

Universality in Block Dependent Linear Models With Applications to Nonlinear Regression

Samriddha Lahiry and Pragya Sur[✉]

Abstract—Over the past decade, characterizing the precise asymptotic risk of regularized estimators in high-dimensional regression has emerged as a prominent research area. This literature focuses on the proportional asymptotics regime, where the number of features and samples diverge proportionally. Much of this work assumes i.i.d. Gaussian entries in the design. Concurrently, researchers have explored the universality of these findings, discovering that results based on the i.i.d. Gaussian assumption extend to other settings, including i.i.d. sub-Gaussian designs. However, universality results examining dependent covariates have predominantly focused on correlation-based dependence or structured forms of dependence allowed by right-rotationally-invariant designs. In this paper, we challenge this limitation by investigating dependence structures beyond these established classes. We identify a class of designs characterized by a block dependence structure where results based on i.i.d. Gaussian designs persist. Formally, we establish that the optimal values of regularized empirical risk and the risk associated with convex regularized estimators, such as the Lasso and the ridge, converge to the same limit under block-dependent designs as for i.i.d. Gaussian entry designs. Our dependence structure differs significantly from correlation-based dependence and enables, for the first time, asymptotically exact risk characterization in prevalent high-dimensional nonlinear regression problems.

Index Terms—Block dependent design, nonlinear regression, universality.

I. INTRODUCTION

OVER the past decade, a novel regime to exploring high-dimensional asymptotics has emerged in supervised learning. This paradigm posits that the number of features (p) and the number of samples (n) both diverge with the ratio p/n converging to a positive constant. For natural reasons, this paradigm is referred to as the proportional asymptotics regime. Numerous parallel techniques have emerged to address statistical and probabilistic questions within this framework, including random matrix theory [1], approximate message passing theory [2], [3], [4], [5], [6], [7], [8], [9], [10], [11].

Received 5 March 2024; revised 31 July 2024; accepted 1 October 2024. Date of publication 16 October 2024; date of current version 22 November 2024. The work of Pragya Sur was supported by NSF under Grant DMS-2113426. (Corresponding author: Pragya Sur.)

Samriddha Lahiry is with the Department of Statistics and Data Science, National University of Singapore, Singapore 119077 (e-mail: slahiry@nus.edu.sg).

Pragya Sur is with the Department of Statistics, Harvard University, Cambridge, MA 02138 USA (e-mail: pragya@fas.harvard.edu).

Communicated by Y. Chen, Associate Editor for Machine Learning and Statistics.

Color versions of one or more figures in this article are available at <https://doi.org/10.1109/TIT.2024.3481956>.

Digital Object Identifier 10.1109/TIT.2024.3481956

[12], [13], [14], the leave-one-out/cavity method [7], [9], [15], [16], [17], [18], [19], [20], [21], [22], convex Gaussian min-max theorem [23], [24], [25], [26], [27], [28], etc. The widespread adoption of this paradigm in statistical research can, in part, be attributed to its precise asymptotic results, which showcase exceptional performance when tested on finite samples (cf. [7], [9], [13], [21], [29], [30], [31], [32], [33], [34] and references cited therein). Moreover, this paradigm liberates researchers from stringent sparsity assumptions on underlying signals, accommodating structured classes of both dense and sparse signals.

However, a limitation arises within this paradigm—conventional tools often assume Gaussianity in the distribution of observed covariates. Extensive prior literature [33], [35], [36], [37], [38], [39], [40], [41], [42], [43], [44], [45], [46], [47], [48], [49], [50] suggests that this assumption is hardly a conceptual constraint. The results derived under this paradigm are often insensitive to the specific form of the covariate distribution, and rely more on the distribution's tails or on quantities such as the spectrum of the sample covariance matrix. In particular, prior work demonstrates the seamless extension of results derived under i.i.d. Gaussian designs to i.i.d. sub-Gaussian designs with suitable moments. Nevertheless, universality results under dependent design entries are relatively scarce. This paper addresses this challenge, revealing that under an appropriate notion of block dependence, universality results persist.

Our universality result substantially broadens the scope of statistical models considered within the proportional asymptotics regime. While the extension of universality beyond the independent model is intrinsically fascinating, we demonstrate that our results apply to widely used statistical models that had thus far remained outside the purview of this literature. Specifically, we leverage our universality result to provide precise risk characterization of estimators arising in nonlinear regression in the presence of high-dimensional covariates (Section II). To the best of our knowledge, this class of problems had so far remained uncharted. Our result relies on a block dependence structure and we exhibit three popular examples where such dependence emerges naturally. The first example pertains to the additive model, well-studied in statistics, and applied in diverse fields such as biomedical research [51], [52], agronomy [53], environmental science [54] etc. The second example, scalar- on-functional regression, arises from the expansive field of functional data analysis, motivated by longitudinal data, growth curves, time series, and more [55].

The third and final example is from genomics, particularly linkage-disequilibrium models [56], where block dependence structures arise naturally.

To introduce our setup formally, let $(Y_i, \mathbb{X}_i)_{i=1}^n$ be i.i.d. random variables with $Y_i \in \mathbb{R}$ and $\mathbb{X}_i \in \mathbb{R}^p$. We assume that Y_i is a function of \mathbb{X}_i corrupted by an additive noise i.e.

$$Y_i = g^*(\mathbb{X}_i) + \epsilon_i$$

with the most popular choice of g^* being the linear model. In the next section, we will see that many nonlinear models can be reduced to linear models by choosing a suitable basis, and doing so induces interesting block dependent structures. To capture such nonlinear regression models, it suffices to focus on linear models with appropriate dependencies among covariates i.e. we will denote $g^*(\mathbb{X}_i) = \mathcal{X}_i^T \beta$ for a suitable choice of \mathcal{X}_i and β . This would be the main focus of our paper.

We consider a penalized empirical risk minimization problem of the form $\min_{\beta} \hat{R}_n(Y, \mathcal{X}, \beta)$ where

$$\hat{R}_n(Y, \mathcal{X}, \beta) = \frac{1}{2n} \sum_{i=1}^n \|Y_i - \mathcal{X}_i^T \beta\|^2 + \lambda f(\beta). \quad (\text{I.1})$$

Here $\mathcal{X} = (\mathcal{X}_1, \dots, \mathcal{X}_n)^T$, $Y = (Y_1, \dots, Y_n)^T$, $\mathcal{X}_i \in \mathbb{R}^p$, $\beta \in \mathbb{R}^p$ and f is a convex penalty function.

We denote the minimum and the minimizer of the empirical risk by

$$\begin{aligned} \hat{R}_n(Y, \mathcal{X}) &= \min_{\beta \in \mathbb{R}^p} \hat{R}_n(Y, \mathcal{X}, \beta) \\ \hat{\beta}_{\mathcal{X}} &\in \arg \min_{\beta \in \mathbb{R}^p} \hat{R}_n(Y, \mathcal{X}, \beta) \end{aligned} \quad (\text{I.2})$$

respectively. For the rest of this paper we will assume that the covariates are centred i.e. $E[\mathcal{X}_i] = 0$.

Regularized estimators have been a subject of intense study over the decades. Extensive prior work characterizes the risk of such estimators in the proportional asymptotics regime when the covariates are drawn from a Gaussian distribution. To extend to non-Gaussian settings, the typical strategy involves proving universality results that answer the following question: can rows of the data matrix \mathcal{X}_i^T be replaced by Gaussian vectors with matching moments and yield the same asymptotic results?

Among existing universality results, [57] established that when the design matrix has i.i.d. entries and finite sixth moment, the optimal value of the risk of a box constrained Lasso is the same as the optimal value when the entries are i.i.d. Gaussian. Subsequently, [40] showed an analogous result in the context of elastic net. See also [47] for a comprehensive set of universality results that hold when the design entries are i.i.d. Despite this stylized setting, results based on i.i.d. Gaussian designs capture many new high-dimensional phenomena accurately in a qualitative sense. Further, this assumption allows one to characterize precise limits of objects of statistical relevance, such as risks of estimators [4], [8], [9], [18], [19], [20], behaviour of finite-dimensional marginals [9], [21], and prediction error of machine learning algorithms [33], [58], [59], [60], as well as prove \sqrt{n} -consistency of estimators for average treatment effects [34]. In light of the elegant theory

and novel high-dimensional phenomena this allows to capture, it is imperative to understand the limits of these results. In this paper, we seek to understand the extent to which results based on i.i.d. Gaussian designs continue to hold in presence of dependence among the covariates.

In the context of dependent covariates, [48] established that when rows of the design are of the form $\mathcal{X}_i = \Sigma^{1/2} Z_i$ with entries of Z_i i.i.d. sub-Gaussian, the optimal value of the risk matches that of the design whose rows are of the form $\Sigma^{1/2} G_i$, where G_i has i.i.d. Gaussian entries. However, in a broad class of problems, the dependence structure among covariates may not be in the precise form $\Sigma^{1/2} Z_i$ with Z_i containing i.i.d. entries. On the other hand, [45], [46], [61] studied universality under curated forms of dependence, as allowed by right-rotationally-invariant designs, that are critical in compressed sensing applications [62], [63], [64]. See [65] for detailed discussion and examples on the nature of dependent problems this class can handle. These forms of dependence studied in prior work fails to capture general sub-Gaussian designs that may arise naturally in important statistical models—this is the main focus of our paper.

