39})$$

C.3 Determining a candidate primary problem, auxiliary problem and its solution.

We start with a reformulation of the problem (1.2-1.3) in order to obtain an acceptable primary problem in the framework of Theorem 4. Partitioning the Gaussian distribution, we can rewrite the matrices \mathbf{U} and \mathbf{V} in the following way, introducing the standard normal vector:

$$\begin{bmatrix} \mathbf{a} \\ \mathbf{b} \end{bmatrix} \in \mathbb{R}^{p+d} \sim \mathcal{N}(0, \mathbf{I}_{p+d}) \quad (\text{C.40})$$

We can then rewrite the vectors \mathbf{u} , \mathbf{v} and matrices U , V as:

$$\mathbf{u} = \Psi^{1/2} \mathbf{a}, \quad U = A \Psi^{1/2} \quad (\text{C.41})$$

$$\mathbf{v} = \Phi^\top \Psi^{-1/2} \mathbf{a} + \left(\Omega - \Phi^\top \Psi^{-1} \Phi \right)^{1/2} \mathbf{b}, \quad V = A \Psi^{-1/2} \Phi + B \left(\Omega - \Phi^\top \Psi^{-1} \Phi \right)^{1/2} \quad (\text{C.42})$$

where the matrices A and B have independent standard normal entries and are independent of θ_0 . The learning problem then becomes equivalent to :

$$\text{Generate labels according to : } \mathbf{y} = f_0 \left(\frac{1}{\sqrt{p}} A \Psi^{1/2} \theta_0 \right) \quad (\text{C.43})$$

$$\text{Learn according to : } \arg \min_{\mathbf{w}} g \left(\frac{1}{\sqrt{d}} \left(A \Psi^{-1/2} \Phi + B \left(\Omega - \Phi^\top \Psi^{-1} \Phi \right)^{1/2} \right) \mathbf{w}, \mathbf{y} \right) + r(\mathbf{w}) \quad (\text{C.44})$$

We are then interested in the optimal cost of the following problem

$$\min_{\mathbf{w}} \frac{1}{d} \left[g \left(\frac{1}{\sqrt{d}} \left(A \Psi^{-1/2} \Phi + B \left(\Omega - \Phi^\top \Psi^{-1} \Phi \right)^{1/2} \right) \mathbf{w}, \mathbf{y} \right) + r(\mathbf{w}) \right] \quad (\text{C.45})$$

Introducing the auxiliary variable \mathbf{z} :

$$\min_{\mathbf{w}} g \left(\frac{1}{\sqrt{d}} \left(A \Psi^{-1/2} \Phi + B \left(\Omega - \Phi^\top \Psi^{-1} \Phi \right)^{1/2} \right) \mathbf{w}, \mathbf{y} \right) + r(\mathbf{w}) \quad (\text{C.46})$$

$$\iff \min_{\mathbf{w}, \mathbf{z}} g(\mathbf{z}, \mathbf{y}) + r(\mathbf{w})$$

$$\text{s.t. } \mathbf{z} = \frac{1}{\sqrt{d}} \left(A \Psi^{-1/2} \Phi + B \left(\Omega - \Phi^\top \Psi^{-1} \Phi \right)^{1/2} \right) \mathbf{w} \quad (\text{C.47})$$

Introducing the corresponding Lagrange multiplier $\boldsymbol{\lambda} \in \mathbb{R}^n$ and using strong duality, the problem is equivalent to :

$$\min_{\mathbf{w}, \mathbf{z}} \max_{\boldsymbol{\lambda}} \boldsymbol{\lambda}^\top \frac{1}{\sqrt{d}} \left(A \Psi^{-1/2} \Phi + B \left(\Omega - \Phi^\top \Psi^{-1} \Phi \right)^{1/2} \right) \mathbf{w} - \boldsymbol{\lambda}^\top \mathbf{z} + g(\mathbf{z}, \mathbf{y}) + r(\mathbf{w}) \quad (\text{C.48})$$

In the remainder of the proof, the preceding cost function will be denoted

$$\mathbf{C}(\mathbf{w}, \mathbf{z}) = \max_{\boldsymbol{\lambda}} \boldsymbol{\lambda}^\top \frac{1}{\sqrt{d}} \left(A \Psi^{-1/2} \Phi + B \left(\Omega - \Phi^\top \Psi^{-1} \Phi \right)^{1/2} \right) \mathbf{w} - \boldsymbol{\lambda}^\top \mathbf{z} + g(\mathbf{z}, \mathbf{y}) + r(\mathbf{w}) \quad (\text{C.49})$$

such that the problem reads $\min_{\mathbf{w}, \mathbf{z}} \mathbf{C}(\mathbf{w}, \mathbf{z})$. Theorem 4 requires working with compact feasibility sets. Adopting similar approaches to the ones from [26, 31], the next lemma shows that the optimization problem (C.48) can be equivalently recast as one over compact sets.

Lemma 8. (*Compactness of feasibility set*) Let $\mathbf{w}^*, \mathbf{z}^*, \boldsymbol{\lambda}^*$ be optimal in (C.48). Then there exists positive constants $C_{\mathbf{w}}, C_{\mathbf{z}}$ and $C_{\boldsymbol{\lambda}}$ such that

$$\mathbb{P} \left(\|\mathbf{w}^*\|_2 \leq C_{\mathbf{w}} \sqrt{d} \right) \xrightarrow{d \rightarrow \infty} 1, \quad \mathbb{P} \left(\|\mathbf{z}^*\|_2 \leq C_{\mathbf{z}} \sqrt{n} \right) \xrightarrow{n \rightarrow \infty} 1, \quad \mathbb{P} \left(\|\boldsymbol{\lambda}^*\|_2 \leq C_{\boldsymbol{\lambda}} \sqrt{n} \right) \xrightarrow{n \rightarrow \infty} 1 \quad (\text{C.50})$$

Proof of Lemma 8: consider the initial minimisation problem:

$$\hat{\mathbf{w}} = \arg \min_{\mathbf{w} \in \mathbb{R}^d} g \left(\frac{1}{\sqrt{d}} V \mathbf{w}, \mathbf{y} \right) + r(\mathbf{w}) \quad (\text{C.51})$$

From assumption (A3), the cost function $g + r$ is coercive, proper and lower semi-continuous. Since it is proper, there exists $\mathbf{w}_0 \in \mathbb{R}^d$ such that $g \left(\frac{1}{\sqrt{d}} V \mathbf{w}, \mathbf{y} \right) + r(\mathbf{w}) \in \mathbb{R}$. The coercivity implies that there exists $\eta \in]0, +\infty[$ such that, for every $\mathbf{w} \in \mathbb{R}^d$ satisfying $\|\mathbf{w} - \mathbf{w}_0\| \geq \eta$, $g \left(\frac{1}{\sqrt{d}} V \mathbf{w}, \mathbf{y} \right) + r(\mathbf{w}) \geq g \left(\frac{1}{\sqrt{d}} V \mathbf{w}_0, \mathbf{y} \right) + r(\mathbf{w}_0)$. Let $S = \{\mathbf{w} \in \mathbb{R}^d \mid \|\mathbf{w} - \mathbf{w}_0\| \leq \eta\}$. Then $S \cap \mathbb{R}^d \neq \emptyset$ and S is compact. Then, there exists $\mathbf{w}^* \in S$ such that $g \left(\frac{1}{\sqrt{d}} V \mathbf{w}^*, \mathbf{y} \right) + r(\mathbf{w}^*) = \inf_{\mathbf{w} \in S} g \left(\frac{1}{\sqrt{d}} V \mathbf{w}, \mathbf{y} \right) + r(\mathbf{w}) \leq g \left(\frac{1}{\sqrt{d}} V \mathbf{w}_0, \mathbf{y} \right) + f(\mathbf{w}_0)$. Thus $g \left(\frac{1}{\sqrt{d}} V \mathbf{w}^*, \mathbf{y} \right) + r(\mathbf{w}^*) \in \inf_{\mathbf{w} \in \mathbb{R}^d} g \left(\frac{1}{\sqrt{d}} V \mathbf{w}, \mathbf{y} \right) + r(\mathbf{w})$ and the set of minimisers is bounded. Closure is immediately checked by considering a sequence of minimisers converging to \mathbf{w}^* .

We conclude that the set of minimisers of problem (C.51) is a non-empty compact set. Then there exists a constant $C_{\mathbf{w}}$ independent of the dimension d , such that:

$$\|\mathbf{w}\|_2 \leq C_{\mathbf{w}}\sqrt{d} \quad (\text{C.52})$$

Now consider the equivalent formulation of problem (C.51):

$$\min_{\mathbf{w}, \mathbf{z}} \max_{\boldsymbol{\lambda}} \boldsymbol{\lambda}^\top \frac{1}{\sqrt{d}} V \mathbf{w} - \boldsymbol{\lambda}^\top \mathbf{z} + g(\mathbf{z}, \mathbf{y}) + r(\mathbf{w}) \quad (\text{C.53})$$

Its optimality condition reads :

$$\nabla_{\boldsymbol{\lambda}} : \frac{1}{\sqrt{d}} V \mathbf{w} = \mathbf{z}, \quad \nabla_{\mathbf{z}} : \boldsymbol{\lambda} \in \partial g(\mathbf{z}, \mathbf{y}), \quad \nabla_{\mathbf{w}} : \frac{1}{\sqrt{d}} V^\top \boldsymbol{\lambda} \in \partial r(\mathbf{w}) \quad (\text{C.54})$$

The optimality condition in $\boldsymbol{\lambda}$ gives:

$$\begin{aligned} \|\mathbf{z}\|_2 &\leq \left\| \frac{1}{\sqrt{d}} V \right\|_{op} \|\mathbf{w}\|_2 \\ &\leq \left\| \frac{1}{\sqrt{d}} \left(A \Psi^{-1/2} \Phi + B \left(\Omega - \Phi^\top \Psi^{-1} \Phi \right)^{1/2} \right) \right\|_{op} \|\mathbf{w}\|_2 \\ &\leq \left[\left\| \Psi^{-1/2} \Phi \right\|_{op} \left\| \frac{1}{\sqrt{d}} A \right\|_{op} + \left\| \left(\Omega - \Phi^\top \Psi^{-1} \Phi \right)^{1/2} \right\|_{op} \left\| \frac{1}{\sqrt{d}} B \right\|_{op} \right] \|\mathbf{w}\|_2 \end{aligned} \quad (\text{C.55})$$

According to assumption (A2), the operator norms of the matrices involving the covariance matrices are almost surely bounded whatever the dimensions and using established results on random matrices, see e.g. [67], the operator norms of $\frac{1}{\sqrt{d}} A$ and $\frac{1}{\sqrt{d}} B$ are bounded by finite constants with high probability when the dimensions go to infinity. Thus there exists a constant $C_{\mathbf{z}}$ also independent of d such that:

$$\mathbb{P} \left(\|\mathbf{z}\|_2 \leq C_{\mathbf{z}} \sqrt{n} \right) \xrightarrow[n \rightarrow \infty]{P} 1 \quad (\text{C.56})$$

Finally, the scaling condition from assumption (A3) directly shows that there exists a constant $C_{\boldsymbol{\lambda}}$ such that

$$\mathbb{P} \left(\|\boldsymbol{\lambda}\|_2 \leq C_{\boldsymbol{\lambda}} \sqrt{n} \right) \xrightarrow[n \rightarrow \infty]{P} 1 \quad (\text{C.57})$$

This concludes the proof of Lemma 8. \square

Defining the sets $\mathcal{S}_{\mathbf{w}} = \{\mathbf{w} \in \mathbb{R}^d \mid \|\mathbf{w}\|_2 \leq C_{\mathbf{w}} \sqrt{d}\}$, $\mathcal{S}_{\mathbf{z}} = \{\mathbf{z} \in \mathbb{R}^n \mid \|\mathbf{z}\|_2 \leq C_{\mathbf{z}} \sqrt{n}\}$ and $\mathcal{S}_{\boldsymbol{\lambda}} = \{\boldsymbol{\lambda} \in \mathbb{R}^n \mid \|\boldsymbol{\lambda}\|_2 \leq C_{\boldsymbol{\lambda}} \sqrt{n}\}$, the optimization problem can now be reduced to:

$$\min_{\mathbf{w} \in \mathcal{S}_{\mathbf{w}}, \mathbf{z} \in \mathcal{S}_{\mathbf{z}}} \max_{\boldsymbol{\lambda} \in \mathcal{S}_{\boldsymbol{\lambda}}} \boldsymbol{\lambda}^\top \frac{1}{\sqrt{d}} \left(A \Psi^{-1/2} \Phi + B \left(\Omega - \Phi^\top \Psi^{-1} \Phi \right)^{1/2} \right) \mathbf{w} - \boldsymbol{\lambda}^\top \mathbf{z} + g(\mathbf{z}, \mathbf{y}) + r(\mathbf{w}) \quad (\text{C.58})$$

The rest of this section can then be summarized by the following lemma, the proof of which shows how to find an acceptable (PO) for problem (C.58), the corresponding (AO) and how to reduce the (AO) to a scalar optimization problem. An extra assumption is needed at this point, which will be relaxed at the end of the proof :

(B1) the distribution of the teacher vector $\boldsymbol{\theta}_0$ is a point mass in \mathbb{R}^p , i.e. it is a deterministic vector.

Lemma 9. (Scalar equivalent problem) *In the framework of Theorem 4, acceptable (AO)s of problem (C.58) can be reduced to the following scalar optimization problems*

$$\text{For } \boldsymbol{\theta}_0 \notin \text{Ker}(\Phi^\top) : \max_{\kappa, \nu, \tau_2} \min_{m, \eta, \tau_1} \mathcal{E}_n(\tau_1, \tau_2, \kappa, \eta, \nu, m) \quad (\text{C.59})$$

$$\text{For } \boldsymbol{\theta}_0 \in \text{Ker}(\Phi^\top) : \max_{\kappa, \tau_2} \min_{\eta, \tau_1} \mathcal{E}_n^0(\tau_1, \tau_2, \kappa, \eta) \quad (\text{C.60})$$

where

$$\begin{aligned} \mathcal{E}_n(\tau_1, \tau_2, \kappa, \eta, \nu, m) &= \frac{\kappa\tau_1}{2} - \frac{\eta\tau_2}{2} + m\nu\sqrt{\gamma} - \frac{\tau_2}{2\eta} \frac{m^2}{\rho} - \frac{\eta}{2\tau_2 d} (\nu\mathbf{v} + \kappa\Omega^{1/2}\mathbf{g})^\top \Omega^{-1} (\nu\mathbf{v} + \kappa\Omega^{1/2}\mathbf{g}) \\ &\quad - \kappa\mathbf{g}^\top \left(\Sigma^{1/2} - \Omega^{1/2} \right) \frac{m\sqrt{\gamma}}{\|\mathbf{v}\|_2^2} \mathbf{v} + \frac{1}{d} \mathcal{M}_{\frac{\tau_1}{\kappa}g(\cdot, \mathbf{y})} \left(\frac{m}{\sqrt{\rho}} \mathbf{s} + \eta\mathbf{h} \right) + \frac{1}{d} \mathcal{M}_{\frac{\eta}{\tau_2}r(\Omega^{-1/2, \cdot})} \left(\frac{\eta}{\tau_2} (\nu\Omega^{-1/2}\mathbf{v} + \kappa\mathbf{g}) \right) \end{aligned} \quad (\text{C.61})$$

$$\mathcal{E}_n^0(\tau_1, \tau_2, \kappa, \nu) = -\frac{\eta\tau_2}{2} + \frac{\kappa\tau_1}{2} + \frac{1}{d} \mathcal{M}_{\frac{\tau_1}{\kappa}g(\cdot, \mathbf{y})}(\eta\mathbf{h}) + \frac{1}{d} \mathcal{M}_{\frac{\eta}{\tau_2}f(\Omega^{-1/2, \cdot})} \left(\frac{\eta}{\tau_2} \kappa\mathbf{g} \right) - \frac{\eta}{2\tau_2 d} \kappa^2 \mathbf{g}^\top \mathbf{g} \quad (\text{C.62})$$

and

$$\Sigma = \Omega - \frac{\mathbf{v}\mathbf{v}^\top}{\rho p} \quad \mathbf{v} = \Phi^\top \boldsymbol{\theta}_0 \quad \rho = \frac{1}{p} \boldsymbol{\theta}_0^\top \Psi \boldsymbol{\theta}_0 \quad (\text{C.63})$$

Proof of Lemma 9: We need to find an i.i.d. Gaussian matrix independent from the rest of the problem in order to use Theorem 4. We thus decompose the mixing matrix \mathbf{A} by taking conditional expectations w.r.t. \mathbf{y} , which amounts to conditioning on a linear subset of the Gaussian space generated by \mathbf{A} . Dropping the feasibility sets for confort of notation in the following lines:

$$\begin{aligned} \min_{\mathbf{w}, \mathbf{z}} \max_{\boldsymbol{\lambda}} \boldsymbol{\lambda}^\top \frac{1}{\sqrt{d}} \left((\mathbb{E}[A|\mathbf{y}] + A - \mathbb{E}[A|\mathbf{y}]) \Psi^{-1/2} \Phi + B \left(\Omega - \Phi^\top \Psi^{-1} \Phi \right)^{1/2} \right) \mathbf{w} \\ - \boldsymbol{\lambda}^\top \mathbf{z} + g(\mathbf{z}, \mathbf{y}) + r(\mathbf{w}) \end{aligned} \quad (\text{C.64})$$

$$\begin{aligned} \iff \min_{\mathbf{w}, \mathbf{z}} \max_{\boldsymbol{\lambda}} \boldsymbol{\lambda}^\top \frac{1}{\sqrt{d}} \left((\mathbb{E}[A|A\Psi^{1/2}\boldsymbol{\theta}_0] + A - \mathbb{E}[A|A\Psi^{1/2}\boldsymbol{\theta}_0]) \Psi^{-1/2} \Phi + B \left(\Omega - \Phi^\top \Psi^{-1} \Phi \right)^{1/2} \right) \mathbf{w} \\ - \boldsymbol{\lambda}^\top \mathbf{z} + g(\mathbf{z}, \mathbf{y}) + r(\mathbf{w}) \end{aligned} \quad (\text{C.65})$$

Conditioning in Gaussian spaces amounts to doing orthogonal projections. Denoting $\tilde{\boldsymbol{\theta}}_0 = \Psi^{1/2}\boldsymbol{\theta}_0$ and \tilde{A} a copy of A independent of \mathbf{y} , the minimisation problem then becomes:

$$\begin{aligned} \min_{\mathbf{w}, \mathbf{z}} \max_{\boldsymbol{\lambda}} \boldsymbol{\lambda}^\top \frac{1}{\sqrt{d}} \left((A\mathbf{P}_{\tilde{\boldsymbol{\theta}}_0} + \tilde{A}\mathbf{P}_{\tilde{\boldsymbol{\theta}}_0}^\perp) \Psi^{-1/2} \Phi + B \left(\Omega - \Phi^\top \Psi^{-1} \Phi \right)^{1/2} \right) \mathbf{w} - \boldsymbol{\lambda}^\top \mathbf{z} + g(\mathbf{z}, \mathbf{y}) + r(\mathbf{w}) \end{aligned} \quad (\text{C.66})$$

$$\begin{aligned} \iff \min_{\mathbf{w}, \mathbf{z}} \max_{\boldsymbol{\lambda}} \boldsymbol{\lambda}^\top \frac{1}{\sqrt{d}} A\mathbf{P}_{\tilde{\boldsymbol{\theta}}_0} \Psi^{-1/2} \Phi \mathbf{w} + \boldsymbol{\lambda}^\top \frac{1}{\sqrt{d}} \tilde{A}\mathbf{P}_{\tilde{\boldsymbol{\theta}}_0}^\perp \Psi^{-1/2} \Phi \mathbf{w} + \boldsymbol{\lambda}^\top \frac{1}{\sqrt{d}} B \left(\Omega - \Phi^\top \Psi^{-1} \Phi \right)^{1/2} \mathbf{w} \\ - \boldsymbol{\lambda}^\top \mathbf{z} + g(\mathbf{z}, \mathbf{y}) + r(\mathbf{w}) \end{aligned} \quad (\text{C.67})$$

$$\begin{aligned} \iff \min_{\mathbf{w}, \mathbf{z}} \max_{\boldsymbol{\lambda}} \boldsymbol{\lambda}^\top \frac{1}{\sqrt{d}} \mathbf{s} \frac{\tilde{\boldsymbol{\theta}}_0^\top}{\|\tilde{\boldsymbol{\theta}}_0\|_2} \Psi^{-1/2} \Phi \mathbf{w} + \boldsymbol{\lambda}^\top \frac{1}{\sqrt{d}} \tilde{A}\mathbf{P}_{\tilde{\boldsymbol{\theta}}_0}^\perp \Psi^{-1/2} \Phi \mathbf{w} + \boldsymbol{\lambda}^\top \frac{1}{\sqrt{d}} B \left(\Omega - \Phi^\top \Psi^{-1} \Phi \right)^{1/2} \mathbf{w} \\ - \boldsymbol{\lambda}^\top \mathbf{z} + g(\mathbf{z}, \mathbf{y}) + r(\mathbf{w}) \end{aligned} \quad (\text{C.68})$$

where we used $\mathbf{P}_{\tilde{\boldsymbol{\theta}}_0} = \frac{\tilde{\boldsymbol{\theta}}_0 \tilde{\boldsymbol{\theta}}_0^\top}{\|\tilde{\boldsymbol{\theta}}_0\|_2^2}$ and $\mathbf{s} = A \frac{\tilde{\boldsymbol{\theta}}_0}{\|\tilde{\boldsymbol{\theta}}_0\|_2}$. Knowing that \tilde{A}, B are independent standard Gaussian matrices, and independent from $\mathbf{A}, \mathbf{y}, f_0$, we can rewrite the problem as :

$$\min_{\mathbf{w}, \mathbf{z}} \max_{\boldsymbol{\lambda}} \boldsymbol{\lambda}^\top \frac{1}{\sqrt{d}} \mathbf{s} \frac{\boldsymbol{\theta}_0^\top}{\|\Psi^{1/2}\boldsymbol{\theta}_0\|} \Phi \mathbf{w} + \boldsymbol{\lambda}^\top \frac{1}{\sqrt{d}} Z \Sigma^{1/2} \mathbf{w} - \boldsymbol{\lambda}^\top \mathbf{z} + g(\mathbf{z}, \mathbf{y}) + r(\mathbf{w}) \quad (\text{C.69})$$

where $\Sigma = \Phi^\top \Psi^{-1/2} \mathbf{P}_{\tilde{\boldsymbol{\theta}}_0}^\perp \Psi^{-1/2} \Phi + \Omega - \Phi^\top \Psi^{-1} \Phi = \Omega - \Phi^\top \Psi^{-1/2} \mathbf{P}_{\tilde{\boldsymbol{\theta}}_0} \Psi^{-1/2} \Phi$, and Z is a standard Gaussian matrix independent of $\mathbf{A}, \mathbf{y}, f_0$. Recall $\rho = \frac{1}{p} \boldsymbol{\theta}_0^\top \Psi \boldsymbol{\theta}_0$ from the main text. Replacing with the expression of $\tilde{\boldsymbol{\theta}}_0$

and letting $\mathbf{v} = \Phi^\top \boldsymbol{\theta}_0$, we have

$$\Sigma = \Omega - \frac{\phi^\top \Psi^{-1/2} \tilde{\boldsymbol{\theta}}_0 \tilde{\boldsymbol{\theta}}_0^\top \Psi^{-1/2} \Phi}{\|\tilde{\boldsymbol{\theta}}_0\|_2^2} = \Omega - \frac{\phi^\top \boldsymbol{\theta}_0 \boldsymbol{\theta}_0^\top \Phi}{\boldsymbol{\theta}_0^\top \Psi \boldsymbol{\theta}_0} \quad (\text{C.70})$$

$$= \Omega - \frac{\mathbf{v} \mathbf{v}^\top}{p\rho} \quad (\text{C.71})$$

The problem then becomes

$$\min_{\mathbf{w}, \mathbf{z}} \max_{\lambda} \lambda^\top \frac{1}{\sqrt{dp}} \mathbf{s} \frac{\mathbf{v}^\top}{\sqrt{\rho}} \mathbf{w} + \lambda^\top \frac{1}{\sqrt{d}} Z \Sigma^{1/2} \mathbf{w} - \lambda^\top \mathbf{z} + g(\mathbf{z}, \mathbf{y}) + r(\mathbf{w}) \quad (\text{C.72})$$

Two cases must now be considered, $\boldsymbol{\theta}_0 \notin \text{Ker}(\phi^\top)$ and $\boldsymbol{\theta}_0 \in \text{Ker}(\phi^\top)$. Another possible case is $\Phi = 0_{p \times d}$, however it leads to the same steps as the case $\boldsymbol{\theta}_0 \in \text{Ker}(\Phi^\top)$.