In our conquest to understand dependent covariates, we recognize that “too much” dependence among entries of \mathcal{X}_i may break universality, as demonstrated in [20] and [47]. In the latter, the authors show that for isotropic designs, if each entry in a row of the design depends on every other entry, the estimation risk is asymptotically different from the model with i.i.d. Gaussian covariates. However, noteworthy examples of dependence exist where each entry of the design depends on a subset of the entire row (cf. Section II). A natural question then arises: do universality results hold under such dependence structures? We answer in the affirmative, identifying a class of dependence structures that ensure universality results continue to hold. In particular, we show the following

Let \mathcal{X} be a design matrix with i.i.d. mean zero rows. Further assume that each row \mathcal{X}_i is “block dependent” and sub-Gaussian. Let \mathcal{G} be a design matrix with i.i.d. rows distributed as a multivariate Gaussian $\mathcal{N}(0, \Lambda)$ for a diagonal matrix Λ . If $E(\mathcal{X}_i \mathcal{X}_i^T) = \Lambda$ then we have

$$\begin{aligned} \min_{\beta} \hat{R}_n(Y, \mathcal{X}, \beta) &\approx \min_{\beta} \hat{R}_n(Y, \mathcal{G}, \beta) \\ \|\hat{\beta}_{\mathcal{X}} - \beta_0\|^2 &\approx \|\hat{\beta}_{\mathcal{G}} - \beta_0\|^2. \end{aligned}$$

This simple mathematical result opens the avenue for exact risk characterization of regularized estimators in high-dimensional models that had so far evaded the literature (see Section II for concrete examples). We define the notion of block dependence formally in Section II. To the best of our knowledge, our result is the first in the literature on universality in high-dimensional regression that goes beyond correlation-based dependence or right-rotationally-invariant designs. We provide a more detailed literature review in Section VI. For simplicity, we focus on two of the most commonly used penalty functions—the ℓ_1 (Lasso) or the ℓ_2 (ridge). These differ drastically in the geometry they induce in high dimensions. In this light, we believe that most other popular examples of separable penalties can be handled by extending our current proof techniques.

Finally we note that a crucial step in establishing the universality result, is showing that the optimizers are bounded in ℓ_∞ norm (see Lemma (7.2) for details). To prove such a bound in the case of independent design one usually relies on a leave-one-out method (see [47] for e.g.), which uses high probability upper bound conditioned on a given column. However this method does not work in the dependent case and to circumvent this issue we use a leave-d-out method (see proof of Lemma 7.2), where we condition on an entire block to prove the high probability bounds.

A. Outline of the Results

To illustrate the importance of considering block dependent designs, we present various examples in Section II. The initial example revolves around the commonly used additive model, where we demonstrate that by considering block dependent settings, one can extend results from linear models to nonlinear models. Additionally, we explore an example of scalar-on-function regression where such dependence may emerge. The third example delves into a genomics model where block dependent designs naturally come into play. Section III presents preliminaries and Section IV presents our main results. In Section V, we present simulations corroborating our theory while Section VI reviews literature related to our work. Section VII provides the proof of the main theorems and corollaries. Finally, we conclude in Section VIII with a discussion on possible generalizations of our results, while the proofs of the auxiliary lemmas are given in the Appendix.

B. Notations

We use the $O(\cdot)$ and $o(\cdot)$ to denote asymptotically bounded and asymptotically goes to 0 respectively i.e we say $a_n = O(\gamma_n)$ if $\limsup |a_n/\gamma_n| \leq C$ and $a_n = o(\gamma_n)$ if $\limsup |a_n/\gamma_n| = 0$. Similarly $O_p(\cdot)$ and $o_p(\cdot)$ are reserved for the stochastic versions of the same quantities. We will use \xrightarrow{P} to denote convergence in probability. The ℓ_1 and ℓ_2 norms of vectors are denoted by $\|\cdot\|_1$ and $\|\cdot\|$ respectively unless otherwise mentioned. For matrices the $\|A\|$ will always denote the largest singular value of A (operator norm) and hence the same notation $\|\cdot\|$ will be reserved for the ℓ_2 norm of a vector and the operator of norm of the matrix, the difference being understood from the context. We will denote the set $\{1, \dots, p\}$ by $[p]$. For any subset $S \subset [p]$ and a matrix X we will denote the submatrix whose column indices are restricted to S by X_S . Similarly subvector of a vector v , whose coordinates are restricted to S , will be called v_S . Throughout the notation C will denote a generic constant unless otherwise specified. The Kronecker delta function will be denoted as $\delta_{ij} = \mathbb{1}_{i=j}$. For μ in \mathbb{R}^m ($m > 0$) we will also use δ_μ to define the delta measure which is 1 at μ and 0 otherwise. Since the same notation is used they will appear in separate context and the one referred to will be clear from the context. For (μ_1, \dots, μ_p) with $\mu_i \in \mathbb{R}^m$ ($m, p > 0, 1 \leq i \leq p$) we will define the empirical distribution by $\frac{1}{p} \sum_{i=1}^p \delta_{\mu_i}$. For any sequences of probability measures μ_n and another probability measure μ , we say that $\mu_n \xrightarrow{W_2} \mu$ if the following holds: there exists a sequence of

couplings Π_n with marginals μ_n and μ respectively, so that if $(X_n, X) \sim \Pi_n$, then $E[(X_n - X)^2] \rightarrow 0$ as $n \rightarrow \infty$.

Next we fix some notations for the optimizers and the optimum value of the risk that will help us state our results in a compact form. Let $\varphi = \beta - \beta_0$ and $\hat{\varphi}_\mathcal{X} = \hat{\beta}_\mathcal{X} - \beta_0$ where $\hat{\beta}_\mathcal{X}$ is the minimizer of the empirical risk in equation (I.1). Consider the expression of empirical risk in equation (I.1). For the sake of brevity, we will denote $\hat{R}_n(Y, \mathcal{X}, \beta)$ as $R(\varphi, \mathcal{X})$, where the dependence on β_0 and ξ (and hence Y) is suppressed. Similarly, we will denote the minimum risk as $R(\mathcal{X})$. Further we will denote the general form as $R^K(\varphi, \mathcal{X})$ where $K = L$ or $K = R$ is used to distinguish between the Lasso and ridge setups i.e whether $f(\beta) = \|\beta\|_1/\sqrt{n}$ or $f(\beta) = \|\beta\|_2^2/2$ respectively. The minimum value of the empirical risk in this setup will be denoted as $R^K(\mathcal{X})$ and the error by $\hat{\varphi}_\mathcal{X}^K := \hat{\beta}_\mathcal{X}^K - \beta_0$. Finally we will abuse notation and call both $\hat{\varphi}_\mathcal{X}$ and $\hat{\beta}_\mathcal{X}$ optimizers when there is no chance of confusion and the quantity referred to will be understood from the context.

II. EXAMPLES OF BLOCK DEPENDENT DESIGNS

As discussed in the preceding section, we seek to study models with block dependencies among the design entries. To be precise, we define our formal notion of block dependence below.

Definition 2.1: A p -dimensional random vector W is called block dependent (with dependence parameter d) if there exist a partition $\{S_j\}_{j=1}^{k_p}$ of $[p]$ (i.e. S_j are disjoint and $\cup_{j=1}^{k_p} S_j = [p]$) such that $|S_j| \leq d \forall j$ and $W_l \perp W_k$ if $l \in S_i, k \in S_j, i \neq j$.

The design matrices whose rows are i.i.d. block dependent random vectors will be called block dependent designs. As a warm up toward our results, we present here three concrete settings where block dependent designs arise naturally.

A. Penalized Additive Regression Model

Let $(Y_i, \mathbb{X}_i)_{i=1}^n$ be i.i.d. with $Y_i \in \mathbb{R}$ and $\mathbb{X}_i \in [0, 1]^{p_0}$. We consider the following additive model

$$Y_i = h^*(\mathbb{X}_i) + \epsilon_i,$$

with $h^*(\mathbb{X}_i) = \sum_{j=1}^{p_0} h_j^*(\mathbb{X}_i^j)$ and \mathbb{X}_i^j denoting the j^{th} component of \mathbb{X}_i . This model features in [66], where non-parametric estimation of h^* is considered under an L^2 loss with suitable smoothness restrictions on the function classes. A high-dimensional version of the same problem can be found in [67], where the authors consider penalized estimators of the form $\hat{h} = \sum_{j=1}^{p_0} \hat{h}_j$ that minimizes

$$\frac{1}{2n} \sum_{i=1}^n (Y_i - h(\mathbb{X}_i))^2 + \mathcal{D}(h). \quad (\text{II.1})$$

Here $\mathcal{D}(h)$ penalizes the complexity of the function h and the h_j 's are the components of the additive model. In this context, [67] consider a more general model where each component h_j is allowed to be multivariate. For convenience, we first describe the model under a univariate setup. In [67], the authors consider an ultra-high-dimensional regime where $\log(p_0) = o(n)$ but with sparsity assumptions on the number

of active components h_j (see Section VI for a detailed discussion). We switch gears to a different regime where $p_0/n \rightarrow \kappa$ where κ can be greater than 1. Further we do not assume any sparsity on the underlying model and obtain an exact characterization of the estimation risk. In other words, our work encompasses dense as well as sparse signals. In contrast to prior work where the assumed sparsity is often sub-linear, we are able to allow linear sparse models (see [67] for other forms of sparsity, such as weak sparsity, that takes a slightly different form, and we refrain from comparing it further here).

To make the setting more concrete, we express each component h_j in terms of the trigonometric basis of $L^2[0, 1]$, denoted by $\phi_j(t)$, and given by

$$\begin{aligned}\phi_0(t) &= 1, & \phi_{2k-1}(t) &= \sqrt{2}\cos(2\pi kt) \\ \phi_{2k}(t) &= \sqrt{2}\sin(2\pi kt), & k &= 1, 2, \dots\end{aligned}$$

In nonparametric estimation theory (see [68]) one considers estimation when the functions are restricted to certain smoothness classes with a popular choice being the periodic Sobolev class which is defined as follows

$$\begin{aligned}W(\beta, Q) &= \{f \in L^2[0, 1] : \theta = \{\theta_l\} \in \Theta(\beta, Q), \\ \text{where } \theta_l &= \int_0^1 \phi_l(t) f(t) dt\},\end{aligned}\quad (\text{II.2})$$

where $\Theta(\beta, Q)$ is the Sobolev ellipsoid defined by

$$\Theta(\beta, Q) = \{\theta = \{\theta_l\} \in \ell^2(\mathbb{N}) : \sum_{l=1}^{\infty} \alpha_l \theta_l^2 \leq Q\}, \quad (\text{II.3})$$

with $\alpha_l = l^{2\beta}$ or $(l-1)^{2\beta}$ for even and odd l respectively.

Let h_j be a function of the form

$$h_j(t) = \sum_{k=1}^{\infty} \theta_k^j \phi_k(t). \quad (\text{II.4})$$

Then observing that θ_k^j are Fourier coefficients one can consider estimation of h_j when $h_j \in W(\beta, Q)$ or equivalently $\{\theta_k^j\}_{k=1}^{\infty} \in \Theta(\beta, Q)$. Thus the estimation of the coefficients θ_k^j is equivalent to the estimation of the function h_j and the number of non-zero coefficients is a measure of complexity of the function.

In this paper we additionally impose the condition that h_j is a finite linear combination of the basis functions i.e. it is a function of the form

$$h_j(t) = \sum_{k=1}^d \theta_k^j \phi_k(t). \quad (\text{II.5})$$

Indeed such finite linear combinations can be interpreted as a subclass within the periodic Sobolev class with finitely many Fourier coefficients and has been widely studied, most notably in [69]. Using (II.1) and (II.5), we arrive at the penalized estimation problem

$$\frac{1}{2n} \sum_{i=1}^n (Y_i - \sum_{j=1}^{p_0} \sum_{k=1}^d \theta_k^j \phi_k(\mathbb{X}_i^j))^2 + \lambda D(\theta),$$

where $D(\theta)$ penalizes the complexity of the function h in terms of the parameter θ .

We observe the following facts. Since one can always center Y , we may drop the intercept term arising from $\phi_0(t)$. Further, if we assume $\mathbb{X}_i^j \sim \text{Unif}(0, 1)$, then using the fact that $\{\phi_k\}$ is an orthonormal basis, we have $E[\phi_k(\mathbb{X}_i^j)] = 0$ and $E[\phi_k(\mathbb{X}_i^j)\phi_l(\mathbb{X}_i^j)] = \delta_{kl}$. Furthermore we assume \mathbb{X}_i^j s to be independent across j . Thus we can reformulate the above problem to estimation of β in the following linear regression problem

$$Y = \mathcal{X}\beta + \xi, \quad (\text{II.6})$$

where the rows of \mathcal{X} (a $n \times p_0 d$ matrix) are of the form

$$\begin{aligned}\mathcal{X}_i = & (\phi_1(\mathbb{X}_i^1), \dots, \phi_d(\mathbb{X}_i^1), \dots, \phi_1(\mathbb{X}_i^{p_0}), \dots \\ & \dots, \phi_d(\mathbb{X}_i^{p_0}))^T\end{aligned}$$

and $\beta = (\theta_k^j)_{j=1, k=1}^{p_0, d}$. Thus the rows are independent isotropic sub-Gaussian vectors with a block dependent structure with each block having size d . We note that letting $p = p_0 d$ we may capture properties of estimators in this model by developing theory under the setup where $p_0 d/n \rightarrow \kappa$ (see Section III-A).

Further note that for any estimator $\hat{h}_j(x)$ of the form $\sum_{k=1}^d \hat{\theta}_k^j \phi_k(x^j)$, the L_2 estimation risk translates to

$$\|\hat{h}_j - h_j\|^2 := \int_0^1 (\hat{h}_j(t) - h_j(t))^2 dt = \sum_{k=1}^d |\hat{\theta}_k^j - \theta_k^j|^2.$$

Thus we have

$$\sum_{j=1}^{p_0} \|h_j - \hat{h}_j\|^2 = \sum_{j=1}^{p_0} \sum_{k=1}^d |\hat{\theta}_k^j - \theta_k^j|^2 \quad (\text{II.7})$$

and it suffices to analyze the estimation risk $\|\hat{\beta} - \beta\|_2^2$ in the model (II.6).

B. Penalized Functional Regression

Let $(Y_i, \mathbb{X}_i(t), \epsilon_i)_{i=1}^n$ be i.i.d. with $Y_i \in \mathbb{R}$ and $\mathbb{X}_i(t) \in \mathbb{R}$ for $t \in (0, 1)$. We consider the following functional regression model

$$Y_i = \int_0^1 \mathbb{X}_i(t) \gamma(t) + \epsilon_i, \quad (\text{II.8})$$

with $\gamma(t) \in \mathbb{R}$ and ϵ_i being sub-Gaussian random variables.

This model and its variants, called functional linear regression or scalar-on-function regression, has been widely studied in functional data analysis [70], [71], [72]. The Kosambi–Karhunen–Loëve theorem states that under standard assumptions the stochastic process $\mathbb{X}_i(t)$ admits an expansion with respect to an orthonormal basis i.e

$$\mathbb{X}_i(t) = \sum_{k=1}^{\infty} \zeta_{ik} \phi_k(t),$$

where the random variables ζ_{ik} s have mean 0 and $E[\zeta_{il} \zeta_{ik}] = \lambda_k \delta_{kl}$ i.e. they are uncorrelated with variance λ_k . When $\mathbb{X}_i(t)$ is expanded upto a finite number of terms as is used in practice [70], the regression model can be written as

$$Y_i = \sum_{k=1}^d \zeta_{ik} \theta_k + \epsilon_i,$$

where we have also expanded the function $\gamma(t)$ as $\gamma(t) = \sum_{k=1}^d \theta_k \phi_k(t)$. In the multivariate setup the data is of the form $(Y_i, \mathbb{X}_i(t), \epsilon_i)_{i=1}^n$ where $Y_i \in \mathbb{R}$ and $\mathbb{X}_i(t) \in \mathbb{R}^p$ for $t \in (0, 1)$. We denote the coordinates of the multivariate function $\mathbb{X}_i(t)$ as $\mathbb{X}_i^j(t)$ and we assume the coordinate functions to be independent. We have the following expansion

$$\mathbb{X}_i^j(t) = \sum_{k=1}^d \zeta_{ik}^j \phi_k(t).$$

The random variables ζ_{ik}^j satisfies the relations $E[\zeta_{ik}^j] = 0$ and $E[\zeta_{ik}^j \zeta_{il}^j] = \lambda_k^j \delta_{kl}$. We consider the multivariate analogue of the model (II.8)

$$Y_i = \sum_{j=1}^{p_0} \int_0^1 \mathbb{X}_i^j(t) \gamma^j(t) + \epsilon_i.$$

Finally letting $\gamma^j(t) = \sum_{k=1}^d \theta_k^j \phi_k(t)$ we obtain

$$Y_i = \sum_{j=1}^{p_0} \sum_{k=1}^d \zeta_{ik}^j \theta_k^j + \epsilon_i. \quad (\text{II.9})$$

We note that since $\mathbb{X}_i^1(t), \dots, \mathbb{X}_i^{p_0}(t)$ are independent stochastic processes, (II.9) can also be written in the form $Y = \mathcal{X}\beta + \xi$ with \mathcal{X}_i of the form

$$\mathcal{X}_i = (\zeta_{i1}^1, \dots, \zeta_{id}^1, \dots, \zeta_{i1}^{p_0}, \dots, \zeta_{id}^{p_0})^T$$

and $\beta = (\theta_k^j)_{j=1, k=1}^{p_0, d}$. The design matrix \mathcal{X} has independent rows whose entries have the block dependent structure as given in Definition 2.1. Indeed for each i , $\mathbb{X}_i^j(t)$ and $\mathbb{X}_i^m(t)$ being independent for $m \neq j$, ζ_{ik}^j and ζ_{il}^m are independent as well. Further we have $\mathcal{X} = X\Lambda^{1/2}$ where X_{ij} s have variance 1. We will assume that the vector $(\zeta_{ik}^j)_{k=1}^d$ is sub-Gaussian random vector (II.9) and that $E[|\zeta_{ik}^j|^4] \leq \mathcal{K}$ for some constant \mathcal{K} . In particular if ζ_{ik}^j are bounded and d remains constant the above assumptions trivially holds. Also, note that the matrix Λ is a diagonal matrix and we will assume that the diagonal entries λ_i satisfy $c \leq \lambda_i \leq C$ for some $c > 0$ for all $i \in \{1, \dots, p_0\}$.

A common approach towards estimating the vector β is to consider the following penalized estimation problem (see [72] for example)

$$\hat{\beta} = \arg \min_{\beta} \frac{1}{2n} \|Y - \mathcal{X}\beta\|^2 + f(\beta)$$

and then construct $\hat{\gamma}^j(t) = \sum_{k=1}^d \hat{\theta}_k^j \phi_k(t)$. As in the case of Example II.1 we have

$$\sum_{j=1}^{p_0} \|\gamma_j - \hat{\gamma}_j\|^2 = \sum_{j=1}^{p_0} \sum_{k=1}^d |\hat{\theta}_k^j - \theta_k^j|^2, \quad (\text{II.10})$$

and we are reduced to the familiar setup of finding an asymptotic expression for $\|\hat{\beta} - \beta\|_2^2$ in the model (II.6).