Case 1: $\boldsymbol{\theta}_0 \notin \text{Ker}(\Phi^\top)$

It is tempting to invert the matrix $\Sigma^{1/2}$ to make the change of variable $\mathbf{w}_\perp = \Sigma^{1/2} \mathbf{w}$ and continue the calculation. However there is no guarantee that Σ is invertible : it is only semi-positive definite. Taking identities everywhere gives for examples $\mathbf{P}_{\boldsymbol{\theta}_0}^\perp$ which coupled optimization problems in the rest of the computation.:

$$\min_{\mathbf{w}, \mathbf{z}, \mathbf{p}} \max_{\lambda, \mu} \lambda^\top \frac{1}{\sqrt{dp}} \mathbf{s} \frac{\mathbf{v}^\top}{\sqrt{\rho}} \mathbf{w} + \lambda^\top \frac{1}{\sqrt{d}} Z \mathbf{p} - \lambda^\top \mathbf{z} + g(\mathbf{z}, \mathbf{y}) + r(\mathbf{w}) + \boldsymbol{\mu}^\top \left(\Sigma^{1/2} \mathbf{w} - \mathbf{p} \right) \quad (\text{C.73})$$

Here the minimisation on f and g is linked by the bilinear form $\lambda^\top \mathbf{s} \mathbf{v}^\top \mathbf{w}$. We wish to separate them in order for the Moreau envelopes to appear later on in simple fashion. To do so, we introduce the orthogonal decomposition of \mathbf{w} on the direction of \mathbf{v} :

$$\begin{aligned} \mathbf{w} &= \left(\mathbf{P}_\mathbf{v} + \mathbf{P}_\mathbf{v}^\perp \right) \mathbf{w} = \frac{\mathbf{v}^\top \mathbf{w}}{\|\mathbf{v}\|_2^2} \mathbf{v} + \mathbf{P}_\mathbf{v}^\perp \mathbf{w} \\ &= \frac{\mathbf{v}^\top \mathbf{w}}{\|\mathbf{v}\|_2^2} \mathbf{v} + \mathbf{w}_\perp \text{ where } \mathbf{w}_\perp \perp \mathbf{v} \\ &= \frac{m\sqrt{dp}}{\|\mathbf{v}\|_2^2} \mathbf{v} + \mathbf{w}_\perp \quad \text{where } m = \frac{1}{\sqrt{dp}} \mathbf{v}^\top \mathbf{w} \end{aligned} \quad (\text{C.74})$$

where the parameter m corresponds to the one defined in (2.5). This gives the following, after introducing the scalar Lagrange multiplier $\nu \in \mathbb{R}$ to enforce the constraint $\mathbf{w}_\perp \perp \mathbf{v}$. Note that several methods can be used to express the orthogonality constraint, as in e.g. [31], but the one chosen here allows to complete the proof and match the replica prediction. Reintroducing the normalization, we then have the equivalent form for (C.45):

$$\begin{aligned} \min_{m, \mathbf{w}_\perp, \mathbf{z}, \mathbf{p}} \max_{\lambda, \mu, \nu} \frac{1}{d} \left[\lambda^\top \frac{m}{\sqrt{\rho}} \mathbf{s} + \lambda^\top \frac{1}{\sqrt{d}} Z \mathbf{m} - \lambda^\top \mathbf{z} + g(\mathbf{z}, \mathbf{y}) + r \left(\frac{m\sqrt{dp}}{\|\mathbf{v}\|_2^2} \mathbf{v} + \mathbf{w}_\perp \right) \right. \\ \left. + \boldsymbol{\mu}^\top \left(\Sigma^{1/2} \left(\frac{m\sqrt{dp}}{\|\mathbf{v}\|_2^2} \mathbf{v} + \mathbf{w}_\perp \right) - \mathbf{p} \right) - \nu \mathbf{v}^\top \mathbf{w}_\perp \right] \end{aligned} \quad (\text{C.75})$$

A straightforward follow-up of the previous equations shows that the feasibility set now reads :

$$\begin{aligned} \mathcal{S}_{m, \mathbf{w}_\perp, \mathbf{z}, \mathbf{p}, \lambda, \mu, \nu} = \left\{ m \in \mathbb{R}, \mathbf{w}_\perp \in \mathbb{R}^{d-1}, \mathbf{z} \in \mathbb{R}^n, \mathbf{p} \in \mathbb{R}^d, \lambda \in \mathbb{R}^n, \mu \in \mathbb{R}^d, \nu \in \mathbb{R} \mid \right. \\ \left. \sqrt{m^2 + \frac{\|\mathbf{w}_\perp\|_2^2}{d}} \leq C_{\mathbf{w}}, \|\mathbf{z}\|_2 \leq C_{\mathbf{z}} \sqrt{n}, \|\mathbf{p}\|_2 \leq \sigma_{\max}(\Sigma^{1/2}) C_{\mathbf{w}} \sqrt{d}, \|\lambda\|_2 \leq C_{\lambda} \sqrt{n} \right\} \end{aligned} \quad (\text{C.76})$$

where the boundedness of $\|\mathbf{p}\|_2$ follows immediately from the assumptions on the covariance matrices and Lemma 8. We denote $\mathcal{S}_{\mathbf{p}} = \{\mathbf{p} \in \mathbb{R}^d \mid \|\mathbf{p}\|_2 \leq C_{\mathbf{p}}\}$ for some constant $C_{\mathbf{p}} \geq \sigma_{\max}(\Sigma^{1/2})C_{\mathbf{w}}$.

The set $\mathcal{S}_{\mathbf{p}} \times \mathcal{S}_{\lambda}$ is compact and the matrix Z is independent of all other random quantities of the problem, thus problem (C.75) is an acceptable (PO). We can now write the auxiliary optimization problem (AO) corresponding to the primary one (C.75), dropping the feasibility sets again for convenience:

$$\begin{aligned} \min_{m, \mathbf{w}_{\perp}, \mathbf{z}, \mathbf{p}} \max_{\lambda, \mu, \nu} \frac{1}{d} & \left[\lambda^{\top} \frac{m}{\sqrt{\rho}} \mathbf{s} + \frac{1}{\sqrt{d}} \|\lambda\|_2 \mathbf{g}^{\top} \mathbf{p} + \frac{1}{\sqrt{d}} \|\mathbf{p}\|_2 \mathbf{h}^{\top} \lambda - \lambda^{\top} \mathbf{z} + g(\mathbf{z}, \mathbf{y}) + r \left(\frac{m\sqrt{dp}}{\|\mathbf{v}\|_2^2} \mathbf{v} + \mathbf{w}_{\perp} \right) \right. \\ & \left. + \mu^{\top} \left(\Sigma^{1/2} \left(\frac{m\sqrt{dp}}{\|\mathbf{v}\|_2^2} \mathbf{v} + \mathbf{w}_{\perp} \right) - \mathbf{p} \right) - \nu \mathbf{v}^{\top} \mathbf{w}_{\perp} \right] \end{aligned} \quad (\text{C.77})$$

We now turn to the simplification of this problem.

The variable λ only appears in linear terms, we can thus directly optimize over its direction, introducing the positive scalar variable $\kappa = \|\lambda\|_2 / \sqrt{d}$:

$$\begin{aligned} \min_{m, \mathbf{w}_{\perp}, \mathbf{z}, \mathbf{p}} \max_{\kappa, \mu, \nu} \frac{1}{d} & \left[\kappa \mathbf{g}^{\top} \mathbf{p} + \kappa \left\| \frac{m}{\sqrt{\rho}} \sqrt{d} \mathbf{s} + \|\mathbf{p}\|_2 \mathbf{h} - \sqrt{d} \mathbf{z} \right\|_2 + g(\mathbf{z}, \mathbf{y}) + r \left(\frac{m\sqrt{dp}}{\|\mathbf{v}\|_2^2} \mathbf{v} + \mathbf{w}_{\perp} \right) \right. \\ & \left. + \mu^{\top} \left(\Sigma^{1/2} \left(\frac{m\sqrt{dp}}{\|\mathbf{v}\|_2^2} \mathbf{v} + \mathbf{w}_{\perp} \right) - \mathbf{p} \right) - \nu \mathbf{v}^{\top} \mathbf{w}_{\perp} \right] \end{aligned} \quad (\text{C.78})$$

The previous expression may not be convex-concave because of the term $\|\mathbf{p}\|_2 \mathbf{h}$. However, it was shown in [26] that the order of the min and max can still be inverted in this case, because of the convexity of the original problem. As the proof would be very similar, we do not reproduce it. Inverting the max-min order and performing the linear optimization on \mathbf{p} with $\eta = \|\mathbf{p}\|_2 / \sqrt{d}$:

$$\begin{aligned} \max_{\kappa, \mu, \nu} \min_{m, \mathbf{w}_{\perp}, \mathbf{z}, \eta} & \left\{ -\frac{\eta}{\sqrt{d}} \|\mu + \kappa \mathbf{g}\|_2 + \frac{\kappa}{\sqrt{d}} \left\| \frac{m}{\sqrt{\rho}} \mathbf{s} + \eta \mathbf{h} - \mathbf{z} \right\|_2 + \right. \\ & \left. + \frac{1}{d} \left[g(\mathbf{z}, \mathbf{y}) + r \left(\frac{m\sqrt{dp}}{\|\mathbf{v}\|_2^2} \mathbf{v} + \mathbf{w}_{\perp} \right) + \mu^{\top} \Sigma^{1/2} \left(\frac{m\sqrt{dp}}{\|\mathbf{v}\|_2^2} \mathbf{v} + \mathbf{w}_{\perp} \right) - \nu \mathbf{v}^{\top} \mathbf{w}_{\perp} \right] \right\} \end{aligned} \quad (\text{C.79})$$

using the following representation of the norm, as in [26], for any vector t , $\|t\|_2 = \min_{\tau > 0} \frac{\tau}{2} + \frac{\|t\|_2^2}{2\tau}$:

$$\begin{aligned} \max_{\kappa, \mu, \nu, \tau_2} \min_{m, \mathbf{w}_{\perp}, \mathbf{z}, \eta, \tau_1} & \left\{ \frac{\kappa \tau_1}{2} - \frac{\eta \tau_2}{2} - \frac{\eta}{2\tau_2 d} \|\mu + \kappa \mathbf{g}\|_2^2 + \frac{\kappa}{2\tau_1 d} \left\| \frac{m}{\sqrt{\rho}} \mathbf{s} + \eta \mathbf{h} - \mathbf{z} \right\|_2^2 \right. \\ & \left. + \frac{1}{d} \left[g(\mathbf{z}, \mathbf{y}) + r \left(\frac{m\sqrt{dp}}{\|\mathbf{v}\|_2^2} \mathbf{v} + \mathbf{w}_{\perp} \right) + \mu^{\top} \Sigma^{1/2} \left(\frac{m\sqrt{dp}}{\|\mathbf{v}\|_2^2} \mathbf{v} + \mathbf{w}_{\perp} \right) - \nu \mathbf{v}^{\top} \mathbf{w}_{\perp} \right] \right\} \end{aligned} \quad (\text{C.80})$$

performing the minimisation over \mathbf{z} and recognizing the Moreau envelope of $g(\cdot, \mathbf{y})$:

$$\begin{aligned} \max_{\kappa, \mu, \nu, \tau_2} \min_{m, \mathbf{w}_{\perp}, \eta, \tau_1} & \left\{ \frac{\kappa \tau_1}{2} - \frac{\eta \tau_2}{2} + \frac{1}{d} \mathcal{M}_{\frac{\tau_1}{\kappa}} g(\cdot, \mathbf{y}) \left(\frac{m}{\sqrt{\rho}} \mathbf{s} + \beta \mathbf{h} \right) - \frac{\eta}{2\tau_2 d} \|\mu + \kappa \mathbf{g}\|_2^2 \right. \\ & \left. + \frac{1}{d} \left[r \left(\frac{m\sqrt{dp}}{\|\mathbf{v}\|_2^2} \mathbf{v} + \mathbf{w}_{\perp} \right) + \mu^{\top} \Sigma^{1/2} \left(\frac{m\sqrt{dp}}{\|\mathbf{v}\|_2^2} \mathbf{v} + \mathbf{w}_{\perp} \right) - \nu \mathbf{v}^{\top} \mathbf{w}_{\perp} \right] \right\} \end{aligned} \quad (\text{C.81})$$

At this point we have a convex-concave problem. Inverting the min-max order, $\boldsymbol{\mu}$ appears in a well defined strictly convex least-square problem.

$$\begin{aligned} & \max_{\kappa, \nu, \tau_2} \min_{m, \mathbf{w}_\perp, \eta, \tau_1} \frac{\kappa\tau_1}{2} - \frac{\eta\tau_2}{2} + \frac{1}{d} \mathcal{M}_{\frac{\tau_1}{\kappa} g(\cdot, \mathbf{y})} \left(\frac{m}{\sqrt{\rho}} \mathbf{s} + \eta \mathbf{h} \right) - \frac{\nu}{d} \mathbf{v}^\top \mathbf{w}_\perp + \frac{1}{d} r \left(\frac{m\sqrt{dp}}{\|\mathbf{v}\|_2^2} \mathbf{v} + \mathbf{w}_\perp \right) \\ & + \frac{1}{d} \max_{\boldsymbol{\mu}} \left\{ -\frac{\eta}{2\tau_2} \|\boldsymbol{\mu} + \kappa \mathbf{g}\|_2^2 + \boldsymbol{\mu}^\top \Sigma^{1/2} \left(\frac{m\sqrt{dp}}{\|\mathbf{v}\|_2^2} \mathbf{v} + \mathbf{w}_\perp \right) \right\} \end{aligned} \quad (\text{C.82})$$

Solving it:

$$\begin{aligned} & \max_{\boldsymbol{\mu}} \left\{ -\frac{\eta}{2\tau_2} \|\boldsymbol{\mu} + \kappa \mathbf{g}\|_2^2 + \boldsymbol{\mu}^\top \Sigma^{1/2} \left(\frac{m\sqrt{dp}}{\|\mathbf{v}\|_2^2} \mathbf{v} + \mathbf{w}_\perp \right) \right\} \\ & \boldsymbol{\mu}^* = \frac{\tau_2}{\eta} \Sigma^{1/2} \left(\frac{m\sqrt{dp}}{\|\mathbf{v}\|_2^2} \mathbf{v} + \mathbf{w}_\perp \right) - \kappa \mathbf{g} \\ & \text{with optimal cost } \frac{\tau_2}{2\eta} \left\| \Sigma^{1/2} \left(\frac{m\sqrt{dp}}{\|\mathbf{v}\|_2^2} \mathbf{v} + \mathbf{w}_\perp \right) \right\|_2^2 - \kappa \mathbf{g}^\top \Sigma^{1/2} \left(\frac{m\sqrt{dp}}{\|\mathbf{v}\|_2^2} \mathbf{v} + \mathbf{w}_\perp \right) \end{aligned} \quad (\text{C.83})$$

remembering that $\Sigma = \Omega - \mathbf{v}\mathbf{v}^\top / (p\rho)$ and $\mathbf{w}_\perp \perp \mathbf{v}$, the optimal cost of this least-square problem simplifies to:

$$c^* = \frac{\tau_2}{2\eta} \left(\left\| \Omega^{1/2} \left(\frac{m\sqrt{dp}}{\|\mathbf{v}\|_2^2} \mathbf{v} + \mathbf{w}_\perp \right) \right\|_2^2 - \frac{m^2}{\rho} d \right) - \kappa \mathbf{g}^\top \Sigma^{1/2} \left(\frac{m\sqrt{dp}}{\|\mathbf{v}\|_2^2} \mathbf{v} + \mathbf{w}_\perp \right) \quad (\text{C.84})$$

The (AO) then reads :

$$\begin{aligned} & \max_{\kappa, \nu, \tau_2} \min_{m, \mathbf{w}_\perp, \eta, \tau_1} \left\{ \frac{\kappa\tau_1}{2} - \frac{\eta\tau_2}{2} + \frac{1}{d} \mathcal{M}_{\frac{\tau_1}{\kappa} g(\cdot, \mathbf{y})} \left(\frac{m}{\sqrt{\rho}} \mathbf{s} + \eta \mathbf{h} \right) - \frac{\tau_2}{2\eta} \frac{m^2}{\rho} - \frac{\nu}{d} \mathbf{v}^\top \mathbf{w}_\perp \right. \\ & \left. + \frac{1}{d} r \left(\frac{m\sqrt{dp}}{\|\mathbf{v}\|_2^2} \mathbf{v} + \mathbf{w}_\perp \right) + \frac{\tau_2}{2\eta d} \left\| \Omega^{1/2} \left(\frac{m\sqrt{dp}}{\|\mathbf{v}\|_2^2} \mathbf{v} + \mathbf{w}_\perp \right) \right\|_2^2 - \frac{\kappa}{d} \mathbf{g}^\top \Sigma^{1/2} \left(\frac{m\sqrt{dp}}{\|\mathbf{v}\|_2^2} \mathbf{v} + \mathbf{w}_\perp \right) \right\} \end{aligned} \quad (\text{C.85})$$

We now need to solve in \mathbf{w}_\perp . To do so, we can replace r with its convex conjugate and solve the least-square problem in \mathbf{w}_\perp . This will lead to a Moreau envelope of r^* in the introduced dual variable, which can be linked to the Moreau envelope of r by Moreau decomposition. Intuitively, it is natural to think that the corresponding primal variable will be $\frac{m\sqrt{dp}}{\|\mathbf{v}\|_2^2} \mathbf{v} + \mathbf{w}_\perp = \mathbf{w}$ for any feasible m, \mathbf{w}_\perp . However, we would like to have an explicit follow-up of the variables we optimize on, as we had for the Moreau envelope of g which is defined with \mathbf{z} , so we prefer to introduce a slack variable $\mathbf{w}' = \frac{m\sqrt{dp}}{\|\mathbf{v}\|_2^2} \mathbf{v} + \mathbf{w}_\perp$ with corresponding dual parameter $\boldsymbol{\eta}$ to show that the (AO) can be reformulated in terms of the original variable \mathbf{w} . Note that the feasibility set on \mathbf{w}' is almost surely compact.

$$\begin{aligned} & \max_{\kappa, \nu, \tau_2, \boldsymbol{\eta}} \min_{m, \mathbf{w}_\perp, \mathbf{w}', \eta, \tau_1} \frac{\kappa\tau_1}{2} - \frac{\eta\tau_2}{2} + \frac{1}{d} \mathcal{M}_{\frac{\tau_1}{\kappa} g(\cdot, \mathbf{y})} \left(\frac{m}{\sqrt{\rho}} \mathbf{s} + \eta \mathbf{h} \right) + \frac{1}{d} r(\mathbf{w}') - \frac{1}{d} \boldsymbol{\eta}^\top \mathbf{w}' - \frac{\tau_2}{2\eta} \frac{m^2}{\rho} \\ & - \frac{\nu}{d} \mathbf{v}^\top \mathbf{w}_\perp + \frac{\tau_2}{2\eta d} \left\| \Omega^{1/2} \left(\frac{m\sqrt{dp}}{\|\mathbf{v}\|_2^2} \mathbf{v} + \mathbf{w}_\perp \right) \right\|_2^2 - \frac{\kappa}{d} \mathbf{g}^\top \Sigma^{1/2} \left(\frac{m\sqrt{dp}}{\|\mathbf{v}\|_2^2} \mathbf{v} + \mathbf{w}_\perp \right) + \frac{1}{d} \boldsymbol{\eta}^\top \left(\frac{m\sqrt{dp}}{\|\mathbf{v}\|_2^2} \mathbf{v} + \mathbf{w}_\perp \right) \end{aligned} \quad (\text{C.86})$$

Isolating the terms depending on \mathbf{w}_\perp , we get a strictly convex least-square problem, remembering that $\Omega \in \mathcal{S}_d^{++}$:

$$\begin{aligned} & \max_{\kappa, \nu, \tau_2, \eta} \min_{m, \mathbf{w}_\perp, \mathbf{w}', \eta, \tau_1} \frac{\kappa\tau_1}{2} - \frac{\eta\tau_2}{2} + \frac{1}{d} \mathcal{M}_{\frac{\tau_1}{\kappa}g(\cdot, \mathbf{y})} \left(\frac{m}{\sqrt{\rho}} \mathbf{s} + \eta \mathbf{h} \right) + \frac{1}{d} r(\mathbf{w}') - \frac{1}{d} \boldsymbol{\eta}^T \mathbf{w}' - \frac{\tau_2}{2\eta} \frac{m^2}{\rho} + \boldsymbol{\eta}^\top \frac{m\sqrt{\kappa_2}}{\|\mathbf{v}\|_2^2} \mathbf{v} \\ & - \kappa \mathbf{g}^\top \Sigma^{1/2} \frac{m\sqrt{\gamma}}{\|\mathbf{v}\|_2^2} \mathbf{v} - \frac{\nu}{d} \mathbf{v}^\top \mathbf{w}_\perp + \frac{\tau_2}{2\eta d} \left\| \Omega^{1/2} \left(\frac{m\sqrt{dp}}{\|\mathbf{v}\|_2^2} \mathbf{v} + \mathbf{w}_\perp \right) \right\|_2^2 - \frac{\kappa}{d} \mathbf{g}^\top \Sigma^{1/2} \mathbf{w}_\perp + \frac{1}{d} \boldsymbol{\eta}^\top \mathbf{w}_\perp \end{aligned} \quad (\text{C.87})$$

$$\begin{aligned} & \max_{\kappa, \nu, \tau_2, \eta} \min_{m, \mathbf{w}', \eta, \tau_1} \frac{\kappa\tau_1}{2} - \frac{\eta\tau_2}{2} + \frac{1}{d} \mathcal{M}_{\frac{\tau_1}{\kappa}g(\cdot, \mathbf{y})} \left(\frac{m}{\sqrt{\rho}} \mathbf{s} + \eta \mathbf{h} \right) + \frac{1}{d} r(\mathbf{w}') - \frac{1}{d} \boldsymbol{\eta}^T \mathbf{w}' - \frac{\tau_2}{2\eta} \frac{m^2}{\rho} + \boldsymbol{\eta}^\top \frac{m\sqrt{\kappa_2}}{\|\mathbf{v}\|_2^2} \mathbf{v} \\ & - \kappa \mathbf{g}^\top \Sigma^{1/2} \frac{m\sqrt{\gamma}}{\|\mathbf{v}\|_2^2} \mathbf{v} + \frac{1}{d} \left[\min_{\mathbf{w}_\perp} \frac{\tau_2}{2\eta} \left\| \Omega^{1/2} \left(\frac{m\sqrt{dp}}{\|\mathbf{v}\|_2^2} \mathbf{v} + \mathbf{w}_\perp \right) \right\|_2^2 - \mathbf{w}_\perp^\top (\kappa \Sigma^{1/2} \mathbf{g} - \boldsymbol{\eta} + \nu \mathbf{v}) \right] \end{aligned} \quad (\text{C.88})$$