Remark 2.1: We note that while the concept that nonlinear models can be studied as linear models via basis expansion is widely studied in the literature, the structure of dependence arising from such models is the focus of our study. Unless Gaussian assumptions are used (for example assuming $\mathbb{X}(t)$ is a Gaussian process) the design matrix contains rows with

dependent entries. In the first and second model, we may assume that d is finite and $n/(p_0 d) \rightarrow \kappa^{-1}$. This assumption can be relaxed to the assumption $c < n/p_0 < C$ for some constants c and C . In the more general setup we can allow the block sizes to vary i.e. there may be p_0 blocks with the j^{th} block having size d_j and $cn < \sum_{j=1}^{p_0} d_j < Cn$.

C. Block Dependent Covariates in GWAS

In Genome Wide Association Studies (GWAS) one often studies the association between phenotypes (Y_i 's) and genetic markers/single nucleotide polymorphism (SNP) \mathbb{X}_{ij} where i denotes the i^{th} individual and $j \in \{1, \dots, p\}$. The markers \mathbb{X}_{ij} typically follow a dependence structure where a particular SNP is dependent on SNPs of nearby loci, a phenomenon called *linkage disequilibrium*. This phenomenon is taken into account in the detection of causal SNPs and we refer the readers to [73] and the references therein for a detailed account.

While most measures of linkage disequilibrium is based on correlation (see [56]), other measures of association has been considered as well. For example, [74] describes a measure of association in terms of Kullback-Leibler distance between distributions. The same reference also points out that situations of dependence may arise where the associated variables may be uncorrelated despite being dependent and that correlation is rarely useful beyond the case of Gaussian designs. Furthermore in GWAS, the markers \mathbb{X}_{ij} typically take on three different discrete values, encoding allele frequencies, and therefore this application rarely encounters Gaussian designs. These observations are perfectly in tandem with the general model that we consider in this paper. To elucidate the occurrence of uncorrelated but dependent random variables with discrete support we consider the following toy example:

\mathbb{X}	A	B	C
a_1	1	0	0
a_2	-1	0	0
a_3	0	1	0
a_4	0	-1	0
a_5	0	0	1
a_6	0	0	-1

Let \mathbb{X} be a random variable taking six distinct values a_1, \dots, a_6 with equal probabilities. Further let the three random variables take values according to the entries of the i^{th} row i.e. when $\mathbb{X} = a_1$ we have $A = 1, B = 0, C = 0$, and so forth. It can be easily verified that the marginal distribution of A, B, C are same and they have mean 0. Further their covariance matrix is cI_3 for some constant c , so that the vector (A, B, C) can be easily scaled to obtain an isotropic random vector. Further because of the underlying generation mechanism, they are dependent. We note that this can be easily generalized to d random variables from a table with $2d$ rows.

On the other hand, the other aspect of the dependence structure arising in GWAS is the concept of neighborhood dependence that tells us that a particular SNP depends only on its neighbors. Of particular importance is the block structure considered in [75], where there is linkage disequilibrium if and only if two SNPs belong to the same block. Our

block dependence structure in this paper thus encompasses this critical GWAS example. In Subsection V-C, we provide simulations on an example where the SNPs are generated from a Hidden Markov model, which has been extensively utilized within the GWAS community (cf. [76], [77]) to simulate and impute genomic data.

After discussing our motivation and relevant examples, we now proceed to present the preliminaries and assumptions necessary for our main results.

III. PRELIMINARIES

A. Assumptions

Throughout we assume that we observe n i.i.d. samples $\{Y_i, \mathcal{X}_i, 1 \leq i \leq n\}$. We will work in a high-dimensional setup where the covariate dimension p is allowed to grow with the sample size n . To make this rigorous, we consider a sequence of problem instances $\{Y_{i,n}, \mathcal{X}_{i,n}, \xi_{i,n}, 1 \leq i \leq n, \beta_0(n), \beta(n)\}_{n \geq 1}$ and work under the setting of a linear model

$$Y_{i,n} = \mathcal{X}_{i,n}^T \beta_0(n) + \xi_{i,n},$$

with $Y_{i,n} \in \mathbb{R}$, $\mathcal{X}_{i,n} \in \mathbb{R}^{p(n)}$, and $\beta_0(n) \in \mathbb{R}^{p(n)}$. We allow $p(n), n \rightarrow \infty$ with $p(n)/n \rightarrow \kappa \in (0, \infty)$ and in the sequel we will drop the dependence on n whenever it is clear from the context. Further we let ξ_i s be independent uniformly sub-Gaussian random variables with variance σ^2 . Note they do not need to arise from the same distribution. In a more compact form the above model will be written as

$$Y = \mathcal{X}\beta_0 + \xi, \quad (\text{III.1})$$

where $\mathcal{X} \in \mathbb{R}^{n \times p}$, $Y, \xi \in \mathbb{R}^n$.

Our assumptions on the design \mathcal{X} , the signal β_0 , and the tuning parameter λ are as follows.

Assumption 1 (Assumptions on the Design): The rows of the design \mathcal{X} are i.i.d. random vectors each of which is block dependent with dependence parameter d . We assume that $\mathcal{X} = X\Lambda^{1/2}$ where the rows of X are centered and isotropic i.e. if X_i denotes the i -th row of X then $E[X_i] = 0$, $E[X_{ij}X_{ik}] = \delta_{jk}$. Further we assume that the entries have bounded fourth moment ($E|X_{ij}|^4 \leq \mathcal{K}$) and the row vector X_i is a sub-Gaussian random vector. Furthermore, we assume that Λ is a diagonal matrix with the diagonal entries λ_i satisfying $c \leq \lambda_i \leq C$ for some positive constants c, C .

As mentioned in the remark at the end of Subsection II-B, the dependence occurs in blocks and the block sizes should be bounded by the dependence parameter d . In Section VII, we establish that our universality holds as long as $d = o\left(\left(\frac{n}{\log n}\right)^{1/5}\right)$. Regarding the uncorrelation assumption, we note that the proofs rely on establishing the ℓ_∞ and ℓ_2 norm bounds on the optimizers (see the proofs of Lemma VII.1. and VII.2.), which in turn is highly dependent on the structure of the penalty as well as the uncorrelation assumption in Assumption 1. In particular, adding correlation renders the penalty inseparable in which case our leave-d-out method fails. Further, we demonstrate universality of the optimizers using the convex Gaussian minmax theorem (CGMT) and subsequently analyzing a scalar optimization problem (see proof

of Lemma VII.3). Analysis of the same in the presence of correlation, if possible, would necessitate further assumptions and entirely different proof techniques, which is outside the scope of this paper. Finally, the assumption $E[X_{ij}^4] \leq \mathcal{K}$ can be weakened to an assumption involving d, n and the moments, but for simplicity, we will state our theorems under this simplified assumption.

Assumption 2: The signal satisfies the bound $\beta_0^T \Lambda \beta_0 \leq C$. Let $v_i = \sqrt{n\lambda_i} \beta_{0,i}$ for $1 \leq i \leq p$. Then we assume that

$$\frac{1}{p} \sum_{i=1}^p \delta_{\lambda_i, v_i} \xrightarrow{W_2} \mu,$$

for some measure μ supported on $\mathbb{R}^{\geq 0} \times \mathbb{R}$. Further we assume that $\|\mu\|_\infty \leq C$.

Assumption 3: The tuning parameter λ satisfies $\lambda \geq \lambda_0$ for a constant $\lambda_0 > 0$.

We remark that we require this assumption only for the case of the Lasso and is the analog of restricted strong convexity type assumptions in the Lasso literature (see Section VII-A for further details).

B. Choice of Penalty Function

As will be shown in Section VII, universality of both the empirical risk and the estimation risk depends on establishing apriori that the optimizers are guaranteed to lie in “nice” sets. However establishing this property requires a case-by-case analysis. In this paper, we consider the two most popular choices of penalty used in the statistics literature—the Lasso and the ridge—that can form prototypes for many other popular examples, the ridge serving as a prototype for strong convex penalties and the Lasso serving as a prototype in the absence of strong convexity. Since we are able to handle these prototypes, we believe our method can be extended to handle other separable penalties.

Remark 3.1: Some comments on the scaling of the penalties are also in order. In [47] the entries of \mathcal{X} are $O_p(n^{-1/2})$ and $f(\beta) = \|\beta\|_1/n$ or $f(\beta) = \|\beta\|_2^2/2n$ respectively. However in our case, the entries of \mathcal{X} are $O_p(1)$. Writing $A = \mathcal{X}/\sqrt{n}$ and $\mu = \sqrt{n}\beta$, equation (I.1) can be rewritten in the form

$$\hat{R}_n(y, A, \mu) = \frac{1}{2n} \sum_{i=1}^n \|Y_i - A_i^T \mu\|^2 + \lambda \bar{f}(\mu), \quad (\text{III.2})$$

where $\bar{f}(\mu) = \|\mu\|_1/n$ or $\bar{f}(\mu) = \|\mu\|_2^2/2n$ and we recover the setup in [47]. We have avoided the $n^{-1/2}$ scaling in the design matrix since it helps to state the assumptions in a concise fashion and to borrow certain results from the random matrix literature. In our setup we end up with the non-standard scaling of the penalty i.e. $f(\beta) = \|\beta\|_1/\sqrt{n}$ or $f(\beta) = \|\beta\|_2^2/2$ respectively.

C. Fixed Point Equations for the Estimation Risk

Finally, our universality result states that the estimation risk of regularized estimators under our setting is asymptotically the same as that in a setting where the design has i.i.d. Gaussian entries. Extensive prior work characterizes the risk of

convex regularized estimators in the Gaussian case [3], [8], [18]. In this characterization, a certain system of fixed point equations plays a critical role. We review this system below for convenience.

Let

$$\eta_k(z, \lambda) := \arg \min_{x \in \mathbb{R}^p} \left\{ \frac{1}{2} \|z - x\|^2 + \lambda \frac{\|x\|_k^k}{k} \right\}.$$

We will focus on two particular values of k in the above expression of $\eta_k(\cdot, \cdot)$; k be taken as 2 when $K = R$ and 1 when $K = L$ (i.e. ridge and Lasso respectively).

Now let $\mu_{0,i} = \sqrt{n\lambda_i}\beta_{0,i}$ and $\omega_i = \lambda_i^{-1/2}$. Also define Q_n as the law of the 3-tuple of random variables $(M, \Omega, Z) \sim (\frac{1}{p} \sum_{i=1}^p \delta_{\mu_{0,i}, \omega_i}) \otimes N(0, 1)$. When the entries of the design matrix are i.i.d. Gaussian, using convex Gaussian minmax theorem (see Appendix for details) the optimum value in (I.2) can be well approximated by $\min_{\beta} \max_{\gamma} \psi^K(\beta, \gamma)$ where the function $\psi^K(\beta, \gamma)$ is defined as follows:

$$\begin{aligned} \psi^K(\beta, \gamma) = & \left(\frac{\sigma^2}{\gamma} + \gamma \right) \frac{\beta}{2} - \frac{\beta^2}{2} + \frac{1}{n/p} E_{Q_n} \min_w \left\{ \frac{w^2 \beta}{2\gamma} \right. \\ & \left. - \beta Z w + \frac{\lambda}{k} (|\Omega(w + \mu_0)|^k - |\mu_0|^k) \right\}. \end{aligned} \quad (\text{III.3})$$

The saddle point of the above equation is given by the fixed point of the following system of equations:

$$\begin{aligned} (\gamma_*^K)^2 = & \sigma^2 + \frac{p}{n} E_{Q_n} \left[\eta_k \left(M + \gamma_*^K Z, \frac{\gamma_*^K \lambda \Omega^k}{\beta_*^K} \right) \right. \\ & \left. - M \right]^2, \\ \beta_*^K = & - \frac{\gamma_*^K p}{n} E_{Q_n} \left[\eta'_k \left(M + \gamma_*^K Z, \frac{\gamma_*^K \lambda \Omega^k}{\beta_*^R} \right) \right] \\ & + \gamma_*^K. \end{aligned} \quad (\text{III.4})$$

We note that under Assumptions 1 and 2 there will be a unique solution of the fixed point equations (see proof of Lemma 7.3 for details). Consistent with our notation, (β_*^R, γ_*^R) will be the solution of the fixed point equations (III.4) for ridge and (β_*^L, γ_*^L) will be the corresponding solution for the Lasso. Next we define ‘‘error’’ terms which are functions of the solutions of the above equation system. Let $g \sim N(0, I_p)$. Define the vectors φ_*^R and φ_*^L for the ridge and Lasso problems in terms of their i^{th} entries as follows

$$\varphi_{*,i}^R = \eta_2 \left(\beta_{0,i} + \frac{\gamma_*^R \omega_i g_i}{\sqrt{n}}, \frac{\gamma_*^R \lambda \omega_i^2}{\beta_*^R} \right) - \beta_{0,i}, \quad (\text{III.5})$$

$$\varphi_{*,i}^L = \eta_1 \left(\beta_{0,i} + \frac{\gamma_*^L \omega_i g_i}{\sqrt{n}}, \frac{\gamma_*^L \lambda \omega_i^2}{\sqrt{n} \beta_*^L} \right) - \beta_{0,i}, \quad (\text{III.6})$$

for $i \in \{1, \dots, p\}$. We note that (III.5) and (III.6) describes the asymptotic behavior of the quantities $\hat{\beta}_{\mathcal{G},i}^R - \beta_{0,i}$ and $\hat{\beta}_{\mathcal{G},i}^L - \beta_{0,i}$ respectively and will play a crucial role in the asymptotic characterization of the estimation risk. We are now ready to state the main results of our paper.

IV. MAIN RESULT

In this section we state the main results of our paper. We consider the linear model (III.1) and seek to study universality properties of the optimum value of the empirical risk

and the estimation risk of the optimizer. In particular, we aim to establish that

$$\begin{aligned} \min_{\beta} \hat{R}_n(Y, \mathcal{X}, \beta) & \approx \min_{\beta} \hat{R}_n(Y, \mathcal{G}, \beta) \\ \|\hat{\beta}_{\mathcal{X}} - \beta_0\|^2 & \approx \|\hat{\beta}_{\mathcal{G}} - \beta_0\|^2, \end{aligned} \quad (\text{IV.1})$$

where $\mathcal{G} = G\Lambda^{1/2}$, G a matrix with i.i.d. standard Gaussian entries.

Recall that we have defined the shifted versions of the optimizers as $\hat{\varphi}_{\mathcal{X}}^K := \hat{\beta}_{\mathcal{X}}^K - \beta_0$. Before we state our main theorems, we describe a set restricted to which our universality result holds.

$$\mathcal{T}_n := \{\varphi : \|\varphi\|_{\infty} \leq C \sqrt{\frac{d \log n}{n}}, \|\varphi\| \leq C\}.$$

Theorem 4.1: Suppose the design matrix $\mathcal{X} = X\Lambda^{1/2}$ satisfies Assumption 1 and β_0 satisfies Assumption 2. Further $\mathcal{G} = G\Lambda^{1/2}$, where $G \in \mathbb{R}^{n \times p}$ contains i.i.d. standard Gaussian entries, and that ψ is a bounded differentiable function with bounded Lipschitz derivative. Then $\exists \gamma_n = o(1)$ such that for any subset $\mathcal{S}_n \subset \mathbb{R}^p$, we have

$$\begin{aligned} |E(\psi(\min_{\varphi \in \mathcal{S}_n \cap \mathcal{T}_n} R^R(\varphi, \mathcal{X}))) - E(\psi(\min_{\varphi \in \mathcal{S}_n \cap \mathcal{T}_n} R^R(\varphi, \mathcal{G})))| \\ \leq \gamma_n \end{aligned}$$

Further if we assume Assumption 3 in addition to Assumptions 1-2, then we have

$$\begin{aligned} |E(\psi(\min_{\varphi \in \mathcal{S}_n \cap \mathcal{T}_n} R^L(\varphi, \mathcal{X}))) - E(\psi(\min_{\varphi \in \mathcal{S}_n \cap \mathcal{T}_n} R^L(\varphi, \mathcal{G})))| \\ \leq \gamma_n. \end{aligned}$$

In the next step we show that the Lasso and ridge optimizers denoted by $\hat{\varphi}_{\mathcal{X}}^L$ and $\hat{\varphi}_{\mathcal{X}}^R$ lie in the subset \mathcal{T}_n with high probability. Thus the restriction over \mathcal{T}_n can be easily lifted and we have the following result.

Corollary 4.1: Assume the conditions of Theorem 4.1. Then we have

$$|E(\psi(\min_{\varphi} R^K(\varphi, \mathcal{X}))) - E(\psi(\min_{\varphi} R^K(\varphi, \mathcal{G})))| \leq \gamma_n,$$

for some $\gamma_n = o(1)$. In particular if $R^K(\mathcal{G}) \rightarrow \rho$, then $R^K(\mathcal{X}) \rightarrow \rho$ as well.

Next we turn to characterizing the estimation risk of the optimizers. The difficulty and the methods used to prove this result is thoroughly discussed in Section VII. Finally, we have our main result.

Theorem 4.2: Assume the conditions in Theorem 4.1. For all Lipschitz functions $\varrho : \mathbb{R}^p \rightarrow \mathbb{R}$ and $\epsilon > 0$ we have

$$P(|\varrho(\varphi_{\mathcal{X}}^K) - \varrho(\hat{\varphi}_{\mathcal{G}}^K)| \geq \epsilon) \leq \varepsilon_n, \quad (\text{IV.2})$$

for some $\varepsilon_n \rightarrow 0$.

We will establish in the proof of Theorem 4.2 that both $\varrho(\hat{\varphi}_{\mathcal{X}}^K)$ and $\varrho(\hat{\varphi}_{\mathcal{G}}^K)$ concentrate around the common value $E[\varrho(\varphi_*^K)]$ where φ_*^R and φ_*^L are defined in equations (III.5) and (III.6) respectively. Thus by choosing a suitable ϱ , the above theorem can naturally be specialized to yield the estimation risk of the ridge and the Lasso in our dependent covariates

setup. We obtain the following corollary for dependent sub-Gaussian designs.