The quantity $\mathbf{g}^\top \Sigma^{1/2} \mathbf{w}_\perp$ is a Gaussian random variable with variance $\|\Sigma^{1/2} \mathbf{w}_\perp\|_2^2 = \mathbf{w}_\perp^\top (\Omega - \nu \mathbf{v} \mathbf{v}^\top / (p\rho)) \mathbf{w}_\perp = \mathbf{w}_\perp^\top \Omega \mathbf{w}_\perp = \|\Omega^{1/2} \mathbf{w}_\perp\|_2^2$ using the expression of Σ and the orthogonality of \mathbf{w}_\perp with respect to \mathbf{v} . We can thus change $\Sigma^{1/2}$ for $\Omega^{1/2}$ in front of \mathbf{w}_\perp combined with \mathbf{g} . The least-square problem, its solution and optimal cost then read:

$$\min_{\mathbf{w}_\perp} \frac{\tau_2}{2\eta} \left\| \Omega^{1/2} \left(\frac{m\sqrt{dp}}{\|\mathbf{v}\|_2^2} \mathbf{v} + \mathbf{w}_\perp \right) \right\|_2^2 - \mathbf{w}_\perp^\top (\kappa \Omega^{1/2} \mathbf{g} - \boldsymbol{\eta} + \nu \mathbf{v}) \quad (\text{C.89})$$

$$\mathbf{w}_\perp^* = \frac{\eta}{\tau_2} \Omega^{-1} (\kappa \Omega^{1/2} \mathbf{g} - \boldsymbol{\eta} + \nu \mathbf{v}) - \frac{m\sqrt{dp}}{\|\mathbf{v}\|_2^2} \mathbf{v} \quad (\text{C.90})$$

$$\text{with optimal cost } -\frac{\eta}{2\tau_2} (\kappa \Omega^{1/2} \mathbf{g} - \boldsymbol{\eta} + \nu \mathbf{v})^\top \Omega^{-1} (\kappa \Omega^{1/2} \mathbf{g} - \boldsymbol{\eta} + \nu \mathbf{v}) + \frac{m\sqrt{dp}}{\|\mathbf{v}\|_2^2} \mathbf{v}^\top (\kappa \Omega^{1/2} \mathbf{g} - \boldsymbol{\eta} + \nu \mathbf{v}) \quad (\text{C.91})$$

replacing in the (AO) and simplifying :

$$\begin{aligned} & \iff \max_{\kappa, \nu, \tau_2, \eta} \min_{m, \mathbf{w}', \eta, \tau_1} \frac{\kappa\tau_1}{2} - \frac{\eta\tau_2}{2} + \frac{1}{d} \mathcal{M}_{\frac{\tau_1}{\kappa}g(\cdot, \mathbf{y})} \left(\frac{m}{\sqrt{\rho}} \mathbf{s} + \eta \mathbf{h} \right) + \frac{1}{d} r(\mathbf{w}') - \frac{1}{d} \boldsymbol{\eta}^T \mathbf{w}' - \frac{\tau_2}{2\eta} \frac{m^2}{\rho} \\ & - \kappa \mathbf{g}^\top (\Sigma^{1/2} - \Omega^{1/2}) \frac{m\sqrt{\gamma}}{\|\mathbf{v}\|_2^2} \mathbf{v} - \frac{\eta}{2\tau_2 d} (\kappa \Omega^{1/2} \mathbf{g} - \boldsymbol{\eta} + \nu \mathbf{v})^\top \Omega^{-1} (\kappa \Omega^{1/2} \mathbf{g} - \boldsymbol{\eta} + \nu \mathbf{v}) + m\nu\sqrt{\gamma} \end{aligned} \quad (\text{C.92})$$

Another strictly convex least-square problem appears on $\boldsymbol{\eta}$, the solution and optimal value of which read

$$\boldsymbol{\eta}^* = -\frac{\tau_2}{\eta} \Omega \mathbf{w}' + (\kappa \Omega^{1/2} \mathbf{g} + \nu \mathbf{v}) \quad (\text{C.93})$$

$$\text{with optimal cost } \frac{\tau_2}{2\eta d} \mathbf{w}'^\top \Omega \mathbf{w}' - \mathbf{w}'^\top (\kappa \Omega^{1/2} \mathbf{g} + \nu \mathbf{v}) \quad (\text{C.94})$$

At this point we have expressed feasible solutions of $\boldsymbol{\eta}$, \mathbf{w}_\perp as functions of the remaining variables. For any feasible solution in those variables, \mathbf{w} and \mathbf{w}' are the same. Replacing in the (AO) and a completion of squares leads to

$$\begin{aligned} & \max_{\kappa, \nu, \tau_2} \min_{m, \eta, \tau_1} \frac{\kappa\tau_1}{2} - \frac{\eta\tau_2}{2} + m\nu\sqrt{\gamma} - \frac{\tau_2}{2\eta} \frac{m^2}{\rho} - \frac{\eta}{2\tau_2 d} (\nu \mathbf{v} + \kappa \Omega^{1/2} \mathbf{g})^\top \Omega^{-1} (\nu \mathbf{v} + \kappa \Omega^{1/2} \mathbf{g}) \\ & - \kappa \mathbf{g}^\top (\Sigma^{1/2} - \Omega^{1/2}) \frac{m\sqrt{\gamma}}{\|\mathbf{v}\|_2^2} \mathbf{v} + \min_{\mathbf{w}'} \left\{ r(\mathbf{w}') + \frac{\tau_2}{2\eta} \left\| \Omega^{1/2} \mathbf{w}' - \frac{\eta}{\tau_2} (\nu \Omega^{-1/2} \mathbf{v} + \kappa \mathbf{g}) \right\|_2^2 \right\} \end{aligned} \quad (\text{C.95})$$

Recognizing the Moreau envelope of f and introducing the variable $\tilde{\mathbf{w}} = \Omega^{1/2}\mathbf{w}' = \Omega^{1/2}\mathbf{w}$, it follows:

$$\begin{aligned} & \max_{\kappa, \nu, \tau_2} \min_{m, \eta, \tau_1} \frac{\kappa\tau_1}{2} - \frac{\eta\tau_2}{2} + m\nu\sqrt{\gamma} - \frac{\tau_2}{2\eta} \frac{m^2}{\rho} - \frac{\eta}{2\tau_2 d} (\nu\mathbf{v} + \kappa\Omega^{1/2}\mathbf{g})^\top \Omega^{-1} (\nu\mathbf{v} + \kappa\Omega^{1/2}\mathbf{g}) \\ & - \kappa\mathbf{g}^\top \left(\Sigma^{1/2} - \Omega^{1/2} \right) \frac{m\sqrt{\gamma}}{\|\mathbf{v}\|_2^2} \mathbf{v} + \frac{1}{d} \mathcal{M}_{\frac{\tau_1}{\kappa}g(\cdot, \mathbf{y})} \left(\frac{m}{\sqrt{\rho}} \mathbf{s} + \eta\mathbf{h} \right) + \frac{1}{d} \mathcal{M}_{\frac{\eta}{\tau_2}r(\Omega^{-1/2, \cdot})} \left(\frac{\eta}{\tau_2} \left(\nu\Omega^{-1/2}\mathbf{v} + \kappa\mathbf{g} \right) \right) \end{aligned} \quad (\text{C.96})$$

where the Moreau envelopes of f and g are respectively defined w.r.t. the variables \mathbf{w}'' and \mathbf{z} . At this point we have reduced the initial high-dimensional minimisation problem (C.77) to a scalar problem over six parameters. Another straightforward follow-up of the feasibility set shows that there exist positive constants C_m, C_κ, C_η independent of n, p, d such that $0 \leq \kappa \leq C_\kappa, 0 \leq \eta \leq C_\eta$ and $0 \leq m \leq C_m$.

Case 2: $\theta_0 \in \mathbf{Ker}(\Phi^\top)$ In this case, the min-max problem (C.72) becomes:

$$\min_{\mathbf{w}, \mathbf{z}} \max_{\lambda} \lambda^\top \frac{1}{\sqrt{d}} Z \Omega^{1/2} \mathbf{w} - \lambda^\top \mathbf{z} + g(\mathbf{z}, \mathbf{y}) + f(\mathbf{w}) \quad (\text{C.97})$$

Since Ω is positive definite, we can define $\tilde{\mathbf{w}} = \Omega^{1/2}\mathbf{w}$ and write the equivalent problem:

$$\min_{\tilde{\mathbf{w}}, \mathbf{z}} \max_{\lambda} \lambda^\top \frac{1}{\sqrt{d}} Z \tilde{\mathbf{w}} - \lambda^\top \mathbf{z} + g(\mathbf{z}, \mathbf{y}) + f(\Omega^{-1/2}\tilde{\mathbf{w}}) \quad (\text{C.98})$$

where the compactness of the feasibility set is preserved almost surely from the almost sure boundedness of the eigenvalues of Ω . We can thus write the corresponding auxiliary optimization problem, reintroducing the normalization by d :

$$\min_{\tilde{\mathbf{w}}, \mathbf{z}} \max_{\lambda} \frac{1}{d} \left[\|\lambda\|_2 \frac{1}{\sqrt{d}} \mathbf{g}^\top \tilde{\mathbf{w}} + \|\tilde{\mathbf{w}}\|_2 \frac{1}{\sqrt{d}} \mathbf{h}^\top \lambda - \lambda^\top \mathbf{z} + g(\mathbf{z}, \mathbf{y}) + f(\Omega^{-1/2}\tilde{\mathbf{w}}) \right] \quad (\text{C.99})$$

introducing the convex conjugate of f with dual parameter η :

$$\min_{\tilde{\mathbf{w}}, \mathbf{z}} \max_{\lambda, \eta} \frac{1}{d} \left[\|\lambda\|_2 \frac{1}{\sqrt{d}} \mathbf{g}^\top \tilde{\mathbf{w}} + \|\mathbf{w}_\perp\|_2 \frac{1}{\sqrt{d}} \mathbf{h}^\top \lambda - \lambda^\top \mathbf{z} + g(\mathbf{z}, \mathbf{y}) + \eta^\top \Omega^{-1/2} \tilde{\mathbf{w}} - f^*(\eta) \right] \quad (\text{C.100})$$

We then define the scalar quantities $\kappa = \frac{\|\lambda\|_2}{\sqrt{d}}$ and $\eta = \frac{\|\tilde{\mathbf{w}}\|_2}{\sqrt{d}}$ and perform the linear optimization on $\lambda, \tilde{\mathbf{w}}$, giving the equivalent:

$$\min_{\mathbf{z}, \eta \geq 0} \max_{\kappa \geq 0} -\frac{\eta}{\sqrt{d}} \left\| \kappa \mathbf{g} - \Omega^{-1/2} \eta \right\|_2 + \frac{\kappa}{\sqrt{d}} \|\eta \mathbf{h} - \mathbf{z}\|_2 + \frac{1}{d} g(\mathbf{z}, \mathbf{y}) - \frac{1}{d} f^*(\eta) \quad (\text{C.101})$$

Using the square root trick with parameters τ_1, τ_2 :

$$\min_{\tau_1 > 0, \mathbf{z}, \eta \geq 0} \max_{\tau_2 > 0, \eta, \kappa \geq 0} -\frac{\eta\tau_2}{2} - \frac{\eta}{2\tau_2 d} \left\| \kappa \mathbf{g} - \Omega^{-1/2} \eta \right\|_2^2 + \frac{\kappa\tau_1}{2} + \frac{\kappa}{2\tau_1 d} \|\eta \mathbf{h} - \mathbf{z}\|_2^2 + \frac{1}{d} g(\mathbf{z}, \mathbf{y}) - \frac{1}{d} f^*(\eta) \quad (\text{C.102})$$

performing the optimizations on \mathbf{z}, η and recognizing the Moreau envelopes, the problem becomes:

$$\min_{\tau_1 > 0, \eta \geq 0} \max_{\tau_2 > 0, \kappa \geq 0} -\frac{\eta\tau_2}{2} + \frac{\kappa\tau_1}{2} + \frac{1}{d} \mathcal{M}_{\frac{\tau_1}{\kappa}g(\cdot, \mathbf{y})}(\eta \mathbf{h}) - \frac{1}{d} \mathcal{M}_{\frac{\tau_2}{\eta}f^*(\Omega^{1/2, \cdot})}(\kappa \mathbf{g}) \quad (\text{C.103})$$

$$\iff \min_{\tau_1 > 0, \eta \geq 0} \max_{\tau_2 > 0, \kappa \geq 0} -\frac{\eta\tau_2}{2} + \frac{\kappa\tau_1}{2} + \frac{1}{d} \mathcal{M}_{\frac{\tau_1}{\kappa}g(\cdot, \mathbf{y})}(\eta \mathbf{h}) + \frac{1}{d} \mathcal{M}_{\frac{\eta}{\tau_2}f(\Omega^{-1/2, \cdot})} \left(\frac{\eta}{\tau_2} \kappa \mathbf{g} \right) - \frac{\eta}{2\tau_2 d} \kappa^2 \mathbf{g}^\top \mathbf{g} \quad (\text{C.104})$$

This concludes the proof of Lemma 9. \square

C.4 Study of the scalar equivalent problem : geometry and asymptotics.

Here we study in detail the geometry, solutions and asymptotics of the scalar optimization problem (C.96). We will focus on the case $\theta_0 \notin \text{Ker}(\Phi^\top)$ as the other case simply shows that no learning is performed (see the remark at the end of this section). The following lemma characterizes the continuity and geometry of the cost function \mathcal{E}_n .

Lemma 10. (Geometry of \mathcal{E}_n) Recall the function:

$$\begin{aligned} \mathcal{E}_n(\tau_1, \tau_2, \kappa, \eta, \nu, m) &= \frac{\kappa\tau_1}{2} - \frac{\eta\tau_2}{2} + m\nu\sqrt{\gamma} - \frac{\tau_2}{2\eta} \frac{m^2}{\rho} - \frac{\eta}{2\tau_2 d} (\nu\mathbf{v} + \kappa\Omega^{1/2}\mathbf{g})^\top \Omega^{-1} (\nu\mathbf{v} + \kappa\Omega^{1/2}\mathbf{g}) \\ &\quad - \kappa\mathbf{g}^\top \left(\Sigma^{1/2} - \Omega^{1/2} \right) \frac{m\sqrt{\gamma}}{\|\mathbf{v}\|_2^2} \mathbf{v} + \frac{1}{d} \mathcal{M}_{\frac{\tau_1}{\kappa}g(\cdot, \mathbf{y})} \left(\frac{m}{\sqrt{\rho}} \mathbf{s} + \eta \mathbf{h} \right) + \frac{1}{d} \mathcal{M}_{\frac{\eta}{\tau_2}r(\Omega^{-1/2, \cdot})} \left(\frac{\eta}{\tau_2} \left(\nu\Omega^{-1/2}\mathbf{v} + \kappa\mathbf{g} \right) \right) \end{aligned} \quad (\text{C.105})$$

Then $\mathcal{E}_n(\tau_1, \tau_2, \kappa, \eta, \nu, m)$ is continuous on its domain, jointly convex in (m, η, τ_1) and jointly concave in (κ, ν, τ_2) .

Proof of Lemma 10: $\mathcal{E}_n(\tau_1, \tau_2, \kappa, \eta, \nu, m)$ is a linear combination of linear and quadratic terms with Moreau envelopes, which are all continuous on their domain. Remembering the formulation

$$\begin{aligned} \mathcal{E}_n(\tau_1, \tau_2, \kappa, \eta, \nu, m) &= \frac{\kappa\tau_1}{2} - \frac{\eta\tau_2}{2} + m\nu\sqrt{\gamma} - \frac{\tau_2}{2\eta} \frac{m^2}{\rho} - \kappa\mathbf{g}^\top \left(\Sigma^{1/2} - \Omega^{1/2} \right) \frac{m\sqrt{\gamma}}{\|\mathbf{v}\|_2^2} \mathbf{v} \\ &\quad + \frac{1}{d} \mathcal{M}_{\frac{\tau_1}{\kappa}g(\cdot, \mathbf{y})} \left(\frac{m}{\sqrt{\rho}} \mathbf{s} + \eta \mathbf{h} \right) - \frac{1}{d} \mathcal{M}_{\frac{\tau_2}{\eta}f^*(\Omega^{1/2, \cdot})} \left(\Omega^{-1/2} \left(\nu\mathbf{v} + \kappa\Omega^{1/2}\mathbf{g} \right) \right) \end{aligned} \quad (\text{C.106})$$

and using the properties of Moreau envelopes, $\mathcal{M}_{\frac{\tau_1}{\kappa}g(\cdot, \mathbf{y})} \left(\frac{m}{\sqrt{\rho}} \mathbf{s} + \eta \mathbf{h} \right)$ is jointly convex in $(\kappa, \tau_1, m, \eta)$ as a composition of convex functions of those arguments. The same applies for $\mathcal{M}_{\frac{\tau_2}{\eta}f^*(\Omega^{1/2, \cdot})} \left(\Omega^{-1/2} \left(\nu\mathbf{v} + \kappa\Omega^{1/2}\mathbf{g} \right) \right)$, jointly convex in $(\tau_2, \eta, \nu, \kappa)$, and its opposite is jointly concave in those parameters. The remaining terms being linear in τ_1, τ_2, ν , we conclude that $\mathcal{E}_n(\tau_1, \tau_2, \kappa, \eta, \nu, m)$ is jointly concave in (ν, τ_2) and convex in τ_1 whatever the values of (κ, η, m) . Going back to equation (C.80), we can write

$$\begin{aligned} \mathcal{E}_n(\tau_1, \tau_2, \kappa, \eta, \nu, m) &= \max_{\boldsymbol{\mu}} \min_{\mathbf{z}, \mathbf{w}_\perp} \left[\frac{\kappa\tau_1}{2} - \frac{\eta\tau_2}{2} - \frac{\eta}{2\tau_2 d} \|\boldsymbol{\mu} + \kappa\mathbf{g}\|_2^2 + \frac{\kappa}{2\tau_1 d} \left\| \frac{m}{\sqrt{\rho}} \mathbf{s} + \eta \mathbf{h} - \mathbf{z} \right\|_2^2 \right. \\ &\quad \left. + \frac{1}{d} \left[g(\mathbf{z}, \mathbf{y}) + f \left(\frac{m\sqrt{d\rho}}{\|\mathbf{v}\|_2^2} \mathbf{v} + \mathbf{w}_\perp \right) + \boldsymbol{\mu}^\top \Sigma^{1/2} \left(\frac{m\sqrt{d\rho}}{\|\mathbf{v}\|_2^2} \mathbf{v} + \mathbf{w}_\perp \right) - \nu \mathbf{v}^\top \mathbf{w}_\perp \right] \right] \end{aligned} \quad (\text{C.107})$$

The squared term in m, η, \mathbf{z} can be written as

$$\frac{\kappa}{2\tau_1 d} \left\| \frac{m}{\sqrt{\rho}} \mathbf{s} + \eta \mathbf{h} - \mathbf{z} \right\|_2^2 = \tau_1 \frac{\kappa}{2d} \left\| \frac{m}{\tau_1 \sqrt{\rho}} \mathbf{s} + \frac{\eta}{\tau_1} \mathbf{h} - \frac{\mathbf{z}}{\tau_1} \right\|_2^2 \quad (\text{C.108})$$

which is the perspective function with parameter τ_1 of a function jointly convex in (\mathbf{z}, m, η) . Thus it is jointly convex in $(\tau_1, \mathbf{z}, m, \eta)$. Furthermore, the term $f \left(\frac{m\sqrt{d\rho}}{\|\mathbf{v}\|_2^2} \mathbf{v} + \mathbf{w}_\perp \right)$ is a composition of a convex function with a linear one, thus it is jointly convex in (m, \mathbf{w}_\perp) . The remaining terms in τ_1, η, m are linear. Since minimisation on convex sets preserves convexity, minimizing with respect to $\mathbf{z}, \mathbf{w}_\perp$ will lead to a jointly convex function in (τ_1, η, m) . Similarly, the term $-\frac{\eta}{2\tau_2 d} \|\boldsymbol{\mu} + \kappa\mathbf{g}\|_2^2$ is jointly concave in $\tau_2, \kappa, \boldsymbol{\mu}$, and maximizing over $\boldsymbol{\mu}$ will result in a jointly concave function in (τ_2, ν, κ) . We conclude that $\mathcal{E}_n(\tau_1, \tau_2, \kappa, \eta, \nu, m)$ is jointly convex in (τ_1, m, η) and jointly concave in (κ, ν, τ_2) . \square

The next lemma then characterizes the infinite dimensional limit of the scalar optimization problem (C.96), along with the consistency of its optimal value.

Lemma 11. (Asymptotics of \mathcal{E}_n) Recall the following quantities:

$$\mathcal{L}_g(\tau_1, \kappa, m, \eta) = \frac{1}{n} \mathbb{E} \left[\mathcal{M}_{\frac{\tau_1}{\kappa} g(\cdot, \mathbf{y})} \left(\frac{m}{\sqrt{\rho}} \mathbf{s} + \eta \mathbf{h} \right) \right] \text{ where } \mathbf{y} = f_0(\sqrt{\rho_p} \mathbf{s}), \mathbf{s} \sim \mathcal{N}(0, \mathbb{I}_n) \quad (\text{C.109})$$

$$\mathcal{L}_r(\tau_2, \eta, \nu, \kappa) = \frac{1}{d} \mathbb{E} \left[\mathcal{M}_{\frac{\eta}{\tau_2} r(\Omega^{-1/2} \cdot)} \left(\frac{\eta}{\tau_2} \left(\nu \Omega^{-1/2} \mathbf{v} + \kappa \mathbf{g} \right) \right) \right] \text{ where } \mathbf{v} = \Phi^\top \boldsymbol{\theta}_0 \quad (\text{C.110})$$

$$\chi = \frac{1}{d} \boldsymbol{\theta}_0^\top \Phi \Omega^{-1} \Phi^\top \boldsymbol{\theta}_0 \quad (\text{C.111})$$

$$\rho = \frac{1}{p} \boldsymbol{\theta}_0^\top \Psi \boldsymbol{\theta}_0 \quad (\text{C.112})$$

and the potential:

$$\mathcal{E}(\tau_1, \tau_2, \kappa, \eta, \nu, m) = \frac{\kappa \tau_1}{2} - \frac{\eta \tau_2}{2} + m \nu \sqrt{\gamma} - \frac{\tau_2}{2\eta} \frac{m^2}{\rho} - \frac{\eta}{2\tau_2} (\nu^2 \chi + \kappa^2) + \alpha \mathcal{L}_g(\tau_1, \kappa, m, \eta) + \mathcal{L}_r(\tau_2, \eta, \nu, \kappa) \quad (\text{C.113})$$

Then:

$$\max_{\kappa, \nu, \tau_2} \min_{m, \eta, \tau_1} \mathcal{E}_n(\tau_1, \tau_2, \kappa, \eta, \nu, m) \xrightarrow[n, p, d \rightarrow \infty]{P} \max_{\kappa, \nu, \tau_2} \min_{m, \eta, \tau_1} \mathcal{E}(\tau_1, \tau_2, \kappa, \eta, \nu, m) \quad (\text{C.114})$$

and $\mathcal{E}(\tau_1, \tau_2, \kappa, \eta, \nu, m)$ is continuously differentiable on its domain, jointly convex in (m, η, τ_1) and jointly concave in (κ, ν, τ_2) .