Corollary 4.2: Let γ_*^K and β_*^K be fixed points of the equation system (III.4). In the setup of Theorem 4.1, we have

$$\begin{aligned} \|\hat{\beta}_{\mathcal{X}}^R - \beta_0\|^2 - E_{\bar{Q}_n} \left[\eta_2 \left(M + \frac{\gamma_*^R \Omega Z}{\sqrt{n}}, \frac{\gamma_*^R \lambda \Omega^2}{\beta_*^R} \right) - M \right]^2 \\ \xrightarrow{P} 0, \\ \|\hat{\beta}_{\mathcal{X}}^L - \beta_0\|^2 - E_{\bar{Q}_n} \left[\eta_1 \left(M + \frac{\gamma_*^L \Omega Z}{\sqrt{n}}, \frac{\gamma_*^L \lambda \Omega^2}{\sqrt{n} \beta_*^L} \right) - M \right]^2 \\ \xrightarrow{P} 0. \end{aligned}$$

Here the expectations are taken with respect to the measure \bar{Q}_n which is the law of the 3-tuple of random variables $(M, \Omega, Z) \sim (\frac{1}{p} \sum_{i=1}^p \delta_{\beta_0, i, \omega_i}) \otimes N(0, 1)$ where $\omega_i = \lambda_i^{-1/2}$. Note that the equivalent results for estimation risk in the independent Gaussian design have been obtained in [3] and [20] for ridge and Lasso respectively and thus our result generalizes the exact risk characterization to the dependent sub-Gaussian case. We are now in position to compute the estimation risks of the Lasso and the ridge in the penalized regression models described in Section II. In the following corollaries, we demonstrate these for the two nonlinear regression models we described.

Corollary 4.3: Let

$$Y_i = \sum_{j=1}^{p_0} h_j^*(\mathbb{X}_i^j) + \epsilon_i, \quad i \in \{1, \dots, n\},$$

where \mathbb{X}_i^j are i.i.d. $\text{Unif}(0,1)$, and $h_j^*(t) = \sum_{k=1}^d \theta_{0k}^j \phi_k(t)$. Define $\theta = \{\theta_k^j\}_{j=1, k=1}^{p_0, d}$ and $\beta_0 = \{\theta_{0,k}^j\}_{j=1, k=1}^{p_0, d}$ and let β_0 satisfy Assumption 1 with $\Lambda = I_{p_0 d}$. Further define the optimizers as follows

$$\begin{aligned} \hat{\theta}^L = \arg \min_{\theta} \frac{1}{2n} \sum_{i=1}^n (Y_i - \sum_{j=1}^{p_0} \sum_{k=1}^d \theta_k^j \phi_k(\mathbb{X}_i^j))^2 \\ + \frac{\lambda}{\sqrt{n}} \sum_{j=1}^{p_0} \sum_{k=1}^d |\theta_k^j|, \end{aligned}$$

$$\begin{aligned} \hat{\theta}^R = \arg \min_{\theta} \frac{1}{2n} \sum_{i=1}^n (Y_i - \sum_{j=1}^{p_0} \sum_{k=1}^d \theta_k^j \phi_k(\mathbb{X}_i^j))^2 \\ + \frac{\lambda}{2} \sum_{j=1}^{p_0} \sum_{k=1}^d |\theta_k^j|^2. \end{aligned}$$

Then denoting $\hat{h}_j^L(t) = \sum_{k=1}^d \hat{\theta}_k^j \phi_k(t)$ and $\hat{h}_j^R(t) = \sum_{k=1}^d \hat{\theta}_k^j \phi_k(t)$ we have,

$$\begin{aligned} \sum_{j=1}^{p_0} \|\hat{h}_j^R - h_j^*\|^2 - E_{\bar{Q}_n} \left[\eta_2 \left(M + \frac{\gamma_*^R Z}{\sqrt{n}}, \frac{\gamma_*^R \lambda}{\beta_*^R} \right) - M \right]^2 \\ \xrightarrow{P} 0. \end{aligned} \quad (\text{IV.3})$$

Further if Assumption 3 is satisfied we have

$$\sum_{j=1}^{p_0} \|\hat{h}_j^L - h_j^*\|^2 - E_{\bar{Q}_n} \left[\eta_2 \left(M + \frac{\gamma_*^L Z}{\sqrt{n}}, \frac{\gamma_*^L \lambda}{\sqrt{n} \beta_*^L} \right) - M \right]^2$$

$$\xrightarrow{P} 0. \quad (\text{IV.4})$$

In the above two statements the expectations are taken with respect to the measure \bar{Q}_n which is the law of the 2-tuple of random variables $(M, Z) \sim (\frac{1}{p} \sum_{i=1}^p \delta_{\beta_0, i, \omega_i}) \otimes N(0, 1)$.

Corollary 4.4: Let

$$Y_i = \sum_{j=1}^{p_0} \int_0^1 \mathbb{X}_i^j(t) \gamma_0^j(t) + \epsilon_i, \quad i \in \{1, \dots, n\},$$

where $\mathbb{X}_i^j(t)$ s are independent, $\mathbb{X}_i^j(t) = \sum_{k=1}^d \zeta_{ik}^j \phi_k(t)$, and $\gamma_0^j(t) = \sum_{k=1}^d \theta_{0k}^j \phi_k(t)$. Define $\theta = \{\theta_k^j\}_{j=1, k=1}^{p_0, d}$, $\beta_0 = \{\theta_{0,k}^j\}_{j=1, k=1}^{p_0, d}$, $\text{Var}(\zeta_{ik}^j) = \lambda_{(j-1)d+k}$, and $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_{p_0 d})$. Let the vectors $\zeta_i = (\zeta_{ik}^j)_{j=1, k=1}^{p_0, d}$ satisfy Assumption 1. Further let β_0 satisfy Assumption 2. Define the optimizers as follows

$$\begin{aligned} \hat{\theta}^L = \arg \min_{\theta} \frac{1}{2n} \sum_{i=1}^n (Y_i - \sum_{j=1}^{p_0} \sum_{k=1}^d \theta_k^j \zeta_{ik}^j)^2 \\ + \frac{\lambda}{\sqrt{n}} \sum_{j=1}^{p_0} \sum_{k=1}^d |\theta_k^j|, \\ \hat{\theta}^R = \arg \min_{\theta} \frac{1}{2n} \sum_{i=1}^n (Y_i - \sum_{j=1}^{p_0} \sum_{k=1}^d \theta_k^j \zeta_{ik}^j)^2 \\ + \frac{\lambda}{2} \sum_{j=1}^{p_0} \sum_{k=1}^d |\theta_k^j|^2. \end{aligned}$$

Then denoting $\hat{\gamma}_j^L(t) = \sum_{k=1}^d \hat{\theta}_k^j \phi_k(t)$ and $\hat{\gamma}_j^R(t) = \sum_{k=1}^d \hat{\theta}_k^j \phi_k(t)$ we have

$$\begin{aligned} \sum_{j=1}^{p_0} \|\hat{\gamma}_j^R - \gamma_j^*\|^2 - E_{\bar{Q}_n} \left[\eta_2 \left(M + \frac{\gamma_*^R \Omega Z}{\sqrt{n}}, \frac{\gamma_*^R \lambda \Omega^2}{\beta_*^R} \right) \right. \\ \left. - M \right]^2 \xrightarrow{P} 0. \end{aligned} \quad (\text{IV.5})$$

Further if Assumption 3 is satisfied we have

$$\begin{aligned} \sum_{j=1}^{p_0} \|\hat{\gamma}_j^L - \gamma_j^*\|^2 - E_{\bar{Q}_n} \left[\eta_1 \left(M + \frac{\gamma_*^L \Omega Z}{\sqrt{n}}, \frac{\gamma_*^L \lambda \Omega^2}{\sqrt{n} \beta_*^L} \right) \right. \\ \left. - M \right]^2 \xrightarrow{P} 0. \end{aligned} \quad (\text{IV.6})$$

Here the expectations are taken with respect to the measure \bar{Q}_n which is the law of the 3-tuple of random variables $(M, \Omega, Z) \sim (\frac{1}{p} \sum_{i=1}^p \delta_{\beta_0, i, \omega_i}) \otimes N(0, 1)$ where $\omega_i = \lambda_i^{-1/2}$.

Remark 4.1: Corollaries 4.3 and 4.4 follows directly by applying Corollary 4.2 to the additive and functional regression models respectively. The simpler form of the risks in (IV.3) and (IV.4) in comparison to (IV.5) and (IV.6) is due to the fact the the rows in the former setup are isotropic random vectors and hence $\Lambda = \tilde{\Omega} = I_{p_0 d}$. Similarly, in this simplified setup the measure \bar{Q}_n appearing Corollary 4.2 simplifies to \bar{Q}_n in (IV.3) and (IV.4).

V. ILLUSTRATION

In this section, we demonstrate the efficacy of our result via experiments. In particular, we consider the first two settings in Section II and compare the risk of regularized estimators in each model with that under i.i.d. Gaussian designs. Across our simulations, we observe that universality kicks in for quite moderate sample sizes, a reassuring observation that suggests our current asymptotic abstraction captures reality reasonably well. As in Section II we will denote $p = p_0 d$, i.e. there will be p_0 blocks each having size d .

A. Penalized Additive Regression Model

Recall the setup given in Example II.1,

$$Y_i = \sum_{j=1}^{p_0} h_j^*(\mathbb{X}_i^j) + \epsilon_i, \quad h_j(t) = \sum_{k=1}^d \theta_k^j \phi_k(t)$$

$$\phi_{2l-1}(t) = \sqrt{2} \cos(2\pi l t), \quad \phi_{2l}(t) = \sqrt{2} \sin(2\pi l t)$$

for $l = 1, \dots, d/2$.