Proof of Lemma 11: The strong law of large numbers, see e.g. [68] gives $\frac{1}{d} \mathbf{g}^\top \mathbf{g} \xrightarrow[d \rightarrow \infty]{a.s.} 1$. Additionally, using assumption (A2) on the summability of $\boldsymbol{\theta}_0$ and (A3) on the boundedness of the spectrum of the covariance matrices, the quantity $\chi = \lim_{d \rightarrow \infty} \frac{1}{d} \boldsymbol{\theta}_0^\top \Phi \Omega^{-1} \Phi^\top \boldsymbol{\theta}_0$ exists and is finite. Since $\boldsymbol{\theta}_0 \notin \text{Ker}(\Phi^\top)$ and using the non-vanishing signal hypothesis, the quantity $\rho_{\mathbf{v}} = \lim_{d \rightarrow \infty} \frac{1}{d} \mathbf{v}^\top \mathbf{v}$ exists, is finite and strictly positive. Then $\kappa \mathbf{g}^\top (\Sigma^{1/2} - \Omega^{1/2}) \frac{m\sqrt{\gamma}}{\|\mathbf{v}\|_2^2} \mathbf{v}$ is a centered Gaussian random variable with variance verifying:

$$\begin{aligned} \text{Var} \left[\kappa \mathbf{g}^\top (\Sigma^{1/2} - \Omega^{1/2}) \frac{m\sqrt{\gamma}}{\|\mathbf{v}\|_2^2} \mathbf{v} \right] &\leq \kappa^2 \sigma_{max}^2 (\Sigma^{1/2} - \Omega^{1/2}) \frac{m^2 \gamma}{\|\mathbf{v}\|_2^2} \\ &= \kappa^2 \sigma_{max}^2 (\Sigma^{1/2} - \Omega^{1/2}) \frac{m^2 \gamma}{d \rho_{\mathbf{v}}} \end{aligned} \quad (\text{C.115})$$

Using lemma 8, κ and m are finitely bounded independently of the dimension d . $\gamma, \sigma_{max} (\Sigma^{1/2} - \Omega^{1/2})$ are finite. Thus there exists a finite constant C such that the standard deviation of $\kappa \mathbf{g}^\top (\Sigma^{1/2} - \Omega^{1/2}) \frac{m\sqrt{\gamma}}{\|\mathbf{v}\|_2^2} \mathbf{v}$ is smaller than \sqrt{C}/\sqrt{d} . Then, for any $\epsilon > 0$:

$$\begin{aligned} \mathbb{P} \left(\left| \kappa \mathbf{g}^\top (\Sigma^{1/2} - \Omega^{1/2}) \frac{m\sqrt{\gamma}}{\|\mathbf{v}\|_2^2} \mathbf{v} \right| \geq \epsilon \right) &\leq \mathbb{P} \left(|\mathcal{N}(0, 1)| \geq \epsilon \sqrt{d}/\sqrt{C} \right) \\ &\leq \frac{\sqrt{C}}{\epsilon \sqrt{d}} \frac{1}{\sqrt{2\pi}} \exp \left(-\frac{1}{2} \frac{\epsilon^2 d}{C} \right) \end{aligned} \quad (\text{C.116})$$

using the Gaussian tail. The Borel-Cantelli lemma and summability of this tail gives

$$\kappa \mathbf{g}^\top (\Sigma^{1/2} - \Omega^{1/2}) \frac{m\sqrt{\gamma}}{\|\mathbf{v}\|_2^2} \mathbf{v} \xrightarrow[d \rightarrow \infty]{a.s.} 0 \quad (\text{C.117})$$

Concentration of the Moreau envelopes of both f and g follows directly from lemma 5.

We thus have the pointwise convergence:

$$\mathcal{E}_n(\tau_1, \tau_2, \kappa, \eta, \nu, m) \xrightarrow[n, p, d \rightarrow \infty]{P} \mathcal{E}(\tau_1, \tau_2, \kappa, \eta, \nu, m) \quad (\text{C.118})$$

Since pointwise convergence preserves convexity, $\mathcal{E}(\tau_1, \tau_2, \kappa, \eta, \nu, m)$ is jointly convex in (m, η, τ_1) and jointly concave in (κ, ν, τ_2) .

Now recall the expression of \mathcal{E}

$$\mathcal{E}(\tau_1, \tau_2, \kappa, \eta, \nu, m) = \frac{\kappa\tau_1}{2} - \frac{\eta\tau_2}{2} + m\nu\sqrt{\gamma} - \frac{\tau_2}{2\eta} \frac{m^2}{\rho} - \frac{\eta}{2\tau_2}(\nu^2\chi + \kappa^2) + \alpha\mathcal{L}_g(\tau_1, \kappa, m, \eta) + \mathcal{L}_f(\tau_2, \eta, \nu, \kappa) \quad (\text{C.119})$$

The feasibility sets of κ, η, m are compact from Lemma 8 and the subsequent follow-up of the feasibility sets. Then, using Proposition 12.32 from [64], for fixed $(\tau_2, \kappa, \eta, \nu, m)$, we have:

$$\lim_{\tau_1 \rightarrow +\infty} \frac{1}{d} \mathcal{M}_{\frac{\tau_1}{\kappa} g(\cdot, \mathbf{y})} \left(\frac{m}{\sqrt{\rho}} \mathbf{s} + \eta \mathbf{h} \right) = \frac{1}{d} \inf_{\mathbf{z} \in \mathbb{R}^n} g(\mathbf{z}, \mathbf{y}) \quad (\text{C.120})$$

which is a finite quantity since $g(\cdot, \mathbf{y})$ is a proper, convex function verifying the scaling assumptions C.1. Then, since $\kappa > 0$, we have:

$$\lim_{\tau_1 \rightarrow +\infty} \mathcal{E}_n(\tau_1, \tau_2, \kappa, \eta, \nu, m) = +\infty \quad (\text{C.121})$$

Similarly, for fixed $(\tau_1, \kappa, \eta, \nu, m)$ and noting that composing f with the positive definite matrix $\Omega^{-1/2}$ does not change its convexity, or it being proper and lower semi-continuous, we get:

$$\lim_{\tau_2 \rightarrow +\infty} \frac{1}{d} \mathcal{M}_{\frac{\eta}{\tau_2} f(\Omega^{-1/2} \cdot)} \left(\frac{\eta}{\tau_2} \left(\nu \Omega^{-1/2} \mathbf{v} + \kappa \mathbf{g} \right) \right) = \frac{1}{d} f(0_d) \quad (\text{C.122})$$

which is also a bounded quantity from the scaling assumptions made on f . Since $\beta > 0$, we then have:

$$\lim_{\tau_2 \rightarrow +\infty} \mathcal{E}_n(\tau_1, \tau_2, \kappa, \eta, \nu, m) = -\infty \quad (\text{C.123})$$

Finally, the limit $\lim_{\nu \rightarrow +\infty} \mathcal{E}_n(\tau_1, \tau_2, \kappa, \eta, \nu, m)$ needs to be checked for both $+\infty$ and $-\infty$ since there is no restriction on the sign of ν . From the definition of the Moreau envelope, we can write:

$$\frac{1}{d} \mathcal{M}_{\frac{\eta}{\tau_2} f(\Omega^{-1/2} \cdot)} \left(\frac{\eta}{\tau_2} \left(\nu \Omega^{-1/2} \mathbf{v} + \kappa \mathbf{g} \right) \right) \leq \frac{1}{d} f(0_d) + \frac{\tau_2}{2\eta} \left\| \frac{\eta}{d\tau_2} \left(\nu \Omega^{-1/2} \mathbf{v} + \kappa \mathbf{g} \right) \right\|_2^2 \quad (\text{C.124})$$

Thus, for any fixed $(\tau_1, \tau_2, m, \kappa, \eta)$:

$$\mathcal{E}_n(\tau_1, \tau_2, \kappa, \eta, \nu, m) \leq \frac{\kappa\tau_1}{2} - \frac{\eta\tau_2}{2} + m\nu\sqrt{\gamma} - \kappa \mathbf{g}^\top \left(\Sigma^{1/2} - \Omega^{1/2} \right) \frac{m\sqrt{\gamma}}{\|\mathbf{v}\|_2} \mathbf{v} + \frac{1}{d} \mathcal{M}_{\frac{\tau_1}{\kappa} g(\cdot, \mathbf{y})} \left(\frac{m}{\sqrt{\rho}} \mathbf{s} + \eta \mathbf{h} \right) + \frac{1}{d} f(0_d) \quad (\text{C.125})$$

which immediately gives $\lim_{\nu \rightarrow -\infty} \mathcal{E}_n = -\infty$. Turning to the other limit, remembering that \mathcal{E}_n is continuously differentiable on its domain, we have:

$$\frac{\partial \mathcal{E}_n}{\partial \nu}(\tau_1, \tau_2, \kappa, \eta, \nu, m) = m\sqrt{\gamma} - \frac{1}{d} \mathbf{v}^\top \Omega^{-1/2} \text{prox}_{\frac{\eta}{\tau_2} f(\Omega^{-1/2} \cdot)} \left(\frac{\eta}{\tau_2} \left(\nu \Omega^{-1/2} \mathbf{v} + \kappa \mathbf{g} \right) \right) \quad (\text{C.126})$$

Thus $\lim_{\nu \rightarrow +\infty} \frac{\partial \mathcal{E}_n(\tau_1, \tau_2, \kappa, \eta, \nu, m)}{\partial \nu} \rightarrow -\infty$. Since \mathcal{E}_n is continuously differentiable in ν on $[0, +\infty[$, and from the short argument led above, we have shown

$$\lim_{|\nu| \rightarrow +\infty} \mathcal{E}_n(\tau_1, \tau_2, \kappa, \eta, \nu, m) = -\infty \quad (\text{C.127})$$

Using similar arguments as in the proof of Lemma 8, we can now reduce the feasibility set of τ_1, τ_2, ν to a compact one. Then, using the fact that convergence of convex functions on compact sets implies uniform convergence [69], we obtain

$$\max_{\kappa, \nu, \tau_2} \min_{m, \eta, \tau_1} \mathcal{E}_n(\tau_1, \tau_2, \kappa, \eta, \nu, m) \xrightarrow{P} \max_{\kappa, \nu, \tau_2} \min_{m, \eta, \tau_1} \mathcal{E}(\tau_1, \tau_2, \kappa, \eta, \nu, m) \quad (\text{C.128})$$

which is the desired result. \square

At this point, it is necessary to characterize the set of solutions of the asymptotic minimisation problem (2.8). We start with the explicit form of the optimality condition associated to any solution.

Lemma 12. (*Fixed point equations*) *The zero-gradient condition of the optimization problem (2.8) prescribes the following set of fixed point equations for any feasible solution:*

$$\partial_{\kappa} : \tau_1 = \frac{1}{d} \mathbb{E} \left[\mathbf{g}^\top \text{prox}_{\frac{\eta}{\tau_2} f(\Omega^{-1/2} \cdot)} \left(\frac{\eta}{\tau_2} (\nu \Omega^{-1/2} \mathbf{v} + \kappa \mathbf{g}) \right) \right] \quad (\text{C.129})$$

$$\partial_{\nu} : m\sqrt{\gamma} = \frac{1}{d} \mathbb{E} \left[\mathbf{v}^\top \Omega^{-1/2} \text{prox}_{\frac{\eta}{\tau_2} f(\Omega^{-1/2} \cdot)} \left(\frac{\eta}{\tau_2} (\nu \Omega^{-1/2} \mathbf{v} + \kappa \mathbf{g}) \right) \right] \quad (\text{C.130})$$

$$\partial_{\eta} : \tau_2 = \alpha \frac{\kappa}{\tau_1} \eta - \frac{\kappa \alpha}{\tau_1 n} \mathbb{E} \left[\mathbf{h}^\top \text{prox}_{\frac{\tau_1}{\kappa} g(\cdot, \mathbf{y})} \left(\frac{m}{\sqrt{\rho}} \mathbf{s} + \eta \mathbf{h} \right) \right] \quad (\text{C.131})$$

$$\begin{aligned} \partial_{\tau_2} : \frac{1}{2d} \frac{\tau_2}{\eta} \mathbb{E} \left[\left\| \frac{\eta}{\tau_2} (\nu \Omega^{-1/2} \mathbf{v} + \kappa \mathbf{g}) - \text{prox}_{\frac{\eta}{\tau_2} f(\Omega^{-1/2} \cdot)} \left(\frac{\eta}{\tau_2} (\nu \Omega^{-1/2} \mathbf{v} + \kappa \mathbf{g}) \right) \right\|_2^2 \right] = \\ \frac{\eta}{2\tau_2} (\nu^2 \chi + \kappa^2) - m\nu\sqrt{\gamma} - \kappa\tau_1 + \frac{\eta\tau_2}{2} + \frac{\tau_2}{2\eta} \frac{m^2}{\rho} \end{aligned} \quad (\text{C.132})$$

$$\partial_m : \nu\sqrt{\gamma} = \alpha \frac{\kappa}{n\tau_1} \mathbb{E} \left[\left(\frac{m}{\eta\rho} \mathbf{h} - \frac{\mathbf{s}}{\sqrt{\rho}} \right)^\top \text{prox}_{\frac{\tau_1}{\kappa} g(\cdot, \mathbf{y})} \left(\frac{m}{\sqrt{\rho}} \mathbf{s} + \eta \mathbf{h} \right) \right] \quad (\text{C.133})$$

$$\partial_{\tau_1} : \frac{\tau_1^2}{2} = \frac{1}{2} \alpha \frac{1}{n} \mathbb{E} \left[\left\| \frac{m}{\sqrt{\rho}} \mathbf{s} + \eta \mathbf{h} - \text{prox}_{\frac{\tau_1}{\kappa} g(\cdot, \mathbf{y})} \left(\frac{m}{\sqrt{\rho}} \mathbf{s} + \eta \mathbf{h} \right) \right\|_2^2 \right] \quad (\text{C.134})$$

Proof of Lemma 12: Using arguments similar to the ones in the proof of Lemma 5, Moreau envelopes and their derivatives verify the necessary conditions of the dominated convergence theorem. Additionally, uniform convergence of the sequence of derivatives can be verified in a straightforward manner as all involved functions are firmly non-expansive and integrated w.r.t. a Gaussian measure. We can therefore invert the limits and derivatives, and invert expectations and derivatives. We can now write explicitly the optimality condition for the scalar problem (C.113), using the expressions for derivatives of Moreau envelopes from Appendix C.2. Some algebra and replacing with prescriptions obtained from each partial derivative leads to the set of equations above. \square

Remark : Here we see that the potential function (C.113) can be further studied using the fixed point equations (12) and the relation (C.10). For any optimal $(\tau_1, \tau_2, \kappa, \eta, \nu, m)$, it holds that

$$\begin{aligned} \mathcal{E}(\tau_1, \tau_2, \kappa, \eta, \nu, m) \\ = \alpha \frac{1}{n} \mathbb{E} \left[g \left(\text{prox}_{\frac{\tau_1}{\kappa} g(\cdot, \mathbf{y})} \left(\frac{m}{\sqrt{\rho}} \mathbf{s} + \eta \mathbf{h} \right), \mathbf{y} \right) \right] + \frac{1}{d} \mathbb{E} \left[f \left(\Omega^{-1/2} \text{prox}_{\frac{\eta}{\tau_2} f(\Omega^{-1/2} \cdot)} \left(\frac{\eta}{\tau_2} (\nu \Omega^{-1/2} \mathbf{v} + \kappa \mathbf{g}) \right) \right) \right] \end{aligned} \quad (\text{C.135})$$

Finally, we give a strict-convexity and strict-concavity property of the asymptotic potential \mathcal{E} which will be helpful to prove Lemma 1.

Lemma 13. (*Strict convexity and strict concavity near minimisers*) *Consider the asymptotic potential function $\mathcal{E}(\tau_1, \tau_2, \kappa, \eta, \nu, m)$. Then for any fixed (η, m, τ_1) in their feasibility sets, the function*

$$\tau_2, \kappa, \nu \rightarrow \mathcal{E}(\tau_1, \tau_2, \kappa, \eta, \nu, m) \quad (\text{C.136})$$

is jointly strictly concave in (τ_2, κ, ν) .

Additionally, consider the set $\mathcal{S}_{\partial\nu, \tau_2}$ defined by:

$$\mathcal{S}_{\partial\nu, \tau_2} = \left\{ \tau_1, \tau_2, \kappa, \eta, \nu, m \mid m\sqrt{\gamma} = \frac{1}{d} \mathbb{E} \left[\mathbf{v}^T \Omega^{-1/2} \text{prox}_{\frac{\eta}{\tau_2} f(\Omega^{-1/2} \cdot)} \left(\frac{\eta}{\tau_2} \left(\nu \Omega^{-1/2} \mathbf{v} + \kappa \mathbf{g} \right) \right) \right] \right. \\ \left. \frac{1}{2d} \frac{1}{\eta} \mathbb{E} \left[\left\| \text{prox}_{\frac{\eta}{\tau_2} f(\Omega^{-1/2} \cdot)} \left(\frac{\eta}{\tau_2} \left(\nu \Omega^{-1/2} \mathbf{v} + \kappa \mathbf{g} \right) \right) \right\|_2^2 \right] = \frac{\eta}{2} + \frac{1}{2\eta} \frac{m^2}{\rho} \right\} \quad (\text{C.137})$$

then for any fixed τ_2, κ, ν in $\mathcal{S}_{\partial\nu, \tau_2}$, the function $(\eta, m, \tau_1) \rightarrow \mathcal{E}(\tau_1, \tau_2, \kappa, \eta, \nu, m)$ is jointly strictly convex in (η, m, τ_1) on $\mathcal{S}_{\partial\nu, \tau_2}$

Proof of Lemma 13: We will use the following first order characterization of strictly convex functions: f is strictly convex $\iff \langle \mathbf{x} - \mathbf{y} | \nabla f(\mathbf{x}) - \nabla f(\mathbf{y}) \rangle > 0 \forall \mathbf{x} \neq \mathbf{y} \in \text{dom}(f)$. To simplify notations, we will write, for any fixed (m, η, τ_1)

$$(\nabla_{\kappa, \nu, \tau_2} \mathcal{E}) = ((\partial_\kappa \mathcal{E}, \partial_\nu \mathcal{E}, \partial_{\tau_2} \mathcal{E})(\tau_1, \tau_2, \kappa, \eta, \nu, m))_i \quad (\text{C.138})$$

as the i -th component of the gradient of $\mathcal{E}(\tau_1, \tau_2, \kappa, \eta, \nu, m)$ with respect to (κ, ν, τ_2) for any fixed (m, η, τ_1) in the feasibility set. Then for any distinct triplets $(\kappa, \nu, \tau_2), (\tilde{\kappa}, \tilde{\nu}, \tilde{\tau}_2)$ and fixed (η, m, τ_1) in the feasibility set, determining the partial derivatives of \mathcal{E} in similar fashion as is implied in the proof of Lemma 12, we have:

$$\begin{aligned} & ((\kappa, \nu, \tau_2) - (\tilde{\kappa}, \tilde{\nu}, \tilde{\tau}_2))^\top (\nabla \mathcal{E}_{\kappa, \nu, \tau_2} - \nabla \mathcal{E}_{\tilde{\kappa}, \tilde{\nu}, \tilde{\tau}_2}) \\ &= (\kappa - \tilde{\kappa}) \alpha \frac{1}{2\tau_1} \frac{1}{n} \left(\mathbb{E} \left[\left\| \mathbf{r}_1 - \text{prox}_{\frac{\tau_1}{\kappa} g(\cdot, \mathbf{y})}(\mathbf{r}_1) \right\|_2^2 - \left\| \mathbf{r}_1 - \text{prox}_{\frac{\tau_1}{\tilde{\kappa}} g(\cdot, \mathbf{y})}(\mathbf{r}_1) \right\|_2^2 \right] \right) \\ &+ \left(\text{prox}_{\frac{\eta}{\tau_2} f(\Omega^{-1/2} \cdot)} \left(\frac{\eta}{\tau_2} \mathbf{r}_2 \right) - \text{prox}_{\frac{\eta}{\tilde{\tau}_2} f(\Omega^{-1/2} \cdot)} \left(\frac{\eta}{\tilde{\tau}_2} \tilde{\mathbf{r}}_2 \right) \right)^\top \left(\tilde{\mathbf{r}}_2 - \mathbf{r}_2 \right. \\ &\quad \left. + \frac{\tau_2 - \tilde{\tau}_2}{2\eta d} \left(\text{prox}_{\frac{\eta}{\tau_2} f(\Omega^{-1/2} \cdot)} \left(\frac{\eta}{\tau_2} \mathbf{r}_2 \right) + \text{prox}_{\frac{\eta}{\tilde{\tau}_2} f(\Omega^{-1/2} \cdot)} \left(\frac{\eta}{\tilde{\tau}_2} \tilde{\mathbf{r}}_2 \right) \right) \right) \\ &\leq (\kappa - \tilde{\kappa}) \alpha \frac{1}{2\tau_1} \frac{1}{n} \left(\mathbb{E} \left[\left\| \mathbf{r}_1 - \text{prox}_{\frac{\tau_1}{\kappa} g(\cdot, \mathbf{y})}(\mathbf{r}_1) \right\|_2^2 - \left\| \mathbf{r}_1 - \text{prox}_{\frac{\tau_1}{\tilde{\kappa}} g(\cdot, \mathbf{y})}(\mathbf{r}_1) \right\|_2^2 \right] \right) \\ &\quad + \left(\frac{\tau_2 + \tilde{\tau}_2}{2\eta d} \mathbb{E} \left[- \left\| \text{prox}_{\frac{\eta}{\tau_2} f(\Omega^{-1/2} \cdot)} \left(\frac{\eta}{\tau_2} \mathbf{r}_2 \right) - \text{prox}_{\frac{\eta}{\tilde{\tau}_2} f(\Omega^{-1/2} \cdot)} \left(\frac{\eta}{\tilde{\tau}_2} \tilde{\mathbf{r}}_2 \right) \right\|_2^2 \right] \right) \end{aligned} \quad (\text{C.139})$$

where the last line follows from the inequality in Lemma 4, and we defined the shorthands, $\mathbf{r}_1 = \frac{m}{\sqrt{\rho}} \mathbf{s} + \eta \mathbf{h}$, $\mathbf{r}_2 = \nu \Omega^{-1/2} \mathbf{v} + \kappa \mathbf{g}$, $\tilde{\mathbf{r}}_2 = \tilde{\nu} \Omega^{-1/2} \mathbf{v} + \tilde{\kappa} \mathbf{g}$. Using Lemma 3, the first term of the r.h.s of the last inequality is also negative as an increment of a nonincreasing function. Thus, both expectations are taken on negative functions. If those functions are not zero almost everywhere with respect to the Lebesgue measure, then the result will be strictly negative. Moreover, the functional taking each operator T to its resolvent $(\text{Id} + T)^{-1}$ is a bijection on the set of non-trivial, maximally monotone operators, see e.g. [64] Proposition 23.21 and the subsequent discussion. The subdifferential of a proper, closed, convex function being maximally monotone, for two different parameters the corresponding proximal operator cannot be equal almost everywhere. The previously studied increment $((\kappa, \nu, \tau_2) - (\tilde{\kappa}, \tilde{\nu}, \tilde{\tau}_2))^\top (\nabla \mathcal{E}_{\kappa, \nu, \tau_2} - \nabla \mathcal{E}_{\tilde{\kappa}, \tilde{\nu}, \tilde{\tau}_2})$ is therefore strictly negative, giving the desired strict concavity in (κ, ν, τ_2) . Restricting ourselves to the set $\mathcal{S}_{\partial\nu, \tau_2}$, the increment in (m, η, τ_1) can be written similarly. Note that $\text{Id} - \text{prox}$ will appear in the expressions instead of prox . The appropriate terms can then be brought to the form of the inequality from Lemma 4 using Moreau's decomposition. Using the definitions of the set $\mathcal{S}_{\partial\nu, \tau_2}$ and the increments from Lemma 3, a similar argument as the previous one can be carried out. The lemma is proved. \square

What is now left to do is link the properties of the scalar optimization problem (2.8) to the original learning problem (1.3) using the tight inequalities from Theorem 4.

Remark: in the case $\theta_0 \in \text{Ker}(\Phi^T)$, the cost function \mathcal{E}_n^0 will uniformly converge to the following potential:

$$-\frac{\eta\tau_2}{2} + \frac{\kappa\tau_1}{2} - \frac{\eta}{2\tau_2}\kappa^2 + \frac{\alpha}{n}\mathbb{E}\left[\mathcal{M}_{\frac{\tau_1}{\kappa}g(\cdot, \mathbf{y})}(\eta\mathbf{h})\right] + \frac{1}{d}\mathbb{E}\left[\mathcal{M}_{\frac{\eta}{\tau_2}f(\Omega^{-1/2}\cdot)}\left(\frac{\eta}{\tau_2}\kappa\mathbf{g}\right)\right] \quad (\text{C.140})$$

As we will see in the next section, this will lead to estimators solely based on noise.

C.5 Back to the original problem : proof of Theorem 1 and 2

We begin this part by considering that the "necessary assumptions for exponential rates" from the set of assumptions C.1 are verified. In the end we will discuss how relaxing these assumptions modifies the convergence speed. We closely follow the detailed probabilistic pipeline introduced in [70] and further developed in [28]. The main difference resides in concentration properties of generic Moreau envelopes depending on the regularity of the target function instead of specific instances such as the LASSO. Our analysis is somewhat less precise in the optimization of constants, but the main goal is to broadly characterize the dimension dependencies of the asymptotic results. Since the dimensions n, p, d are linked by multiplicative constants, we can express the rates with any of the three. Recall the original reformulation of the problem defining the student.

$$\max_{\lambda} \min_{\mathbf{w}, \mathbf{z}} g(\mathbf{z}, \mathbf{y}) + f(\mathbf{w}) + \lambda^\top \left(\frac{1}{\sqrt{d}} \left(A\Psi^{-1/2}\Phi + B \left(\Omega - \Phi^\top\Psi^{-1}\Phi \right)^{1/2} \right) \mathbf{w} - \mathbf{z} \right) \quad (\text{C.141})$$

Introducing the variable $\tilde{\mathbf{w}} = \Omega^{1/2}\mathbf{w}$ it can be equivalently written, since Ω is almost surely invertible and the problem is convex concave with a closed convex feasibility set on $\tilde{\mathbf{w}}, \mathbf{z}$.

$$\min_{\tilde{\mathbf{w}}, \mathbf{z}} \max_{\lambda} g(\mathbf{z}, \mathbf{y}) + f(\Omega^{-1/2}\tilde{\mathbf{w}}) + \lambda^\top \left(\frac{1}{\sqrt{d}} \left(A\Psi^{-1/2}\Phi + B \left(\Omega - \Phi^\top\Psi^{-1}\Phi \right)^{1/2} \right) \Omega^{-1/2}\tilde{\mathbf{w}} - \mathbf{z} \right) \quad (\text{C.142})$$

Recall the equivalent scalar auxiliary problem at finite dimension \mathcal{E}_n and its asymptotic counterpart \mathcal{E} both defined on the same variables as the original problem $\tilde{\mathbf{w}}, \mathbf{z}$.

$$\mathcal{E}(\tau_1, \tau_2, \kappa, \eta, \nu, m) = \frac{\kappa\tau_1}{2} - \frac{\eta\tau_2}{2} + m\nu\sqrt{\gamma} - \frac{\tau_2}{2\eta} \frac{m^2}{\rho} - \frac{\eta}{2\tau_2}(\nu^2\chi + \kappa^2) + \alpha\mathcal{L}_g(\tau_1, \kappa, m, \eta) + \mathcal{L}_f(\tau_2, \eta, \nu, \kappa) \quad (\text{C.143})$$

$$\begin{aligned} \mathcal{E}_n(\tau_1, \tau_2, \kappa, \eta, \nu, m) &= \frac{\kappa\tau_1}{2} - \frac{\eta\tau_2}{2} + m\nu\sqrt{\gamma} - \frac{\tau_2}{2\eta} \frac{m^2}{\rho} - \frac{\eta}{2\tau_2 d} (\nu\mathbf{v} + \kappa\Omega^{1/2}\mathbf{g})^\top \Omega^{-1} (\nu\mathbf{v} + \kappa\Omega^{1/2}\mathbf{g}) \\ &- \kappa\mathbf{g}^\top \left(\Sigma^{1/2} - \Omega^{1/2} \right) \frac{m\sqrt{\gamma}}{\|\mathbf{v}\|_2^2} \mathbf{v} + \frac{1}{d} \mathcal{M}_{\frac{\tau_1}{\kappa}g(\cdot, \mathbf{y})} \left(\frac{m}{\sqrt{\rho}} \mathbf{s} + \eta\mathbf{h} \right) + \frac{1}{d} \mathcal{M}_{\frac{\eta}{\tau_2}f(\Omega^{-1/2}\cdot)} \left(\frac{\eta}{\tau_2} (\nu\Omega^{-1/2}\mathbf{v} + \kappa\mathbf{g}) \right) \end{aligned} \quad (\text{C.144})$$

Recall the variables:

$$\tilde{\mathbf{w}}^* = \text{prox}_{\frac{\eta}{\tau_2}f(\Omega^{-1/2}\cdot)} \left(\frac{\eta}{\tau_2} (\nu\mathbf{t} + \kappa\mathbf{g}) \right), \quad \mathbf{z}^* = \text{prox}_{\frac{\tau_1}{\kappa}g(\cdot, \mathbf{y})} \left(\frac{m}{\sqrt{\rho}} \mathbf{s} + \eta\mathbf{h} \right) \quad (\text{C.145})$$

Denote $(\tau_1^*, \tau_2^*, \kappa^*, \eta^*, \nu^*, m^*)$ the unique solution to the optimization problem (2.8) and \mathcal{E}^* the corresponding optimal cost. Similarly, denote $(\tau_{1,n}^*, \tau_{2,n}^*, \kappa_n^*, \eta_n^*, \nu_n^*, m_n^*)$ any solution to the optimization problem on \mathcal{E}_n and \mathcal{E}_n^* the corresponding optimal value. Finally, we write $E_n(\tilde{\mathbf{w}}, \mathbf{z})$ the cost function of the optimization problem on $\tilde{\mathbf{w}}, \mathbf{z}$ defined by \mathcal{E}_n^* for any optimal solution $(\tau_{1,n}^*, \tau_{2,n}^*, \kappa_n^*, \eta_n^*, \nu_n^*, m_n^*)$, such that:

$$\mathcal{E}_n^* = \min_{\tilde{\mathbf{w}}, \mathbf{z}} E_n(\tilde{\mathbf{w}}, \mathbf{z}) \quad (\text{C.146})$$

By the definition of Moreau envelopes, we have that $E_n(\tilde{\mathbf{w}}, \mathbf{z})$ is $\frac{\kappa}{2d\tau_1}$ strongly convex in \mathbf{z} and $\frac{\tau_2}{2d\eta}$ strongly convex in $\tilde{\mathbf{w}}$. The following lemma ensures that these strong convexity constants are non-zero for any finite n .

Lemma 14. Consider the finite size scalar optimization problem

$$\max_{\kappa, \nu, \tau_2} \min_{m, \eta, \tau_1} \mathcal{E}_n(\tau_1, \tau_2, \kappa, \eta, \nu, m) \quad (\text{C.147})$$

where the feasibility set of $(\tau_1, \tau_2, \kappa, \eta, \nu, m)$ is compact and $\tau_1 > 0, \tau_2 > 0$. Then any optimal values κ^*, τ_2^* verify:

$$\kappa^* \neq 0 \quad \tau_2^* \rightarrow 0 \quad (\text{C.148})$$

Proof of Lemma 14: from the analysis carried out in the proof of Lemma 11, the feasibility set of the optimization problem is compact. Suppose $\kappa^* = 0$. Then the value of m minimizing the cost function is $-\infty$, which contradicts the compactness of the feasibility set. A similar argument holds for τ_2 . \square

The next lemma characterizes the speed of convergence of the optimal value of the finite dimensional scalar optimization problem to its asymptotic counterpart, which has a unique minimizer.

Lemma 15. For any $\epsilon > 0$, there exist constants C, c, γ such that:

$$\mathbb{P}(|\mathcal{E}_n^* - \mathcal{E}^*| \geq \epsilon) \leq \frac{C}{\epsilon} \exp^{-c n \epsilon^2} \quad (\text{C.149})$$

which is equivalent to

$$\mathbb{P}\left(\left|\min_{\tilde{\mathbf{w}}, \mathbf{z}} E_n(\tilde{\mathbf{w}}, \mathbf{z}) - \mathcal{E}^*\right| \geq \epsilon\right) \leq \frac{C}{\epsilon} \exp^{-c n \epsilon^2} \quad (\text{C.150})$$

Proof of Lemma 15: for any fixed $(\tau_1, \tau_2, \kappa, \nu, \eta, m)$, we can determine the rates of convergence of all the random quantities in \mathcal{E}_n . The linear terms involving $\frac{1}{d} \mathbf{g}^T \mathbf{v}$ are sub-Gaussian with sub-Gaussian norm bounded by C/d for some constant $C > 0$. Thus we can find constants, $C, c > 0$ such that, for any $\epsilon > 0$:

$$\mathbb{P}\left(\left|\frac{1}{d} \mathbf{g}^T \mathbf{v}\right| \geq \epsilon\right) \leq C e^{-c n \epsilon^2} \quad (\text{C.151})$$

The term involving $\mathbf{v}^T \Omega \mathbf{v}$ is deterministic in this setting. We will see in section C.6 how a random $\boldsymbol{\theta}_0$ affects the convergence rates. The term involving $\frac{1}{d} \mathbf{g}^T \mathbf{g}$ is a weighted sum of sub-exponential random variables, the tail of which can be determined using Bernstein's inequality, see e.g. [71] Corollary 2.8.3, which gives a sub-Gaussian tail for small deviations and a sub-exponential tail for large deviations. Parametrizing the deviation ϵ with a scalar variable c' , we thus get the following bound: for any $\epsilon > 0$, there exists constants $C, c, c' > 0$ such that:

$$\mathbb{P}\left(\left|\frac{1}{d} \mathbf{g}^T \mathbf{g} - 1\right| \geq c' \epsilon\right) \leq C e^{-c n \epsilon^2} \quad (\text{C.152})$$

Since, in this case, we assume that the eigenvalues of the covariance matrices are bounded with probability one, multiplications by these matrices do not change these two previous rates. The remaining convergence rates that need to be determined are those of the Moreau envelopes. By assumption, the function g is separable, and pseudo-Lipschitz of order two. Moreover, the argument $\frac{m}{\sqrt{\rho}} \mathbf{s} + \eta \mathbf{h}$ is an i.i.d. Gaussian random vector with finite variance. The Moreau envelope $\frac{1}{d} \mathcal{M}_{\frac{\eta}{\tau_2} r(\Omega^{-1/2})} \left(\frac{\eta}{\tau_2} (\nu \Omega^{-1/2} \mathbf{v} + \kappa \mathbf{g})\right)$ is therefore a sum of pseudo-Lipschitz functions of order 2 of scalar Gaussian random variables. Using the concentration Lemma 7, we can find constants $C, c, c' > 0$ such that, for any $\epsilon > 0$, the following holds:

$$\mathbb{P}\left(\left|\alpha \frac{1}{n} \mathcal{M}_{\frac{\tau_1}{\kappa} g(\cdot, \mathbf{y})} \left(\frac{m}{\sqrt{\rho}} \mathbf{s} + \eta \mathbf{h}\right) - \mathbb{E} \left[\alpha \frac{1}{n} \mathcal{M}_{\frac{\tau_1}{\kappa} g(\cdot, \mathbf{y})} \left(\frac{m}{\sqrt{\rho}} \mathbf{s} + \eta \mathbf{h}\right) \right]\right| \geq c' \epsilon\right) \leq C e^{-c n \epsilon^2} \quad (\text{C.153})$$

For the second Moreau envelope, the argument $\frac{\eta}{\tau_2} (\nu \Omega^{-1/2} \mathbf{v} + \kappa \mathbf{g})$ is not separable. If the regularization is a square, it is straightforward to see that concentration will reduce to that of the terms $\frac{1}{d} \mathbf{g}^T \mathbf{v}$ and $\frac{1}{d} \mathbf{g}^T \mathbf{g}$. If the regularization is a Lipschitz function, then the Moreau envelope is also Lipschitz from Lemma 2. Furthermore,

since the eigenvalues of the covariance matrix Ω are bounded with probability one, the composition with the deterministic term $\nu\Omega^{1/2}\mathbf{v}$ does not change the Lipschitz property. Gaussian concentration of Lipschitz functions then gives an exponential decay independent of the magnitude of the deviation. Taking the loosest bound, which is the one obtained with the square penalty, we obtain that, for any $\epsilon > 0$, there exist constants $C, c, c' > 0$ such that the event

$$\left\{ \left| \frac{1}{d} \mathcal{M}_{\frac{\eta}{\tau_2} r(\Omega^{-1/2})} \left(\frac{\eta}{\tau_2} \left(\nu\Omega^{-1/2}\mathbf{v} + \kappa\mathbf{g} \right) \right) - \mathbb{E} \left[\frac{1}{d} \mathcal{M}_{\frac{\eta}{\tau_2} r(\Omega^{-1/2})} \left(\frac{\eta}{\tau_2} \left(\nu\Omega^{-1/2}\mathbf{v} + \kappa\mathbf{g} \right) \right) \right] \right| \geq c'\epsilon \right\} \quad (\text{C.154})$$

has probability at most $Ce^{-cn\epsilon^2}$. Combining these bounds gives the exponential rate for the convergence of \mathcal{E}_n to \mathcal{E} for any fixed $(\tau_1, \tau_2, \kappa, \nu, \eta, m)$. An ϵ -net argument can then be used to obtain the bound on the minmax values. \square

The next lemma shows that the function E_n evaluated at $\tilde{\mathbf{w}}^*, \mathbf{z}^*$ is close to the optimal value \mathcal{E}^* . Intuitively, this tell should be close to $\tilde{\mathbf{w}}^*, \mathbf{z}^*$. We will make this exact using the strong convexity of E_n .

Lemma 16. *For any $\epsilon > 0$, there exist constants C, c, c' such that:*

$$\mathbb{P} \left(|E_n(\tilde{\mathbf{w}}^*, \mathbf{z}^*) - \mathcal{E}^*| \geq c'\epsilon \right) \leq Ce^{-cn\epsilon^2} \quad (\text{C.155})$$

Proof of Lemma 16: in a similar fashion as was done in the proof of the fixed point equations given in Lemma 12, the optimality conditions corresponding to the finite size problem can be written and used to simplify the optimal value. Doing so, it can be shown, as done for \mathcal{E}^* in the remark following the proof of Lemma 12, that the function E_n is also equal to the original cost:

$$E_n(\tilde{\mathbf{w}}, \mathbf{z}) = \alpha \frac{1}{n} g(\mathbf{z}) + \frac{1}{d} r(\Omega^{-1/2}\tilde{\mathbf{w}}) \quad (\text{C.156})$$

we then have:

$$E_n(\tilde{\mathbf{w}}^*, \mathbf{z}^*) = \alpha \frac{1}{n} g \left(\text{prox}_{\frac{\tau_1^*}{\kappa^*} g(\cdot, \mathbf{y})} \left(\frac{m^*}{\sqrt{\rho}} \mathbf{s} + \eta^* \mathbf{h} \right), \mathbf{y} \right) + \frac{1}{d} f \left(\Omega^{-1/2} \text{prox}_{\frac{\eta^*}{\tau_2^*} f(\Omega^{-1/2})} \left(\frac{\eta^*}{\tau_2^*} \left(\nu^* \Omega^{-1/2} \mathbf{v} + \kappa^* \mathbf{g} \right) \right) \right) \quad (\text{C.157})$$

Note that this is not enough to finish the proof. It is tempting to claim that the optimal estimators are then given by forms similar to $\tilde{\mathbf{w}}^*, \mathbf{z}^*$ where the parameters $(\tau_1^*, \tau_2^*, \kappa^*, \eta^*, \nu^*, m^*)$ are replaced with any optimal solution $(\tau_{1,n}^*, \tau_{2,n}^*, \kappa_n^*, \eta_n^*, \nu_n^*, m_n^*)$. However the solution $(\tau_{1,n}^*, \tau_{2,n}^*, \kappa_n^*, \eta_n^*, \nu_n^*, m_n^*)$ is not unique. The point will be to use the fact that the asymptotic solution $(\tau_1^*, \tau_2^*, \kappa^*, \eta^*, \nu^*, m^*)$ is, this time, unique. Recall that, in the remark following the proof of Lemma 12, we showed that :

$$\mathcal{E}^* = \alpha \frac{1}{n} \mathbb{E} \left[g \left(\text{prox}_{\frac{\tau_1^*}{\kappa^*} g(\cdot, \mathbf{y})} \left(\frac{m^*}{\sqrt{\rho}} \mathbf{s} + \eta^* \mathbf{h} \right), \mathbf{y} \right) + \frac{1}{d} \mathbb{E} \left[f \left(\Omega^{-1/2} \text{prox}_{\frac{\eta^*}{\tau_2^*} f(\Omega^{-1/2})} \left(\frac{\eta^*}{\tau_2^*} \left(\nu^* \Omega^{-1/2} \mathbf{v} + \kappa^* \mathbf{g} \right) \right) \right) \right] \right] \quad (\text{C.158})$$

to prove the lemma, we therefore only need to prove the rates of convergence of the two terms in eq.(C.157). Since g is separable, so is its proximal operator. The proximal operator is a Lipschitz function, and the composition of a pseudo-Lipschitz function of order 2 with a Lipschitz function gives a pseudo-Lipschitz function of order 2. Therefore, this term concentrates like a scaled separable, pseudo-Lipschitz of order 2 function of a Gaussian random vector, which is given by Lemma 7. Regarding the second term, it is shown in [70, 28] that, for covariance matrices with bounded spectrum, and provided the function f is either Lipschitz or a square, the variable $\Omega^{-1/2} \text{prox}_{\frac{\eta^*}{\tau_2^*} f(\Omega^{-1/2})} \left(\frac{\eta^*}{\tau_2^*} \left(\nu^* \Omega^{-1/2} \mathbf{v} + \kappa^* \mathbf{g} \right) \right)$ is sub-gamma. Sub-gamma random variables provide the same tails as Bernstein's inequality : sub-gaussian for small deviation, sub-exponential for large deviation. This is enough to conclude the proof of Lemma 16. \square

We now use the strongly-convex form of E_n to show that minimizers of E_n become asymptotically close to $\tilde{\mathbf{w}}^*, \mathbf{z}^*$.

Lemma 17. For any $\epsilon > 0$, there exists constants $c', c, C > 0$ such that the event

$$\exists (\tilde{\mathbf{w}}, \mathbf{z}) \in \mathbb{R}^{n+d}, \frac{1}{d} \min\left(\frac{\kappa_n^*}{2\tau_{1,n}^*}, \frac{\tau_{2,n}^*}{2\eta_n^*}\right) \|(\tilde{\mathbf{w}}, \mathbf{z}) - (\tilde{\mathbf{w}}^*, \mathbf{z}^*)\|_2^2 > \epsilon \text{ and } \min_{\tilde{\mathbf{w}}, \mathbf{z}} E_n(\tilde{\mathbf{w}}, \mathbf{z}) \leq E_n(\tilde{\mathbf{w}}^*, \mathbf{z}^*) + c'\epsilon \quad (\text{C.159})$$

has probability at most $\frac{C}{\epsilon} e^{-c\epsilon^2}$.

This lemma can be proven using the same arguments as in [70] Appendix B, Theorem B.1. Note that we are normalizing the norm of a vector of size $(n + d)$ with d . This is not a problem, since n, d are linked by a finite multiplicative constant. Intuitively, if two values of a strongly convex function are arbitrarily close, then the corresponding points are arbitrarily close. This shows that any solution outside the ball centered around $\tilde{\mathbf{w}}^*, \mathbf{z}^*$ is sub-optimal. Now define the set:

$$D_{\tilde{\mathbf{w}}, \mathbf{z}, \epsilon} = \left\{ \tilde{\mathbf{w}} \in \mathbb{R}^d, \mathbf{z} \in \mathbb{R}^n : \left| \phi_1\left(\frac{\tilde{\mathbf{w}}}{\sqrt{d}}\right) - \mathbb{E} \left[\phi_1\left(\frac{\tilde{\mathbf{w}}^*}{\sqrt{d}}\right) \right] \right| > \epsilon, \left| \phi_2\left(\frac{\mathbf{z}}{\sqrt{n}}\right) - \mathbb{E} \left[\phi_2\left(\frac{\mathbf{z}^*}{\sqrt{n}}\right) \right] \right| > \epsilon \right\} \quad (\text{C.160})$$

where ϕ_1 is either a square or a Lipschitz function, and ϕ_2 is a separable, pseudo-Lipschitz function of order 2. Using the same arguments as in the proof of Lemma 16, we have that, for any $\epsilon > 0$, there exist constants $C, c > 0$ such that:

$$\mathbb{P}((\tilde{\mathbf{w}}^*, \mathbf{z}^*) \in D_{\tilde{\mathbf{w}}, \mathbf{z}, \epsilon}) \geq 1 - C e^{-c\epsilon^2} \quad (\text{C.161})$$

Also, any solution outside of the ball centered on $\tilde{\mathbf{w}}^*, \mathbf{z}^*$ with radius $c'\epsilon$, with $c' > 0$ a given constant, is not in $D_{\tilde{\mathbf{w}}, \mathbf{z}, \epsilon}$. The proof of Theorem 2 then follows from the exact same argument as the proof of Lemma B.3 from [28]. Theorem 1 is a consequence of Theorem 2.