We set $n = 200, p_0 = 30, d = 10$ so that $(p_0 d)/n = 1.5$. For the i^{th} row of the design matrix, we use the following strategy: Generate i.i.d. $\{\mathbb{X}_i^j\}_{j=1}^{30}$ such that $\mathbb{X}_i^j \sim \text{Unif}(0, 1)$ and construct the following

$$\mathcal{X}_i = (\phi_1(\mathbb{X}_i^1), \dots, \phi_{10}(\mathbb{X}_i^1), \dots, \phi_1(\mathbb{X}_i^{30}), \dots, \phi_{10}(\mathbb{X}_i^{30}))^T.$$

We repeat this construction for $n = 200$ rows to create the design matrix \mathcal{X} and for the Gaussian problem, we construct the 200×300 matrix \mathcal{G} with i.i.d. Gaussian entries. We generate the error entries $\epsilon_i, \bar{\epsilon}_i$ so that they are i.i.d. standard normal. Finally, we generate the signal $\beta_0 = \{\theta_{0,k}^j\}_{j=1, k=1}^{p_0, d}$ as $\beta_{0,j} = n^{-1/2} b_j$, where the b_j 's are i.i.d. Bernoulli random variables, and subsequently, keep β_0 fixed across multiple runs/iterations of the simulation. For each iteration m (we will average over 50 iterations later), we generate data corresponding to our two models under consideration as follows

$$Y_i = \mathcal{X}_i^T \beta_0 + \epsilon_i, \quad Y_i^{\mathcal{G}} = \mathcal{G}_i^T \beta_0 + \bar{\epsilon}_i. \quad (\text{V.1})$$

We recall from (II.7) that in this model

$$\sum_{j=1}^{p_0} \|h_j - \hat{h}_j\|^2 = \|\hat{\beta}_{\mathcal{X}}^L - \beta_0\|^2,$$

if the Lasso is used and $\|\hat{\beta}_{\mathcal{X}}^R - \beta_0\|^2$ in the case of ridge regression. In Figures (1a) and (1b), we plot the estimation risks $\|\hat{\beta}_{\mathcal{X}}^L - \beta_0\|^2$ and $\|\hat{\beta}_{\mathcal{X}}^R - \beta_0\|^2$ averaged over 50 iterations with red curves. We overlay the corresponding risks for the Gaussian problem, $\|\hat{\beta}_{\mathcal{G}}^L - \beta_0\|^2$ and $\|\hat{\beta}_{\mathcal{G}}^R - \beta_0\|^2$, as black dots. Observe the impeccable agreement of the risks under both settings that clearly validates our universality results and demonstrates its impressive efficacy in sample sizes as low as a couple of hundred.

B. Penalized Functional Regression

Recall that in the setup given in Example II-B we have,

$$Y_i = \sum_{j=1}^{p_0} \sum_{k=1}^d \zeta_{ik}^j \beta_k^j + \epsilon_i,$$

where $E[\zeta_{ik}^j] = 0$, $E[\zeta_{ik}^j \zeta_{il}^m] = \delta_{kl} \lambda_k$. Further ζ_{ik}^j and ζ_{il}^m are independent if $j \neq m$. We consider the case when $n = 500$ and $p_0 = 30$, and $d = 10$ i.e. we construct a setup where each row has 30 blocks of size 10 so that so that $(p_0 d)/n = 0.6$. The random variables across the block are independent while the random variables within the blocks are dependent. To construct dependent random variables within a block, we generate 10 independent Rademacher random variables and multiply them by a common random variable, as defined below. We repeat this procedure for every block. Let V be a random variable distributed as follows

$$P(V = 0) = \frac{1}{2}, \quad P(V = -\sqrt{2}) = \frac{1}{4}$$

$$P(V = \sqrt{2}) = \frac{1}{4}.$$

We generate i.i.d. rows X_i with j -th entry X_{ij} distributed as follows. Let $\{U_j\}_{j=1}^{300}$ and $\{V_j\}_{j=1}^{30}$ be two sets of independent random variables with U_j 's distributed as i.i.d. Rademacher and V_j 's as i.i.d. random variables with the distribution same as that of V . Now define $W_j = U_j V_{\lceil \frac{j}{10} \rceil}$ for $j \in \{1, \dots, 300\}$. Next we generate the diagonal matrix Λ with 30 identical blocks: the diagonal elements of the first block $\lambda_i^{1/2}, i \in 1 \dots 10$ are generated independently from $\text{Unif}(1, 2)$ and all other blocks are identical copies of the first block. It can be easily verified that the random variables are pairwise uncorrelated and the random variables within the blocks are not independent. We generate the design matrix \mathcal{X} as $\mathcal{X} = X \Lambda^{1/2}$. For the Gaussian case, we generate a 500×300 matrix \mathcal{G} with i.i.d. Gaussian entries and then define $\mathcal{G} = \mathcal{G} \Lambda^{1/2}$. As in the previous subsection, we generate the signal $\beta_0 = \{\theta_{0,k}^j\}_{j=1, k=1}^{p_0, d}$ as $\beta_{0,j} = n^{-1/2} b_j$, where the b_j 's are i.i.d. Bernoulli random variables, and subsequently, keep β_0 fixed across multiple runs/iterations of the simulation. Finally, we generate the random variables Y and $Y^{\mathcal{G}}$ following the linear models in (V.1) (after adjusting for the dimension).

In the aforementioned setting, we plot the estimation risk of the Lasso and the ridge as a function of λ . Figure (2a) shows Lasso with the red curve being the case of the dependent sub-Gaussian design and the black dots being the independent Gaussian design. Figure (2b) shows the corresponding plot for the ridge. In both cases, the risk curves match, demonstrating the validity of our universality result and the fact that the asymptotics kicks in for quite moderate sample size and dimensions.

C. Regression With Genetic Covariates

In Section II-C we have described a toy model of genetic covariates where the random variables within a block are uncorrelated but dependent. In this section we give another example where the covariates are generated from a hidden Markov model (HMM). The model has been considered in [76]

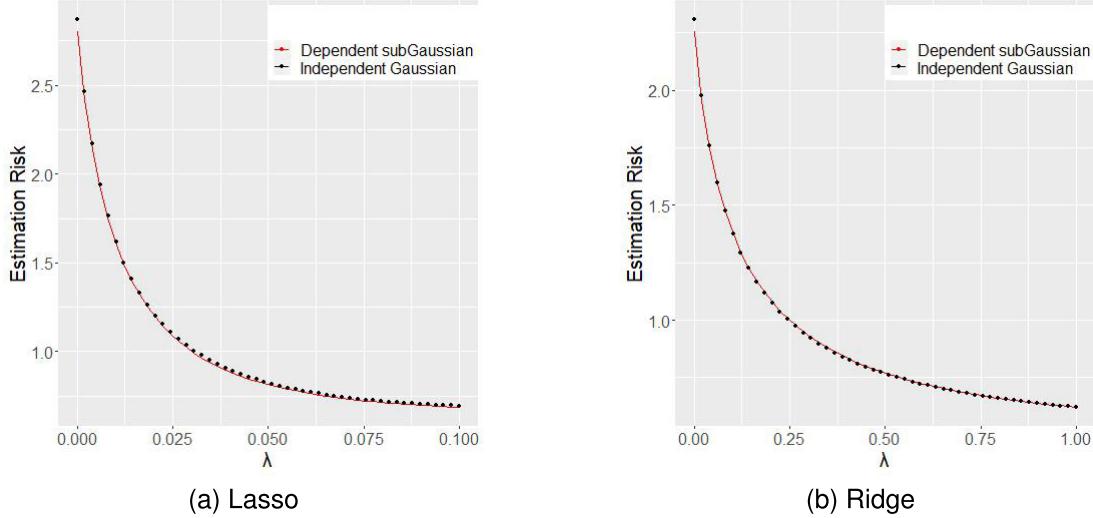


Fig. 1. Estimation risk in the penalized additive model (see Section II-A) averaged over 50 iterations with $n = 200$, $p_0 = 30$, and $d = 10$. The red curves correspond to the risk $\|\beta_X^L - \beta_0\|^2$ and $\|\beta_X^R - \beta_0\|^2$ for the Lasso and the ridge respectively, while the corresponding risks for the Gaussian problem $\|\beta_G^L - \beta_0\|^2$ and $\|\beta_G^R - \beta_0\|^2$ are overlaid as black dots.