If the restriction on f, g, ϕ_1, ϕ_2 are relaxed to any pseudo-Lipschitz functions of finite orders, the exponential rates involving them are lost and become linear following Lemma 5.

C.6 Relaxing assumption (B1)

The entirety of the previous proof has been done with a deterministic vector θ_0 . Now, if θ_0 is assumed to be a random vector independent of all other quantities, as prescribed in the set of assumptions C.1, we can "freeze" the variable θ_0 by conditioning on it. The whole proof can then be understood as studying the value of the cost conditioned on the value of θ_0 . This is sometimes referred to as the "freezing lemma", see e.g. [68] Example 4.1.7. Note that, in the Gaussian case, correlations between the teacher and student are expressed through the covariance matrices, thus leaving the possibility to parametrise the teacher with a vector θ_0 indeed independent of all the rest. To lift the conditioning in the end, one only needs to average out on the distribution of θ_0 , the summability conditions of which are prescribed in the set of assumptions C.1. Thus, random teacher vectors can be treated simply by taking an additional expectation in the expressions of Theorem 2, provided θ_0 is independent of the matrices A, B and the randomness in f_0 .

As mentioned at the end of the previous section, the finite size rates will be determined by the assumptions made on the teacher vector and decay of the eigenvalues of the covariance matrices. We do not investigate in detail the limiting assumptions under which exponential rates still hold regarding the randomness of the teacher or tails of the eigenvalue distributions of covariance matrices.

C.7 The 'vanilla' teacher-student scenario

In this section, we give the explicit forms of the fixed points equations and optimal asymptotic estimators in the case where the teacher and the student are sampled from the same distribution, i.e. $\Omega = \Phi = \Psi = \Sigma$ where Σ is a positive definite matrix with sub-Gaussian eigenvalue decay. This setup was rigorously studied in [28] for

the LASSO and heuristically in [38] for the ridge regularized logistic regression. In this case, the fixed point equations become

$$\tau_1 = \frac{1}{d} \mathbb{E} \left[\mathbf{g}^\top \text{prox}_{\frac{\eta}{\tau_2} f(\Sigma^{-1/2, \cdot})} \left(\frac{\eta}{\tau_2} (\nu \Sigma^{1/2} \boldsymbol{\theta}_0 + \kappa \mathbf{g}) \right) \right] \quad (\text{C.162})$$

$$m\sqrt{\gamma} = \frac{1}{d} \mathbb{E} \left[\mathbf{v}^\top \Sigma^{-1/2} \text{prox}_{\frac{\eta}{\tau_2} f(\Sigma^{-1/2, \cdot})} \left(\frac{\eta}{\tau_2} (\nu \Sigma^{1/2} \boldsymbol{\theta}_0 + \kappa \mathbf{g}) \right) \right] \quad (\text{C.163})$$

$$\tau_2 = \alpha \frac{\kappa}{\tau_1} \eta - \frac{\kappa \alpha}{\tau_1 n} \mathbb{E} \left[\mathbf{h}^\top \text{prox}_{\frac{\tau_1}{\kappa} g(\cdot, \mathbf{y})} \left(\frac{m}{\sqrt{\rho}} \mathbf{s} + \eta \mathbf{h} \right) \right] \quad (\text{C.164})$$

$$\eta^2 + \frac{m^2}{\rho} = \frac{1}{d} \mathbb{E} \left[\left\| \text{prox}_{\frac{\eta}{\tau_2} f(\Sigma^{-1/2, \cdot})} \left(\frac{\eta}{\tau_2} (\nu \Sigma^{1/2} \boldsymbol{\theta}_0 + \kappa \mathbf{g}) \right) \right\|_2^2 \right] \quad (\text{C.165})$$

$$\nu\sqrt{\gamma} = \alpha \frac{\kappa}{n\tau_1} \mathbb{E} \left[\left(\frac{m}{\eta\rho} \mathbf{h} - \frac{\mathbf{s}}{\sqrt{\rho}} \right)^\top \text{prox}_{\frac{\tau_1}{\kappa} g(\cdot, \mathbf{y})} \left(\frac{m}{\sqrt{\rho}} \mathbf{s} + \eta \mathbf{h} \right) \right] \quad (\text{C.166})$$

$$\tau_1^2 = \frac{\alpha}{n} \mathbb{E} \left[\left\| \frac{m}{\sqrt{\rho}} \mathbf{s} + \eta \mathbf{h} - \text{prox}_{\frac{\tau_1}{\kappa} g(\cdot, \mathbf{y})} \left(\frac{m}{\sqrt{\rho}} \mathbf{s} + \eta \mathbf{h} \right) \right\|_2^2 \right] \quad (\text{C.167})$$

and the asymptotic optimal estimators read:

$$\mathbf{w}^* = \Sigma^{-1/2} \text{prox}_{\frac{\eta^*}{\tau_2^*} f(\Sigma^{-1/2, \cdot})} \left(\frac{\eta^*}{\tau_2^*} (\nu^* \Sigma^{1/2} \boldsymbol{\theta}_0 + \kappa^* \mathbf{g}) \right), \quad \mathbf{z}^* = \text{prox}_{\frac{\tau_1^*}{\kappa^*} g(\cdot, \mathbf{y})} \left(\frac{m^*}{\sqrt{\rho}} \mathbf{s} + \eta^* \mathbf{h} \right) \quad (\text{C.168})$$

D Equivalence replica-Gordon

In this Appendix, we show that the rigorous result of Theorem 2 can be used to prove the replica prediction in the case of a separable loss, a ridge penalty. For simplicity, we restrict ourselves to the case of random teacher weights with $\boldsymbol{\theta}_0 \sim \mathcal{N}(0, \mathbf{I}_p)$. We provide an exact analytical matching between the replica prediction and the one obtained with Gordon's theorem. We start by an explicit derivation of the form presented in Corollary 1 from the main result (2.7).

D.1 Solution for separable loss and ridge regularization

Replacing r with a ridge penalty, we can go back to step (C.85) of the main proof and finish the calculation without inverting the matrix Ω . The assumption on the invertibility of Ω can thus be dropped in the case of ℓ_2 regularization. Letting $G = \left(\frac{\tau_2}{\eta}\Omega + \lambda_2\mathbf{I}_d\right)^{-1}$, we get

$$\begin{aligned} \mathcal{E}(\tau_1, \tau_2, \kappa, \eta, \nu, m) &= \frac{\kappa\tau_1}{2} - \frac{\eta\tau_2}{2} + m\nu\sqrt{\gamma} - \frac{\tau_2}{2\eta} \frac{m^2}{\rho} + \alpha \frac{1}{n} \mathbb{E} \left[\mathcal{M}_{\frac{\tau_1}{\kappa}g(\cdot, y)} \left(\frac{m}{\sqrt{\rho}}\mathbf{s} + \eta\mathbf{h} \right) \right] \\ &\quad - \frac{1}{2d} \nu^2 \boldsymbol{\theta}_0^\top \Phi G \Phi^\top \boldsymbol{\theta}_0 - \frac{1}{2d} \kappa^2 \text{Tr} \left(\Omega^{1/2} G \Omega^{1/2} \right) \end{aligned} \quad (\text{D.1})$$

using Lemma 5 with a separable function, the expectation over the Moreau envelope converges to:

$$\frac{1}{n} \mathbb{E} \left[\mathcal{M}_{\frac{\tau_1}{\kappa}g(\cdot, y)} \left(\frac{m}{\sqrt{\rho}}\mathbf{s} + \eta\mathbf{h} \right) \right] = \mathbb{E} \left[\mathcal{M}_{\frac{\tau_1}{\kappa}g(\cdot, y)} \left(\frac{m}{\sqrt{\rho}}s + \eta h \right) \right] \quad (\text{D.2})$$

where s and h are standard normal random variables and $y = f_0(\sqrt{\rho}s)$. The corresponding optimality conditions then reads:

$$\frac{\partial}{\partial \kappa} : \frac{\tau_1}{2} + \frac{1}{2\tau_1} \alpha \mathbb{E} \left[\left(\frac{m}{\sqrt{\rho}}s + \eta h - \text{prox}_{\frac{\tau_1}{\kappa}g(\cdot, y)} \left(\frac{m}{\sqrt{\rho}}s + \eta h \right) \right)^2 \right] - \kappa \frac{1}{d} \text{Tr} \left(\Omega^{1/2} G \Omega^{1/2} \right) = 0 \quad (\text{D.3})$$

$$\frac{\partial}{\partial \nu} : m\sqrt{\gamma} - \frac{1}{d} \nu \boldsymbol{\theta}_0^\top \Phi^\top G \Phi^\top \boldsymbol{\theta}_0 = 0 \quad (\text{D.4})$$

$$\frac{\partial}{\partial \tau_2} : -\frac{\eta}{2} - \frac{m^2}{2\rho\eta} + \frac{1}{2} \frac{\nu^2}{\eta} \left(\Omega^{1/2} \Phi^\top \boldsymbol{\theta}_0 \right)^\top G^2 \Omega^{1/2} \Phi^\top \boldsymbol{\theta}_0 + \frac{\kappa^2}{2\eta} \text{Tr} \left(G^2 \Omega^2 \right) = 0 \quad (\text{D.5})$$

$$\frac{\partial}{\partial m} : \nu\sqrt{\gamma} - \frac{\tau_2}{\rho\eta} m + \alpha \mathbb{E} \left[\frac{\kappa}{\tau_1} \frac{s}{\sqrt{\rho}} \left(\frac{m}{\sqrt{\rho}}s + \eta h - \text{prox}_{\frac{\tau_1}{\kappa}g(\cdot, y)} \left(\frac{m}{\sqrt{\rho}}s + \eta h \right) \right) \right] = 0 \quad (\text{D.6})$$

$$\begin{aligned} \frac{\partial}{\partial \eta} : & -\frac{\tau_2}{2} + \frac{\tau_2 m^2}{2\rho\eta^2} + \alpha \mathbb{E} \left[\frac{\kappa}{\tau_1} h \left(\frac{m}{\sqrt{\rho}}s + \eta h - \text{prox}_{\frac{\tau_1}{\kappa}g(\cdot, y)} \left(\frac{m}{\sqrt{\rho}}s + \eta h \right) \right) \right] \\ & - \frac{1}{2} \frac{\tau_2 \nu^2}{\eta^2} \left(\Omega^{1/2} \Phi^\top \boldsymbol{\theta}_0 \right)^\top G^2 \Omega^{1/2} \Phi^\top \boldsymbol{\theta}_0 - \frac{\tau_2 \kappa^2}{2\eta^2} \text{Tr} \left(G^2 \Omega^2 \right) = 0 \end{aligned} \quad (\text{D.7})$$

$$\frac{\partial}{\partial \tau_1} : \frac{\kappa}{2} - \frac{\kappa}{2\tau_1^2} \alpha \mathbb{E} \left[\left(\frac{m}{\sqrt{\rho}}s + \eta h - \text{prox}_{\frac{\tau_1}{\kappa}g(\cdot, y)} \left(\frac{m}{\sqrt{\rho}}s + \eta h \right) \right)^2 \right] = 0 \quad (\text{D.8})$$

simplifying these equations using Stein's lemma, we get:

$$\frac{\partial}{\partial \kappa} : \frac{\tau_1}{\kappa} = \frac{1}{d} \text{Tr} \left(\Omega^{1/2} \left(\frac{\tau_2}{\eta} \Omega + \lambda_2 \mathbf{I}_d \right)^{-1} \Omega^{1/2} \right) \quad (\text{D.9})$$

$$\frac{\partial}{\partial \nu} : m\sqrt{\gamma} = \frac{1}{d} \nu \boldsymbol{\theta}_0 \Phi \left(\frac{\tau_2}{\eta} \Omega + \lambda_2 \mathbf{I}_d \right)^{-1} \Phi^\top \boldsymbol{\theta}_0 \quad (\text{D.10})$$

$$\frac{\partial}{\partial \tau_2} : \eta^2 + \frac{m^2}{\rho} = \frac{1}{d} \nu^2 \left(\Omega^{1/2} \Phi^\top \boldsymbol{\theta}_0 \right)^\top \left(\frac{\tau_2}{\eta} \Omega + \lambda_2 \mathbf{I}_d \right)^{-2} \left(\Omega^{1/2} \Phi^\top \boldsymbol{\theta}_0 \right) + \frac{1}{d} \kappa^2 \text{Tr} \left(\left(\frac{\tau_2}{\eta} \Omega + \lambda_2 \mathbf{I}_d \right)^{-2} \Omega^2 \right) \quad (\text{D.11})$$

$$\frac{\partial}{\partial m} : \nu\sqrt{\gamma} = \alpha \frac{\kappa}{\sqrt{\rho\tau_1}} \left(\mathbb{E} \left[\text{sprox}_{\frac{\tau_1}{\kappa} g(\cdot, f_0(\sqrt{\rho}s))} \left(\frac{m}{\sqrt{\rho}} + \eta h \right) \right] - \frac{m}{\sqrt{\rho}} \mathbb{E} \left[\text{prox}'_{\frac{\kappa}{\tau_1} g(\cdot, f_0(\sqrt{\rho}s))} \left(\frac{m}{\sqrt{\rho}} s + \eta h \right) \right] \right) \quad (\text{D.12})$$

$$\frac{\partial}{\partial \eta} : \frac{\tau_2}{\eta} = \alpha \frac{\kappa}{\tau_1} \left(1 - \mathbb{E} \left[\text{prox}'_{\frac{\tau_1}{\kappa} g(\cdot, f_0(\sqrt{\rho}s))} \left(\frac{m}{\sqrt{\rho}} s + \eta h \right) \right] \right) \quad (\text{D.13})$$

$$\frac{\partial}{\partial \tau_1} : \kappa^2 = \left(\frac{\kappa}{\tau_1} \right)^2 \alpha \mathbb{E} \left[\left(\frac{m}{\sqrt{\rho}} s + \eta h - \text{prox}_{\frac{\tau_1}{\kappa} g(\cdot, f_0(\sqrt{\rho}s))} \left(\frac{m}{\sqrt{\rho}} s + \eta h \right) \right)^2 \right] \quad (\text{D.14})$$

D.2 Matching with Replica equations

In this section, we show that the fixed point equations obtained from the asymptotic optimality condition of the scalar minimization problem 1 match the ones obtained using the replica method. In what follows we will use the same notations as in [42], and an explicit, clear match with the notations from the proof of the main theorem will be shown. The replica computation, similar to the one from [42], leads to the following fixed point equations, in the replica notations:

$$V = \frac{1}{p} \text{Tr} \left(\lambda \hat{V} I_p + \Omega \right)^{-1} \Omega \quad (\text{D.15})$$

$$q = \frac{1}{p} \text{Tr} \left[\left(\hat{q} \Omega + \hat{m}^2 \Phi^\top \Phi \right) \Omega \left(\lambda \hat{V} I_p + \Omega \right)^{-2} \right] \quad (\text{D.16})$$

$$m = \frac{1}{\sqrt{\gamma}} \frac{\hat{m}}{p} \text{Tr} \left[\Phi^\top \Phi \left(\lambda \hat{V} I_p + \Omega \right)^{-1} \right] \quad (\text{D.17})$$

$$\hat{V} = \alpha \mathbb{E}_\xi \left[\int_{\mathbb{R}} dy \mathcal{Z}_y^0 \left(y, \frac{m}{\sqrt{q}}, \rho - \frac{m^2}{q} \right) \partial_\omega f_g(y, \sqrt{q}\xi, V) \right] \quad (\text{D.18})$$

$$\hat{q} = \alpha \mathbb{E}_\xi \left[\int_{\mathbb{R}} dy \mathcal{Z}_y^0 \left(y, \frac{m}{\sqrt{q}}, \rho - \frac{m^2}{q} \right) f_g(y, \sqrt{q}\xi, V)^2 \right] \quad (\text{D.19})$$

$$\hat{m} = \frac{\alpha}{\sqrt{\gamma}} \mathbb{E}_\xi \left[\int_{\mathbb{R}} dy \partial_\omega \mathcal{Z}_y^0 \left(y, \frac{m}{\sqrt{q}}, \rho - \frac{m^2}{q} \right) f_g(y, \sqrt{q}\xi, V) \right] \quad (\text{D.20})$$

where $f_g(y, \omega, V) = -\partial_\omega \mathcal{M}_{Vg(y, \cdot)}(\omega)$ and \mathcal{Z}_0 is given by:

$$\mathcal{Z}_0(y, \omega, V) = \int \frac{dx}{\sqrt{2\pi V}} e^{-\frac{1}{2V}(x-\omega)^2} \delta(y - f^0(x)). \quad (\text{D.21})$$

In particular we have:

$$\partial_\omega \mathcal{Z}_0(y, \omega, V) = \int \frac{dx}{\sqrt{2\pi V}} e^{-\frac{1}{2V}(x-\omega)^2} \left(\frac{x-\omega}{V} \right) \delta(y - f^0(x)) \quad (\text{D.22})$$

To be explicit with the notation, let's open the equations up. Take for instance the one for \hat{m} . Opening all the integrals:

$$\begin{aligned}\hat{m} &= \int \frac{d\xi}{\sqrt{2\pi}} e^{-\frac{1}{2}\xi^2} \int dy \int \frac{dx}{\sqrt{2\pi}(\rho - m^2/q)} e^{-\frac{1}{2}\frac{(x - \frac{m}{\sqrt{q}}\xi)^2}{\rho - m^2/q}} \left(\frac{x - \frac{m}{\sqrt{q}}\xi}{\rho - m^2/q} \right) f_g(y, \sqrt{q}\xi, V) \\ &\stackrel{(a)}{=} \int \frac{d\xi}{\sqrt{2\pi}} e^{-\frac{1}{2}\xi^2} \int \frac{dx}{\sqrt{2\pi}(\rho - m^2/q)} e^{-\frac{1}{2}\frac{(x - \frac{m}{\sqrt{q}}\xi)^2}{\rho - m^2/q}} \left(\frac{x - \frac{m}{\sqrt{q}}\xi}{\rho - m^2/q} \right) f_g(f_0(x), \sqrt{q}\xi, V)\end{aligned}\tag{D.23}$$

where in (a) we integrated over y explicitly. A direct comparison between the two sets of equations suggests the following mapping to navigate between the replica derivation and the proof using Gaussian comparison theorems. We denote replica quantities with *Rep* indices:

$$\begin{aligned}V_{Rep} &\iff \frac{\tau_1}{\kappa}, & \hat{V}_{Rep} &\iff \frac{\tau_2}{\eta}, & q_{Rep} &\iff \eta^2 + \frac{m^2}{\rho} \\ \hat{q}_{Rep} &\iff \kappa^2, & m_{Rep} &\iff m, & \hat{m}_{Rep} &\iff \nu\end{aligned}\tag{D.24}$$

with these notations, we get :

$$\frac{\partial}{\partial \kappa} : V = \frac{1}{d} \text{Tr}((\hat{V}\Omega + \lambda_2 \mathbf{I}_d)^{-1} \Omega)\tag{D.25}$$

$$\frac{\partial}{\partial \nu} : m = \frac{1}{\sqrt{\gamma}} \frac{\hat{m}}{d} \text{Tr}((\hat{V}\Omega + \lambda_2 \mathbf{I}_d)^{-1} \Phi^\top \Phi)\tag{D.26}$$

$$\frac{\partial}{\partial \tau_2} : q = \frac{1}{d} \text{Tr}((\hat{q}\Omega + \hat{m}^2 \Phi^\top \Phi) \Omega (\hat{V}\Omega + \lambda_2 \mathbf{I}_d)^{-2})\tag{D.27}$$

$$\begin{aligned}\frac{\partial}{\partial m} : \hat{m} &= \frac{\alpha}{\sqrt{\gamma}} \frac{1}{V} \left(\mathbb{E} \left[\frac{s}{\sqrt{\rho}} \text{prox}_{Vg(\cdot, f_0(\sqrt{\rho}s))} \left(\frac{m}{\sqrt{\rho}} s + \sqrt{q - \frac{m^2}{\rho}} h \right) \right] \right. \\ &\quad \left. - \frac{m}{\rho} \mathbb{E} \left[\text{prox}'_{Vg(\cdot, f_0(\sqrt{\rho}s))} \left(\frac{m}{\sqrt{\rho}} s + \sqrt{q - \frac{m^2}{\rho}} h \right) \right] \right)\end{aligned}\tag{D.28}$$

$$\frac{\partial}{\partial \eta} : \hat{V} = \frac{\alpha}{V} \left(1 - \mathbb{E} \left[\text{prox}'_{Vg(\cdot, f_0(\sqrt{\rho}s))} \left(\frac{m}{\sqrt{\rho}} s + \sqrt{q - \frac{m^2}{\rho}} h \right) \right] \right)\tag{D.29}$$

$$\frac{\partial}{\partial \tau_1} : \hat{q} = \left(\frac{\alpha}{V^2} \right) \mathbb{E} \left[\left(\frac{m}{\sqrt{\rho}} s + \sqrt{q - \frac{m^2}{\rho}} h - \text{prox}_{Vg(\cdot, f_0(\sqrt{\rho}s))} \left(\frac{m}{\sqrt{\rho}} s + \sqrt{q - \frac{m^2}{\rho}} h \right) \right)^2 \right]\tag{D.30}$$

The first three equations match the replica prediction, the last three can be exactly matched using the following change of variable and Gaussian integration:

$$\tilde{x} = \frac{x}{\sqrt{\rho}} \quad \tilde{\xi} = \left(\frac{\rho}{\rho - \frac{m^2}{q}} \right)^{1/2} \left(\frac{m}{\sqrt{q\rho}} \tilde{x} - \xi \right)\tag{D.31}$$

E Details on the simulations

In this Appendix we give full details on the numerics used to generate the plots in the main manuscript. An implementation of all the pipelines described below is available at [41].

E.1 Ridge regression on real data

Consider a real data set $\{\mathbf{x}^\mu, y^\mu\}_{\mu=1}^{n_{\text{tot}}}$, where n_{tot} denote the total number of samples available. In Figs. 4 and 4 we work with the MNIST and fashion MNIST data sets for which $n_{\text{tot}} = 6 \times 10^4$ and $D = 28 \times 28 = 784$. In both cases, we center the data and normalise by dividing it by the global standard deviation. We work with binary labels $y^\mu \in \{-1, 1\}$, with $y^\mu = 1$ for even digits (MNIST) or clothes above the waist (fashion MNIST) and $y^\mu = -1$ for odd digits (MNIST) or clothes below the waist (fashion MNIST). In a ridge regression task, we assume $y^\mu = \boldsymbol{\theta}_0^\top \mathbf{u}^\mu$ for a teacher feature map $\mathbf{u}^\mu = \boldsymbol{\varphi}_t(\mathbf{x}^\mu)$ and we are interested in studying the performance of the estimator $\hat{y} = \mathbf{v}^\top \hat{\mathbf{w}}$ where $\mathbf{v} = \boldsymbol{\varphi}_s(\mathbf{x})$ obtained by solving the empirical risk minimisation problem in eq. (A.3) with the squared loss $g(x, y) = \frac{1}{2}(y - x)^2$ and ℓ_2 regularisation $\lambda > 0$.

Simulations: First, we discuss in detail how we conducted the numerical simulations in Figs. 4 and 4 in the main manuscript.

In Fig. 4, the student feature maps $\boldsymbol{\varphi}_s$ is taken to be different transforms used in the literature. For the scattering transform, we have used the out-of-the-box python package `Kymatio` [18] with hyperparameters $J = 3$ and $L = 8$, which defines a feature map $\boldsymbol{\varphi}_s : \mathbb{R}^{28 \times 28} \rightarrow \mathbb{R}^{217 \times 3 \times 3}$, and thus $d = 1953$. For the random features, a random matrix $\mathbf{F} \in \mathbb{R}^{d \times 784}$ with i.i.d. $\mathcal{N}(0, 1/784)$ entries is generated and fixed. Note that the number of features $d = 1953$ is chosen to match the ones for the scattering transform. The random feature map is then applied to the flattened MNIST image as $\boldsymbol{\varphi}_s(\mathbf{x}) = \text{erf}(\mathbf{F}\mathbf{x})$. Finally, we have chosen a kernel corresponding to the limit of this random feature map [56]:

$$K(\mathbf{x}_1, \mathbf{x}_2) = \frac{2}{\pi} \sin^{-1} \left(\frac{2\mathbf{x}_1^\top \mathbf{x}_2}{\sqrt{(1/d + 2\|\mathbf{x}_1\|_2^2)(1/d + 2\|\mathbf{x}_2\|_2^2)}} \right). \quad (\text{E.1})$$

In Fig. 4, the feature $\boldsymbol{\varphi}_s^t$ is taken from a learned neural network at different epochs $t \in \{0, 5, 50, 200\}$ of training. For this experiment, we chose the following architecture implemented in `PyTorch`:

```
Sequential(
  (0): Linear(in_features=784, out_features=2352, bias=False)
  (1): ReLU()
  (2): Linear(in_features=2352, out_features=2352, bias=False)
  (3): ReLU()
  (4): Linear(in_features=2352, out_features=1, bias=False)
)
```

The first two layers of the network therefore defines a feature map $\boldsymbol{\varphi}_s : \mathbb{R}^{784} \rightarrow \mathbb{R}^{2352}$ acting on flattened fashion MNIST images. The network was initialized using the `pyTorch`'s default Kaiming initialisation [72] and was trained on the full data set (n_{tot} samples) with Adam [73] optimiser (learning rate 10^{-3}) on the MSE loss for a total of 500 epochs. Snapshots were taken at epochs $t \in \{0, 5, 50, 200\}$, defining the feature maps $\boldsymbol{\varphi}_s^t(\cdot)$ at each of these epochs.

In both experiments, we ran ridge regression at fixed regularisation $\lambda > 0$ by sub-sampling n samples from the data set $\mathcal{D} = \{\mathbf{v}^\mu, y^\mu\}_{\mu=1}^{n_{\text{tot}}}$, $\mathbf{v}^\mu = \boldsymbol{\varphi}_s(\mathbf{x}^\mu)$, with the estimator given by the closed-form expression:

$$\hat{\mathbf{w}} = \begin{cases} (\lambda \mathbf{I}_d + \mathbf{V}^\top \mathbf{V})^{-1} \mathbf{V}^\top \mathbf{y}, & \text{if } n \geq d \\ \mathbf{V}^\top (\lambda \mathbf{I}_n + \mathbf{V} \mathbf{V}^\top)^{-1} \mathbf{y}, & \text{if } n < d \end{cases} \quad (\text{E.2})$$

where $\mathbf{V} \in \mathbb{R}^{n \times d}$ is the normalised matrix obtained by concatenating $\{\mathbf{v}^\mu / \sqrt{d}\}_{\mu=1}^n$. A similar closed-form expression in terms of the Gram matrix was used in the kernel case. The averaged training and test errors were




Figure 6: Test error as a function of the number of samples for kernel ridge regression task on MNIST odd vs. even data, with $\lambda = 10^{-1}$. The different curves compare the performance of a random features approximation $\varphi_s(\mathbf{x}) = \text{erf}(\mathbf{F}\mathbf{x})$ with the performance of the limiting kernel eq. (E.1). Different curves correspond to different aspect ratios of the Gaussian projection matrix $\mathbf{F} \in \mathbb{R}^{d \times 784}$.

computed over 10 independent draws sub-samples of \mathcal{D} . To reduce the effect spurious correlations due to the sampling of a finite universe \mathcal{D} , we have always evaluated the test error on the whole universe \mathcal{D} . The code for these two experiments is available in [41].

Self-consistent equations: For the theoretical curves, we need to provide the population covariances (Ω, Φ, Ψ) and the teacher weights $\theta_0 \in \mathbb{R}^p$ corresponding to the task of interest. Since when dealing with real data we have a limited number of samples n_{tot} at our disposal, we estimate the population covariances by the empirical covariances on the whole universe:

$$\Psi = \frac{1}{n_{\text{tot}}} \sum_{\mu=1}^{n_{\text{tot}}} \mathbf{u}^{\mu} \mathbf{u}^{\mu \top}, \quad \Phi = \frac{1}{n_{\text{tot}}} \sum_{\mu=1}^{n_{\text{tot}}} \mathbf{u}^{\mu} \mathbf{v}^{\mu \top}, \quad \Omega = \frac{1}{n_{\text{tot}}} \sum_{\mu=1}^{n_{\text{tot}}} \mathbf{v}^{\mu} \mathbf{v}^{\mu \top}. \quad (\text{E.3})$$

In principle, the teacher weights need to be estimated by inverting $\mathbf{y} = \mathbf{U}\theta_0$. However, as explained in n_{tot} in Sec. 3.4, one can avoid doing so by noting the teacher weights only appear in the self-consistent equations 2.17 through $\rho = \frac{1}{k} \theta_0^{\top} \Psi \theta_0$ and $\Phi^{\top} \theta_0$. Therefore, all teacher vector θ_0 and feature map φ_s that linearly interpolate the data set $\{\mathbf{x}^{\mu}, y^{\mu}\}_{\mu=1}^{n_{\text{tot}}}$ are equivalent, since we can write:

$$\rho = \frac{1}{n_{\text{tot}}} \sum_{\mu=1}^{n_{\text{tot}}} (y^{\mu})^2, \quad \Phi^{\top} \theta_0 = \frac{1}{n_{\text{tot}}} \sum_{\mu=1}^{n_{\text{tot}}} \mathbf{v}^{\mu} y^{\mu}. \quad (\text{E.4})$$

which is independent from (φ_t, θ_0) . In particular, note that for our binary labels $y^{\mu} \in \{+1, -1\}$, we have $\rho = 1$. In both Fig. 4 and 4 of the main, we estimated the covariance Ω as in eq. (E.3) by applying the feature maps φ_s described above to the whole data set, took $\rho = 1$ (since in both we have binary labels) and used eq. (E.4) to estimate $\Phi^{\top} \theta_0$. This was then fed to our iterator package [41] to compute the curves. For the kernel curve, we used the random features approximation of eq. (E.1) with a $d = 20 \times 1953$ dimensional feature space to estimate the covariance Ω . We have checked that this indeed provide a good approximation of K for the sample range considered, see Fig.6.

Limitations: As we have discussed above, a key ingredient of our theoretical analysis is the estimation of the population covariances. For real data, this relies on the empirical covariance of the whole data set with n_{tot} samples. We expect this approximation to be good only for $n \ll n_{\text{tot}}$ samples, as it is the case for the ranges

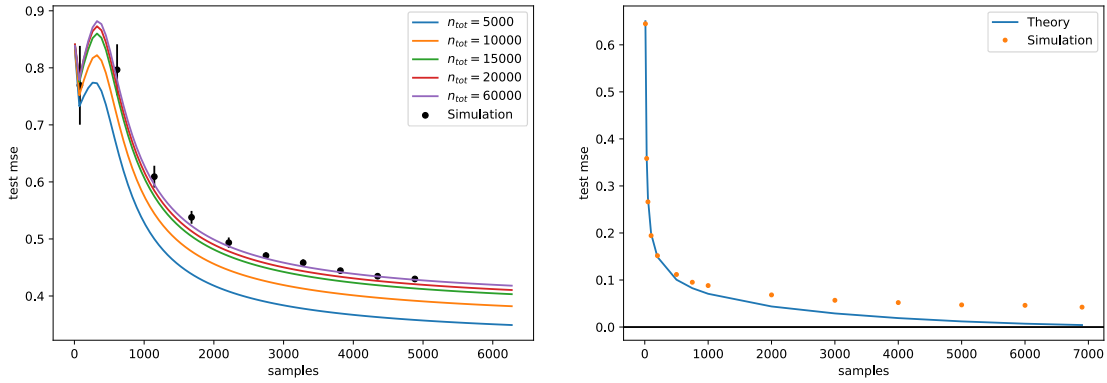


Figure 7: (Left) Test mse for ridge regression on MNIST odd vs. even task and $\lambda = 0.01$. Different curves show the theoretical prediction when the population covariances are estimated using a smaller number of samples n_{tot} in the universe. (Right) Test mse for NTK kernel regression on MNIST 8 vs. 9 task with $\lambda = 0.01$. Note that for this task we have $n_{\text{tot}} = 7000$, and while the theoretical result predicts perfect generalisation as the number of samples approach n_{tot} , the true test error goes to a constant.

plotted in Figs. 4 and 4. Indeed, as $n \approx n_{\text{tot}}$ we start observing deviations between the theoretical prediction and the simulations. In Fig. 7 (right) we show an example of a NTK kernel regression task on 8 vs 9 MNIST digit classification, for which $n_{\text{tot}} = 7000$. Note that while the theoretical prediction reach perfect generalisation at $n \approx n_{\text{tot}}$, the simulated error approaches a plateau. Alternatively, instead of varying the sample range, in Fig. 7 (left) we show how the matching between theory and simulation degrades by varying n_{tot} on a fixed sample range for a MNIST odd vs. even task.

As it was discussed in Sec. 3.4 of the main manuscript, the universality argument sketched above is only valid in the case of a linear student. For instance, applying the same construction to a binary classification task with $f_0(x) = \hat{f}(x) = \text{sign}(x)$ lead to a mismatch between theory and experiments, as exemplified in Fig. 5 of the main for a logistic regression task on CIFAR10 gray-scale images. Interestingly, this is even the case for binary classification with the square loss $g(x, y) = \frac{1}{2}(x - y)^2$, in which the estimator \hat{w} is the same as for ridge regression. In other words, by simply changing the predictor $\hat{f}(x) = \text{sign}(x)$, we have a breakdown of universality, as shown in Fig. 8.




Figure 8: Classification error for binary classification task with the square loss on MNIST odd vs. even task and $\lambda = 0.01$

E.2 Binary classification on GAN generated data

For our purposes, a generative adversarial network (GAN) is a pre-trained neural network defining a map \mathcal{G} taking a Gaussian i.i.d. vector $z \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ (a.k.a. the latent representation) into a realistic looking input image $\mathbf{x} \in \mathbb{R}^D$. In both Figs. 3 and 5, we have used a deep convolutional GAN (dcGAN) [55] with the following architecture

and which has been trained on the full CIFAR10 data set. It therefore takes a 100-dimensional latent vector and returns a $D = 32 \times 32 \times 3 = 3072$ CIFAR10-looking image. The GAN was trained on the original CIFAR10 data set without data augmentation for 50 epochs. Both the discriminator and the generator were trained using Adam, with Adam parameters $\beta_1 = 0.5$ and $\beta_2 = 0.999$. In practice, the advantage of working with a GAN is that we have a generative process to sample as many independent data points as we need, both for the simulations and for the estimation of the population covariances.

Learning the teacher: As discussed in Sec. 3.3 of the main manuscript, to label the GAN generated CIFAR10-looking images we learn a teacher feature map φ_t and weights $\theta_0 \in \mathbb{R}^p$. For the experiments shown in Figs. 3, we have trained with a fully-connected neural network on the full CIFAR10 data set with the following squared architecture:

```
Generator(
  (main): Sequential(
    (0): ConvTranspose2d(100, 512, kernel_size=(4, 4), stride=(1, 1), bias=False)
    (1): BatchNorm2d(512, eps=1e-05, momentum=0.1, affine=True, track_running_stats=True)
    (2): ReLU(inplace=True)
    (3): ConvTranspose2d(512, 256, kernel_size=(4, 4), stride=(2, 2), padding=(1, 1),
      bias=False)
    (4): BatchNorm2d(256, eps=1e-05, momentum=0.1, affine=True, track_running_stats=True)
    (5): ReLU(inplace=True)
    (6): ConvTranspose2d(256, 128, kernel_size=(4, 4), stride=(2, 2), padding=(1, 1),
      bias=False)
    (7): BatchNorm2d(128, eps=1e-05, momentum=0.1, affine=True, track_running_stats=True)
    (8): ReLU(inplace=True)
    (9): ConvTranspose2d(128, 64, kernel_size=(4, 4), stride=(2, 2), padding=(1, 1), bias
      =False)
    (10): BatchNorm2d(64, eps=1e-05, momentum=0.1, affine=True, track_running_stats=True)
    (11): ReLU(inplace=True)
    (12): ConvTranspose2d(64, 3, kernel_size=(1, 1), stride=(1, 1), bias=False)
    (13): Tanh()
  )
)
```

The teacher feature map $\varphi_t : \mathbb{R}^D \rightarrow \mathbb{R}^p$ was then taken to be the first 2-layers, and the teacher weights θ_0 the weights of the last layer, where $D = p = 32 \times 32 \times 3 = 3072$. We used the same architecture for the experiment in Fig. 5, but with $D = p = 32 \times 32 = 1024$ on gray-scale CIFAR10 images. Both teachers were trained on the odd-even discrimination task on CIFAR10 discussed above with the mean-squared error for 50 epochs, starting from pyTorch's default Kaiming initialisation [72]. Optimisation was performed using SGD with momentum 0.9 and weight decay $5 \cdot 10^{-4}$. We started with a learning rate of 0.05, which decayed by a factor 0.1 after 25 and 40 epochs. The resulting trained teacher achieved a 78% classification accuracy on this task. See Fig. 9 for an illustration of this pipeline.

Simulations: The experiment shown in Fig. 3 follow a similar pipeline as the one described in Sec. E.1. The student feature maps φ_s^t are obtained by removing the last layer of a trained a 3-layer student network with architecture:

```
Sequential(
  (0): Linear(in_features=1024, out_features=2304, bias=False)
  (1): ReLU()
  (2): Linear(in_features=2304, out_features=2304, bias=False)
  (3): ReLU()
```

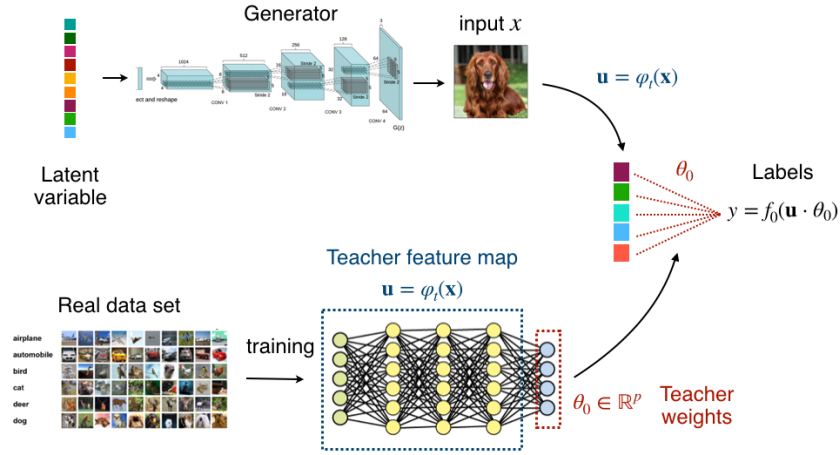



Figure 9: Illustration of the pipeline to generate synthetic realistic data. A dcGAN is first trained to generate CIFAR10-looking images from i.i.d. Gaussian noise. Then, a teacher trained to classify real CIFAR10 images is used to assign labels to the dcGAN generated images.

```
(4 : Linear(in_features=2304, out_features=1, bias=False)
)
```

Training was performed on a data set composed of $n = 30000$ independent samples drawn from the dcGAN described above, with labels $y^\mu \in \{+1, -1\}$ assigned by the learned teacher $y^\mu = \text{sign}(\mathbf{u}^\top \boldsymbol{\theta}_0)$, $\mathbf{u}^\mu = \boldsymbol{\varphi}_t(\mathbf{x}^\mu)$. The network was trained for 300 epochs using Adam optimiser on the MSE loss and `pyTorch`'s default Kaiming initialisation, and snapshots of the weights were extracted at epochs $t \in \{0, 5, 50, 200\}$. Finally, logistic regression was performed on the learned features $\mathbf{v} = \boldsymbol{\varphi}_s(\mathbf{x})$ on fresh pair of dcGAN generated samples and labels using the out-of-the-box `LogisticRegression` solver from `Scikit-learn`. The points and error bars in Fig. 3 were computed by averaging over 10 independent runs. The same pipeline was used for Fig. 5, but for $\boldsymbol{\varphi}_s = \text{id}$ and on dcGAN generated CIFAR10 gray-scale images.

Self-consistent equations: As before, the self-consistent eqs. 2.17 require the population covariances (Ω, Φ, Ψ) and the teacher weights $\boldsymbol{\theta}_0$. For synthetic GAN data, the population covariances of the feature maps $(\boldsymbol{\varphi}_t, \boldsymbol{\varphi}_s^t)$ used in the simulations can be estimated as well as needed with a Monte Carlo sampling algorithm. For the curves shown in Figs. 3 and 5, the covariances were estimated with $n = 10^6$ samples with a precision of the order of 10^{-5} . Together with the teacher weights $\boldsymbol{\theta}_0$ used to generate the labels, this provides everything needed to compute the theoretical learning curves from the self-consistent equations.

F Ridge regression with linear teachers

In this Appendix we discuss briefly random matrix theory, and consider heuristic reasons behind the validity of our asymptotic result beyond Gaussian covariates (\mathbf{u}, \mathbf{v}) in the context of ridge regression, with linear teacher. As is well known, the computation of the training and test MSE for ridge regression can be written as a random matrix theory problem. We do not attempt a rigorous approach, but rather to motivate with simple arguments, many of them actually well known, the observed universality and its limits.

First, let us remind the definition of the model and introduce some simplifications that arise in ridge regression task. We have Gaussian covariates vectors $\mathbf{u} \in \mathbb{R}^p$ and $\mathbf{v} \in \mathbb{R}^d$, with correlations matrices Ψ, Ω and Φ , from which we draw n independent samples:

$$\begin{bmatrix} \mathbf{u} \\ \mathbf{v} \end{bmatrix} \in \mathbb{R}^{p+d} \sim \mathcal{N}\left(0, \begin{bmatrix} \Psi & \Phi \\ \Phi^\top & \Omega \end{bmatrix}\right). \quad (\text{F.1})$$

We assume the existence of a linear teacher generating the labels $\mathbf{y} = \mathbf{U}\boldsymbol{\theta}_0$, and recall the student performs ridge regression on the data matrix \mathbf{V} .

Note that since ridge regression can be performed in any basis, we might as well work in the basis where the population covariance Ψ of the vector \mathbf{u} is diagonal. Additionally, we shall use the fact that one can consider a $\boldsymbol{\theta}_0$ to be an i.i.d. Rademacher vector, i.e. a random vector of ± 1 *without loss of generality*. Indeed, the statistical properties of the random variable $\mathbf{u} \cdot \boldsymbol{\theta}_0$, for a generic $\boldsymbol{\theta}_0$, and of the random variable $\tilde{\mathbf{u}} \cdot \boldsymbol{\theta}$, with $\boldsymbol{\theta}$ a Rademacher vector are identical provided a change in the (diagonal) covariance:

$$\Psi = \begin{pmatrix} \Psi_1(\boldsymbol{\theta}_0)_1^2 & 0 & \dots & 0 \\ 0 & \Psi_2(\boldsymbol{\theta}_0)_2^2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \Psi_p(\boldsymbol{\theta}_0)_p^2 \end{pmatrix} \quad (\text{F.2})$$

The Gaussian model we consider can therefore be rewritten with a Rademacher vector $\boldsymbol{\theta}$ provided we change the correlation matrix Ψ (as well as the cross-correlation Φ) accordingly.

We now come back on the problem. Given the vector \mathbf{y} and the data $\mathbf{V} \in \mathbb{R}^{n \times d}$, the ridge estimator has the following closed-form solution:

$$\hat{\mathbf{w}} = \left(\frac{1}{n}\mathbf{V}^\top\mathbf{V} + \lambda\mathbf{I}_d\right)^{-1} \mathbf{V}^\top\mathbf{y} = (\mathbf{S}_{v,v} + \lambda\mathbf{I}_d) \mathbf{S}_{u,v}^\top \boldsymbol{\theta} \quad (\text{F.3})$$

where we have defined the *empirical* covariance matrices

$$\mathbf{S}_{u,u} \equiv \frac{1}{n}\mathbf{U}^\top\mathbf{U}, \quad \mathbf{S}_{u,v} \equiv \frac{1}{n}\mathbf{U}^\top\mathbf{V}, \quad \mathbf{S}_{v,v} \equiv \frac{1}{n}\mathbf{V}^\top\mathbf{V}. \quad (\text{F.4})$$

Given this vector, one can now readily write the expected value of the training and test losses as follows:

$$\begin{aligned} \mathcal{E}_{\text{train.}} &= \mathbb{E}_{\mathbf{U}, \mathbf{V}, \boldsymbol{\theta}} \left[\frac{1}{n} \|\mathbf{U}\boldsymbol{\theta} - \mathbf{V}\hat{\mathbf{w}}(\mathbf{U}, \mathbf{V})\|_2^2 \right] \\ &= \mathbb{E}_{\mathbf{U}, \mathbf{V}, \boldsymbol{\theta}} \left[\frac{1}{n} \boldsymbol{\theta}^\top \mathbf{U}^\top \mathbf{U} \boldsymbol{\theta} \right] + \mathbb{E}_{\mathbf{U}, \mathbf{V}} \left[\frac{1}{n} \hat{\mathbf{w}}(\mathbf{U}, \mathbf{V})^\top \mathbf{V}^\top \mathbf{V} \hat{\mathbf{w}}(\mathbf{U}, \mathbf{V}) \right] - 2\mathbb{E}_{\mathbf{U}, \mathbf{V}} \left[\frac{1}{n} \boldsymbol{\theta}^\top \mathbf{U}^\top \mathbf{V} \hat{\mathbf{w}}(\mathbf{U}, \mathbf{V}) \right] \\ &= \mathbb{E} [\text{Tr } \mathbf{S}_{u,u}] + \mathbb{E} \left[\text{Tr } \mathbf{S}_{u,v} (\mathbf{S}_{v,v} + \lambda\mathbf{I}_d)^{-1} \mathbf{S}_{v,v} (\mathbf{S}_{v,v} + \lambda\mathbf{I}_d)^{-1} \mathbf{S}_{u,v}^\top \right] - 2\mathbb{E} \left[\text{Tr } \mathbf{S}_{u,v} (\mathbf{S}_{v,v} + \lambda\mathbf{I}_d)^{-1} \mathbf{S}_{u,v}^\top \right] \end{aligned} \quad (\text{F.5})$$

and

$$\begin{aligned} \mathcal{E}_{\text{gen.}} &= \mathbb{E}_{\mathbf{U}, \mathbf{V}, \mathbf{u}, \mathbf{v}, \boldsymbol{\theta}} \left[\frac{1}{n} \|\mathbf{u}^\top \boldsymbol{\theta} - \mathbf{v}^\top \hat{\mathbf{w}}(\mathbf{U}, \mathbf{V})\|_2^2 \right] \\ &= \mathbb{E}_{\mathbf{U}, \mathbf{V}, \mathbf{u}, \mathbf{v}, \boldsymbol{\theta}} \left[\frac{1}{n} \boldsymbol{\theta}^\top \mathbf{u} \mathbf{u}^\top \boldsymbol{\theta} \right] + \mathbb{E}_{\mathbf{U}, \mathbf{V}} \left[\frac{1}{n} \hat{\mathbf{w}}(\mathbf{U}, \mathbf{V})^\top \mathbf{v} \mathbf{v}^\top \hat{\mathbf{w}}(\mathbf{U}, \mathbf{V}) \right] - 2\mathbb{E}_{\mathbf{U}, \mathbf{V}} \left[\frac{1}{n} \boldsymbol{\theta}^\top \mathbf{u} \mathbf{v}^\top \hat{\mathbf{w}}(\mathbf{U}, \mathbf{V}) \right] \\ &= \text{Tr } \Sigma_{u,u} + \mathbb{E} \left[\text{Tr } \mathbf{S}_{u,v} (\mathbf{S}_{v,v} + \lambda\mathbf{I}_d)^{-1} \Sigma_{v,v} (\mathbf{S}_{v,v} + \lambda\mathbf{I}_d)^{-1} \mathbf{S}_{u,v}^\top \right] - 2\mathbb{E} \left[\text{Tr } \Sigma_{u,v} (\mathbf{S}_{v,v} + \lambda\mathbf{I}_d)^{-1} \mathbf{S}_{u,v}^\top \right] \end{aligned} \quad (\text{F.6})$$

where we have denoted the *population* correlation matrices $\Psi \equiv \Sigma_{u,u}$, $\Phi \equiv \Sigma_{u,v}$, $\Omega \equiv \Sigma_{v,v}$ for readability and a direct comparison with their empirical counterpart. The traces appears by the left and right multiplication by the random vector $\boldsymbol{\theta}$.

At this point, the entire problem has been mapped to a random matrix theory exercise: assuming data are indeed Gaussian, one can use RMT to compute the six traces that appears in (F.5,F.6). Indeed, this is the canonical approach used in most rigorous works for the ridge regression task in the teacher-student framework, instance in [32, 12, 34, 33]. Remarkably, the replica (and the rigorous Gordon counterpart) allow to find the same result without the explicit use of RMT.

We now discuss, heuristically, why these results are valid even though the distribution of $[\mathbf{u}, \mathbf{v}]$ is not actually Gaussian, and in some instances even for real data. Indeed, that both $(\mathcal{E}_{\text{train.}}, \mathcal{E}_{\text{gen.}})$ *do not depend* explicitly on the distribution of the data, but —assuming some concentration (or self-averaging)— only on:

1. The spectrum of the population covariances $\Sigma_{u,u}, \Sigma_{u,v}, \Sigma_{v,v}$.
2. The spectrum of the empirical covariances $S_{u,u}, S_{u,v}, S_{v,v}$.
3. The expectation of the trace of products between empirical and population covariances.

We expect that asymptotically the prediction from the theory will thus be valid for much more generic distributions $[\mathbf{u}, \mathbf{v}] \sim P_{\mathbf{u},\mathbf{v}}$, provided they share the same population covariances $\Sigma_{u,u}, \Sigma_{u,v}, \Sigma_{v,v}$ (which we call Ψ, Φ, Ω). To see this, we need to check how this change in distribution would affect points (1),(2) and (3). Fixing the population covariances, the first bullet point (1) is automatically taken into account. Point (2) and (3) are, however, less trivial: in order to have universality we need that a) the spectrum of the *empirical* covariances of the non-Gaussian distribution to converge the one obtained with the Gaussian one; and b) the trace of products between the *empirical* and the *population* covariances also to converge to the universal values computed from Gaussians data.

These two last points have been investigated in RMT [74], and it is a classical result that such quantities are universal and converge to the Gaussian-predicted values for many distribution, way beyond the Gaussian assumption (in which case the spectral densities are known as the Wigner and Wishart model, or Marcenko-Pastur distribution [75]): this powerful universality of RMT is at the origin of the applicability of the model beyond Gaussian data. For instance, [60] showed that these assumptions are verified for any data generated as $\mathbf{u} = \Sigma_{\mathbf{u},\mathbf{u}}^{1/2} \boldsymbol{\omega}$, assuming the components of the vector $\boldsymbol{\omega}$ are drawn i.i.d. from *any* distribution (with some assumption on the larger moments). While this is still restrictive, stronger results can be shown, and [59, 61, 76, 47] extended them (also loosening the independence assumption) for a very generic class of distributions of correlated random vectors \mathbf{u} .

Let us give a concrete example. For simplicity, consider the restricted case where $\mathbf{u} = \mathbf{v}$, i.e. the teacher acts on the same space as the student. In this case, eqs. (F.5,F.6) simplify (this is essentially the analysis in [32]) to:

$$\begin{aligned}
\mathcal{E}_{\text{train.}} &= \mathbb{E}_{\mathbf{U},\boldsymbol{\theta}} \left[\frac{1}{n} \|\mathbf{U}\boldsymbol{\theta} - \mathbf{U}\hat{\mathbf{w}}(\mathbf{U})\|_2^2 \right] \\
&= \mathbb{E}_{\mathbf{U},\boldsymbol{\theta}} \left[\frac{1}{n} \boldsymbol{\theta}^\top \mathbf{U}^\top \mathbf{U} \boldsymbol{\theta} \right] + \mathbb{E}_{\mathbf{U}} \left[\frac{1}{n} \hat{\mathbf{w}}(\mathbf{U})^\top \mathbf{U}^\top \mathbf{U} \hat{\mathbf{w}}(\mathbf{U}) \right] - 2\mathbb{E}_{\mathbf{U}} \left[\frac{1}{n} \boldsymbol{\theta}^\top \mathbf{U}^\top \mathbf{U} \hat{\mathbf{w}}(\mathbf{U}) \right] \\
&= \mathbb{E} [\text{Tr } S] + \mathbb{E} \left[\text{Tr } S (S + \lambda \mathbf{I}_d)^{-1} S (S + \lambda \mathbf{I}_d)^{-1} S^\top \right] - 2\mathbb{E} \left[\text{Tr } S (S + \lambda \mathbf{I}_d)^{-1} S^\top \right] \quad (\text{F.7})
\end{aligned}$$

and

$$\begin{aligned}
\mathcal{E}_{\text{gen.}} &= \mathbb{E}_{\mathbf{U},\mathbf{u},\boldsymbol{\theta}} \left[\frac{1}{n} \|\mathbf{u}^\top \boldsymbol{\theta} - \mathbf{u}^\top \hat{\mathbf{w}}(\mathbf{U})\|_2^2 \right] \\
&= \mathbb{E}_{\mathbf{U},\mathbf{u},\boldsymbol{\theta}} \left[\frac{1}{n} \boldsymbol{\theta}^\top \mathbf{u} \mathbf{u}^\top \boldsymbol{\theta} \right] + \mathbb{E}_{\mathbf{U}} \left[\frac{1}{n} \hat{\mathbf{w}}(\mathbf{U})^\top \mathbf{u} \mathbf{u}^\top \hat{\mathbf{w}}(\mathbf{U}) \right] - 2\mathbb{E}_{\mathbf{U}} \left[\frac{1}{n} \boldsymbol{\theta}^\top \mathbf{u} \mathbf{u}^\top \hat{\mathbf{w}}(\mathbf{U}) \right] \\
&= \text{Tr } \Sigma_{u,u} + \mathbb{E} \left[\text{Tr } S (S + \lambda \mathbf{I}_d)^{-1} \Sigma (S + \lambda \mathbf{I}_d)^{-1} S^\top \right] - 2\mathbb{E} \left[\text{Tr } \Sigma (S + \lambda \mathbf{I}_d)^{-1} S^\top \right] \quad (\text{F.8})
\end{aligned}$$

In the expression of the training loss eq. (F.7), we see terms such as

$$\mathcal{A} = \mathbb{E} \left[\text{Tr} S (S + \lambda \mathbf{I}_d)^{-1} S^\top \right], \quad (\text{F.9})$$

depend only on the limiting distribution of eigenvalues of $S \in \mathbb{R}^{d \times d}$. This is a very well known problem when the dimension d and the number of samples n are sent to infinity with fixed ratio $\alpha = n/d$, and the limiting spectral density is known as the Marcenko-Pastur law. This is a very robust distribution that is valid way beyond the Gaussian hypothesis [59, 61, 76, 47].

In the expression of the generalisation loss eq. (F.8), however, terms such as

$$\mathcal{B} = \mathbb{E} \left[\text{Tr} \Sigma (S + \lambda \mathbf{I}_d)^{-1} S^\top \right]. \quad (\text{F.10})$$

appears. These can be computed using classical RMT results on the concentration of the inverse of the covariance [77, 60]. The strongest result we are aware of for such problems is from the remarkable work of [47]. This universality of random matrix theory is thus at the origin of the surprisingly successful application of our Gaussian theory to real data with arbitrary feature maps. Of course the discussion here is limited to the case where $\mathbf{u} = \mathbf{v}$ and a concrete mathematical statement would require the generalisation of these arguments to the more generic case of eqs.(F.5,F.6), which are closer to the work of [33]. We leave this discussion to future works.

A similar universality has been discussed for kernel methods in very recent works, but for the slightly different setting in which data is drawn from a Mixture of Gaussians [34, 48] (in which case there is no teacher, the label depends on which Gaussian has been chosen). The universality observed here for ridge regression with linear student, albeit different, is of a similar nature, and it would be interesting to discuss the link between these two approaches.

References

- [1] Hyunjune Sebastian Seung, Haim Sompolinsky, and Naftali Tishby. Statistical mechanics of learning from examples. *Physical review A*, 45(8):6056, 1992.
- [2] Timothy LH Watkin, Albrecht Rau, and Michael Biehl. The statistical mechanics of learning a rule. *Reviews of Modern Physics*, 65(2):499, 1993.
- [3] Andreas Engel and Christian Van den Broeck. *Statistical mechanics of learning*. Cambridge University Press, 2001.
- [4] David L Donoho, Arian Maleki, and Andrea Montanari. Message-passing algorithms for compressed sensing. *Proceedings of the National Academy of Sciences*, 106(45):18914–18919, 2009.
- [5] Noureddine El Karoui, Derek Bean, Peter J Bickel, Chinghay Lim, and Bin Yu. On robust regression with high-dimensional predictors. *Proceedings of the National Academy of Sciences*, 110(36):14557–14562, 2013.
- [6] Lenka Zdeborová and Florent Krzakala. Statistical physics of inference: Thresholds and algorithms. *Advances in Physics*, 65(5):453–552, 2016.
- [7] David Donoho and Andrea Montanari. High dimensional robust m-estimation: Asymptotic variance via approximate message passing. *Probability Theory and Related Fields*, 166(3-4):935–969, 2016.
- [8] C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals. Understanding deep learning requires rethinking generalization. In *ICLR*, 2017.
- [9] Mikhail Belkin, Daniel Hsu, Siyuan Ma, and Soumik Mandal. Reconciling modern machine-learning practice and the classical bias–variance trade-off. *Proceedings of the National Academy of Sciences*, 116(32):15849–15854, 2019.
- [10] Mikhail Belkin, Daniel Hsu, and Ji Xu. Two models of double descent for weak features. *SIAM Journal on Mathematics of Data Science*, 2(4):1167–1180, 2020.
- [11] Song Mei and Andrea Montanari. The generalization error of random features regression: Precise asymptotics and double descent curve. *arXiv preprint arXiv:1908.05355*, 2019.
- [12] Trevor Hastie, Andrea Montanari, Saharon Rosset, and Ryan J Tibshirani. Surprises in high-dimensional ridgeless least squares interpolation. *arXiv preprint arXiv:1903.08560*, 2019.
- [13] Emmanuel J Candès, Pragma Sur, et al. The phase transition for the existence of the maximum likelihood estimate in high-dimensional logistic regression. *The Annals of Statistics*, 48(1):27–42, 2020.
- [14] Benjamin Aubin, Florent Krzakala, Yue M Lu, and Lenka Zdeborová. Generalization error in high-dimensional perceptrons: Approaching bayes error with convex optimization. In *Advances in Neural Information Processing Systems*, volume 33, 2020.
- [15] Fariborz Salehi, Ehsan Abbasi, and Babak Hassibi. The performance analysis of generalized margin maximizers on separable data. In *International Conference on Machine Learning*, pages 8417–8426. PMLR, 2020.
- [16] Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In *Advances in neural information processing systems*, pages 1177–1184, 2008.
- [17] Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in neural information processing systems*, pages 8571–8580, 2018.

- [18] Mathieu Andreux, Tomás Angles, Georgios Exarchakis, Roberto Leonarduzzi, Gaspar Rochette, Louis Thiry, John Zarka, Stéphane Mallat, Joakim Andén, Eugene Belilovsky, et al. Kymatio: Scattering transforms in python. *Journal of Machine Learning Research*, 21(60):1–6, 2020.
- [19] Antonia Creswell, Tom White, Vincent Dumoulin, Kai Arulkumaran, Biswa Sengupta, and Anil A Bharath. Generative adversarial networks: An overview. *IEEE Signal Processing Magazine*, 35(1):53–65, 2018.
- [20] Yehoram Gordon. Some inequalities for gaussian processes and applications. *Israel Journal of Mathematics*, 50(4):265–289, 1985.
- [21] Marc Mézard, Giorgio Parisi, and Miguel Virasoro. *Spin glass theory and beyond: An Introduction to the Replica Method and Its Applications*, volume 9. World Scientific Publishing Company, 1987.
- [22] Marc Mézard and Andrea Montanari. *Information, physics, and computation*. Oxford University Press, 2009.
- [23] Elizabeth Gardner and Bernard Derrida. Three unfinished works on the optimal storage capacity of networks. *Journal of Physics A: Mathematical and General*, 22(12):1983, 1989.
- [24] Manfred Opper and Wolfgang Kinzel. Statistical mechanics of generalization. In *Models of neural networks III*, pages 151–209. Springer, 1996.
- [25] Behrooz Ghorbani, Song Mei, Theodor Misiakiewicz, and Andrea Montanari. When do neural networks outperform kernel methods? In *Advances in Neural Information Processing Systems*, volume 33, 2020.
- [26] Christos Thrampoulidis, Ehsan Abbasi, and Babak Hassibi. Precise error analysis of regularized m -estimators in high dimensions. *IEEE Transactions on Information Theory*, 64(8):5592–5628, 2018.
- [27] Andrea Montanari, Feng Ruan, Youngtak Sohn, and Jun Yan. The generalization error of max-margin linear classifiers: High-dimensional asymptotics in the overparametrized regime. *arXiv preprint arXiv:1911.01544*, 2019.
- [28] Michael Celentano, Andrea Montanari, and Yuting Wei. The lasso with general gaussian designs with applications to hypothesis testing. *arXiv preprint arXiv:2007.13716*, 2020.
- [29] Mihailo Stojnic. A framework to characterize performance of lasso algorithms. *arXiv preprint arXiv:1303.7291*, 2013.
- [30] Samet Oymak, Christos Thrampoulidis, and Babak Hassibi. The squared-error of generalized lasso: A precise analysis. In *2013 51st Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 1002–1009. IEEE, 2013.
- [31] Oussama Dhifallah and Yue M Lu. A precise performance analysis of learning with random features. *arXiv preprint arXiv:2008.11904*, 2020.
- [32] Edgar Dobriban, Stefan Wager, et al. High-dimensional asymptotics of prediction: Ridge regression and classification. *The Annals of Statistics*, 46(1):247–279, 2018.
- [33] Denny Wu and Ji Xu. On the optimal weighted ℓ_2 regularization in overparameterized linear regression. In *Advances in Neural Information Processing Systems*, volume 33, 2020.
- [34] Zhenyu Liao, Romain Couillet, and Michael W Mahoney. A random matrix analysis of random fourier features: beyond the gaussian kernel, a precise phase transition, and the corresponding double descent. In *Advances in Neural Information Processing Systems*, volume 33, 2020.
- [35] Fanghui Liu, Zhenyu Liao, and Johan AK Suykens. Kernel regression in high dimension: Refined analysis beyond double descent. *arXiv preprint arXiv:2010.02681*, 2020.

- [36] Cédric Gerbelot, Alia Abbara, and Florent Krzakala. Asymptotic errors for high-dimensional convex penalized linear regression beyond gaussian matrices. In *Conference on Learning Theory*, pages 1682–1713. PMLR, 2020.
- [37] Blake Bordelon, Abdulkadir Canatar, and Cengiz Pehlevan. Spectrum dependent learning curves in kernel regression and wide neural networks. In *International Conference on Machine Learning*, pages 1024–1034. PMLR, 2020.
- [38] Hanwen Huang and Qinglong Yang. Large scale analysis of generalization error in learning using margin based classification methods. *Journal of Statistical Mechanics: Theory and Experiment*, 2020(10):103407, 2020.
- [39] Sebastian Goldt, Galen Reeves, Marc Mézard, Florent Krzakala, and Lenka Zdeborová. The gaussian equivalence of generative models for learning with two-layer neural networks. *arXiv preprint arXiv:2006.14709*, 2020.
- [40] Partha P Mitra. Understanding overfitting peaks in generalization error: Analytical risk curves for l_2 and l_1 penalized interpolation. *arXiv preprint arXiv:1906.03667*, 2019.
- [41] Bruno Loureiro, Hugo Cui, and Sebastian Goldt. Gaussian covariate model and applications. <https://github.com/IdePHICS/GCMPProject>, 2020.
- [42] F. Gerace, B. Loureiro, F. Krzakala, M. Mézard, and L. Zdeborová. Generalisation error in learning with random features and the hidden manifold model. In *37th International Conference on Machine Learning*, 2020.
- [43] S. Goldt, M. Mézard, F. Krzakala, and L. Zdeborová. Modeling the influence of data structure on learning in neural networks: The hidden manifold model. *Phys. Rev. X*, 10(4):041044, 2020.
- [44] Hong Hu and Yue M Lu. Universality laws for high-dimensional learning with random features. *arXiv preprint arXiv:2009.07669*, 2020.
- [45] Noureddine El Karoui et al. The spectrum of kernel random matrices. *Annals of statistics*, 38(1):1–50, 2010.
- [46] Jeffrey Pennington and Pratik Worah. Nonlinear random matrix theory for deep learning. In *Advances in Neural Information Processing Systems*, volume 30, pages 2637–2646, 2017.
- [47] Cosme Louart and Romain Couillet. Concentration of measure and large random matrices with an application to sample covariance matrices. *arXiv preprint arXiv:1805.08295*, 2018.
- [48] Mohamed El Amine Seddik, Cosme Louart, Mohamed Tamaazousti, and Romain Couillet. Random matrix theory proves that deep learning representations of gan-data behave as gaussian mixtures. In *International Conference on Machine Learning*, pages 8573–8582. PMLR, 2020.
- [49] B. Scholkopf and A.J. Smola. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. Adaptive Computation and Machine Learning. MIT Press, 2018.
- [50] Ingo Steinwart, Don R Hush, Clint Scovel, et al. Optimal rates for regularized least squares regression. In *COLT*, pages 79–93, 2009.
- [51] Andrea Caponnetto and Ernesto De Vito. Optimal rates for the regularized least-squares algorithm. *Foundations of Computational Mathematics*, 7(3):331–368, 2007.
- [52] Loucas Pillaud-Vivien, Alessandro Rudi, and Francis Bach. Statistical optimality of stochastic gradient descent on hard learning problems through multiple passes. In *Advances in Neural Information Processing Systems*, volume 31, pages 8114–8124, 2018.

- [53] Stefano Spigler, Mario Geiger, and Matthieu Wyart. Asymptotic learning curves of kernel methods: empirical data versus teacher–student paradigm. *Journal of Statistical Mechanics: Theory and Experiment*, 2020(12):124001, 2020.
- [54] Saharon Rosset, Ji Zhu, and Trevor Hastie. Margin maximizing loss functions. In *NIPS*, pages 1237–1244, 2003.
- [55] A. Radford, L. Metz, and S. Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. In *ICLR*, 2016.
- [56] Christopher K. I. Williams. Computing with infinite networks. In *Proceedings of the 9th International Conference on Neural Information Processing Systems*, NIPS’96, page 295–301, Cambridge, MA, USA, 1996. MIT Press.
- [57] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms, 2017.
- [58] J. Bruna and S. Mallat. Invariant scattering convolution networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8):1872–1886, 2013.
- [59] Zhidong Bai and Wang Zhou. Large sample covariance matrices without independence structures in columns. *Statistica Sinica*, pages 425–442, 2008.
- [60] Olivier Ledoit and Sandrine Péché. Eigenvectors of some large sample covariance matrix ensembles. *Probability Theory and Related Fields*, 151(1):233–264, 2011.
- [61] Noureddine El Karoui et al. Concentration of measure and spectra of random matrices: Applications to correlation matrices, elliptical distributions and beyond. *Annals of Applied Probability*, 19(6):2362–2405, 2009.
- [62] Xiuyuan Cheng and Amit Singer. The spectrum of random inner-product kernel matrices. *Random Matrices: Theory and Applications*, 2(04):1350010, 2013.
- [63] Zhou Fan and Andrea Montanari. The spectral norm of random inner-product kernel matrices. *Probability Theory and Related Fields*, 173(1):27–85, 2019.
- [64] Heinz H Bauschke, Patrick L Combettes, et al. *Convex analysis and monotone operator theory in Hilbert spaces*, volume 408. Springer, 2011.
- [65] Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration inequalities: A nonasymptotic theory of independence*. Oxford university press, 2013.
- [66] Yanting Ma, Cynthia Rush, and Dror Baron. Analysis of approximate message passing with a class of non-separable denoisers. In *2017 IEEE International Symposium on Information Theory (ISIT)*, pages 231–235. IEEE, 2017.
- [67] Roman Vershynin. Introduction to the non-asymptotic analysis of random matrices. In Y. Eldar and G. Kutyniok., editors, *Compressed Sensing, Theory and Applications*. Cambridge University Press, 2012.
- [68] Rick Durrett. *Probability: theory and examples*, volume 49. Cambridge university press, 2019.
- [69] Per Kragh Andersen and Richard D Gill. Cox’s regression model for counting processes: a large sample study. *The annals of statistics*, pages 1100–1120, 1982.
- [70] Léo Miolane and Andrea Montanari. The distribution of the lasso: Uniform control over sparse balls and adaptive parameter tuning. *arXiv preprint arXiv:1811.01212*, 2018.

- [71] Roman Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018.
- [72] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 1026–1034, 2015.
- [73] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *Proceedings of the 3rd International Conference for Learning Representations*, volume 3, 2015.
- [74] Greg W Anderson, Alice Guionnet, and Ofer Zeitouni. *An introduction to random matrices*. Number 118. Cambridge university press, 2010.
- [75] Vladimir Alexandrovich Marchenko and Leonid Andreevich Pastur. Distribution of eigenvalues for some sets of random matrices. *Matematicheskii Sbornik*, 114(4):507–536, 1967.
- [76] Djalil Chafaï and Konstantin Tikhomirov. On the convergence of the extremal eigenvalues of empirical covariance matrices with dependence. *Probability Theory and Related Fields*, 170(3):847–889, 2018.
- [77] Walid Hachem, Philippe Loubaton, Jamal Najim, et al. Deterministic equivalents for certain functionals of large random matrices. *The Annals of Applied Probability*, 17(3):875–930, 2007.