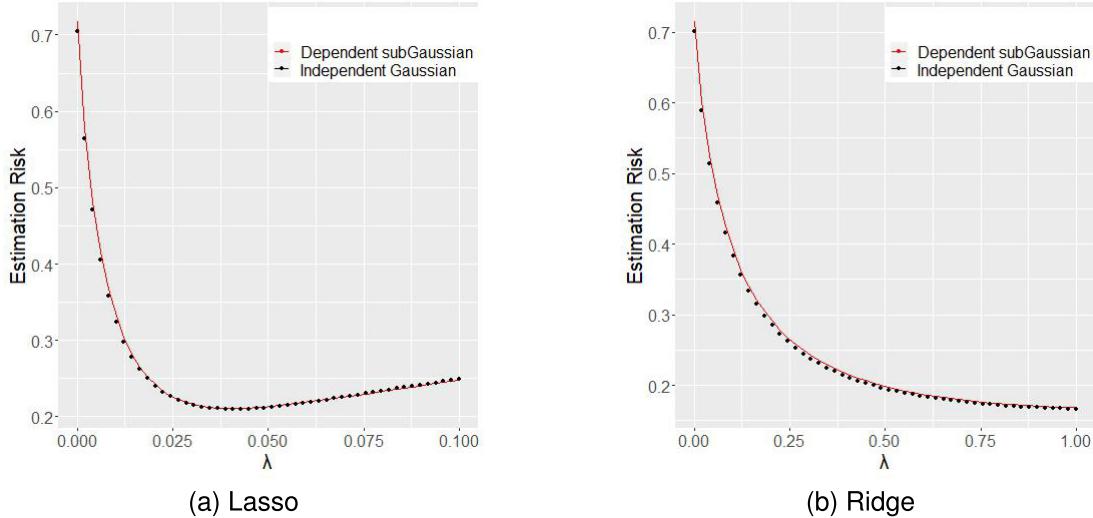


Fig. 2. Estimation risk in the penalized functional regression model (see Section II-B) averaged over 50 iterations with $n = 500$, $p_0 = 30$, and $d = 10$. The red curves corresponds to the risk $\|\beta_X^L - \beta_0\|^2$ and $\|\beta_X^R - \beta_0\|^2$ for Lasso and risk respectively, while the corresponding risks for the Gaussian $\|\beta_G^L - \beta_0\|^2$ and $\|\beta_G^R - \beta_0\|^2$ are overlaid as black dots.

and also forms the basis of the fastPHASE algorithm (see [77]). However, the covariates generated by this model will be correlated and we have to decorrelate the design before proceeding to show universality.

Recall that in GWAS the response Y is a quantitative trait of interest, while the covariates contains the genotypic information. In particular each row \mathbb{X}_i of the design matrix \mathbb{X} , is an element of $\{0, 1, 2\}^p$. Further, we assume block structure, so that the covariates in the j^{th} block form a d -dimensional vector, specifically $\mathbb{X}_i^j \in \{0, 1, 2\}^d$. Now we describe the HMM used to generate this random vector.

We consider a d -dimensional vector $H = (H_1, \dots, H_d)$ (called the vector of haplotypes) where H_l 's are Bernoulli random variables. The H_l 's will be “generated” by Z_l where Z_l form a Markov chain. Let $Z = (Z_1, \dots, Z_d)$ denote the Markov chain with Z_l taking K distinct values and the

transition matrix at l^{th} stage given by

$$Q_l^{hap}(k'|k) = P(Z_l = k'|Z_{l-1} = k)$$

with $k, k' \in \{1, \dots, K\}$. Further the marginal distribution of the first element Z_1 is given by:

$$q_1^{hap}(k) = \alpha_{1,k}$$

for $k = \{1, \dots, K\}$. The HMM can be written compactly as follows:

$$\begin{aligned} Z &\sim MC(q_1^{hap}, Q^{hap}) \\ H_l | Z &\sim f_l^{hap}(H_l | Z_l) \end{aligned}$$

for some distribution $f_l^{hap}(\cdot)$ called the emission distribution.

In [76], the authors further specify the transition matrix and emission matrix in terms of some parameters. The transition

matrix is given by

$$Q_l^{hap}(k'|k) = \begin{cases} e^{-r_l} + (1 - e^{-r_l})\alpha_{l,k'}, & k' = k, \\ (1 - e^{-r_l})\alpha_{l,k'}, & k \neq k' \end{cases}$$

and the emission distribution is parametrized by

$$f_l^{hap}(h_l, z_l, \theta) = \begin{cases} 1 - \theta_{l,z_l}, & h_l = 0, \\ \theta_{l,z_l}, & h_l = 1 \end{cases}$$

We refer the reader to Section V-A of [76] for the interpretation of these quantities. Next the genotype vector of a subject can be described as the element-wise sum of two i.i.d. haplotypes. Since the haplotypes form an HMM, the genotype vector also forms an HMM with the Markov chain having bivariate states (corresponding to unordered pairs of haplotype). The marginal distribution of the Markov chain, transition matrix and the emission probabilities are given as follows:

$$\begin{aligned} q_1^{gen}(\{k_a, k_b\}) &= \begin{cases} (\alpha_{1,k_a})^2, & k_a = k_b \\ 2\alpha_{1,k_a}\alpha_{1,k_b}, & k_a \neq k_b \end{cases} \\ Q_l^{gen}(\{k'_a, k'_b\}|\{k_a, k_b\}) &= \begin{cases} Q_l^{hap}(k'_a|k_a)Q_l^{hap}(k'_b|k_b) \\ + Q_l^{hap}(k'_b|k_a)Q_l^{hap}(k'_a|k_b) & k'_a \neq k'_b \\ Q_l^{hap}(k'_a|k_a)Q_l^{hap}(k'_b|k_b), & \text{otherwise} \end{cases} \\ f_l^{gen}(x_l; \{k_a, k_b\}, \theta) &= \begin{cases} (1 - \theta_{j,k_a})(1 - \theta_{j,k_b}), & x_l = 0 \\ \theta_{j,k_a}(1 - \theta_{j,k_b}) + \theta_{j,k_b}(1 - \theta_{j,k_a}), & x_l = 1 \\ \theta_{j,k_a}\theta_{j,k_b}, & x_l = 2. \end{cases} \end{aligned}$$

To simulate genetic covariate under a simple setup we let $K = 2$ so that $k_a \in \{0, 1\}$. Further, we are not interested in estimating the parameters as in [76], so we fix the parameters as follows:

$$\begin{aligned} \alpha_{l,0} &= 0.2, \alpha_{l,1} = 0.8, \forall l, \\ e^{-r_l} &= 0.5, \forall l, \\ \theta_{l,0} &= 0.4, \theta_{l,1} = 0.5, \forall l. \end{aligned}$$

The marginal distribution for the above choice of parameters are given by

$$\begin{aligned} q^{gen}(\{0, 0\}) &= 0.04, & q^{gen}(\{0, 1\}) &= 0.32, \\ q^{gen}(\{1, 1\}) &= 0.64. \end{aligned}$$

We also have the following transition matrix

$$Q^{gen} = \begin{array}{c|ccc} & \{0, 0\} & \{0, 1\} & \{1, 1\} \\ \hline \{0, 0\} & 0.36 & 0.48 & 0.16 \\ \{0, 1\} & 0.06 & 0.58 & 0.36 \\ \{1, 1\} & 0.01 & 0.18 & 0.81 \end{array}$$

Finally the emission probabilities are given by

$$\begin{aligned} f_l^{gen}(0; \{0, 0\}) &= 0.36, & f_l^{gen}(1; \{0, 0\}) &= 0.48, \\ f_l^{gen}(2; \{0, 0\}) &= 0.16 \\ f_l^{gen}(0; \{0, 1\}) &= 0.3 & f_l^{gen}(1; \{0, 1\}) &= 0.5 \\ f_l^{gen}(2; \{0, 1\}) &= 0.2 \end{aligned}$$

$$\begin{aligned} f_l^{gen}(0; \{1, 1\}) &= 0.25, & f_l^{gen}(1; \{1, 1\}) &= 0.5 \\ f_l^{gen}(2; \{1, 1\}) &= 0.25. \end{aligned}$$

We set $n = 400, p_0 = 30, d = 10$ so that $(p_0 d)/n = 0.75$. For the i^{th} row of the design matrix, we generate the data as follows

$$\mathcal{X}_i = (\tilde{\mathbb{X}}_i^{(1)T}, \dots, \tilde{\mathbb{X}}_i^{(30)T})^T.$$

where

$$\begin{aligned} \tilde{\mathbb{X}}_i^{(j)} &= \mathbb{V}^{-1/2}(\mathbb{X}_i^j - \mathbf{m}), & \mathbf{m} &= E[\mathbb{X}_i^j] \\ \mathbb{V} &= E[(\mathbb{X}_i^j - \mathbf{m})(\mathbb{X}_i^j - \mathbf{m})^T] \end{aligned}$$

with each $\mathbb{X}_i^j = (\mathbb{X}_{i,l}^j)_{l=1}^d$ generated according to the HMM described above. Further, for the purpose of simulation we estimate the vector \mathbf{m} and the matrix V from a separate sample of this HMM. We repeat this construction for $n = 400$ rows to create the design matrix \mathcal{X} and for the Gaussian problem, we construct the 400×300 matrix \mathcal{G} with i.i.d. Gaussian entries. We generate the error entries $\epsilon_i, \bar{\epsilon}_i$ so that they are i.i.d. standard normal. Finally, we generate the signal $\beta_0 = \{\beta_{0,k}^j\}_{j=1, k=1}^{p_0, d}$ as $\beta_{0,j} = n^{-1/2}\mathbf{b}_j$, where the \mathbf{b}_j s are i.i.d. Bernoulli random variables, and subsequently, keep β_0 fixed across multiple runs/iterations of the simulation. As in the previous two examples we generate data corresponding to our two models:

$$Y_i = \mathcal{X}_i^T \beta_0 + \epsilon_i, \quad Y_i^{\mathcal{G}} = \mathcal{G}_i^T \beta_0 + \bar{\epsilon}_i. \quad (\text{V.2})$$

Figures (3a) and (3b), shows the estimation risks $\|\hat{\beta}_{\mathcal{X}}^L - \beta_0\|^2$ with red curves for Lasso and ridge respectively. We overlay the corresponding risks for the Gaussian problem, $\|\hat{\beta}_{\mathcal{G}}^L - \beta_0\|^2$ and $\|\hat{\beta}_{\mathcal{G}}^R - \beta_0\|^2$, with black dots and the perfect agreement of the risk curves showcases universality.

VI. RELATED LITERATURE

In this section, we provide a detailed discussion of existing universality results in the literature, and situate our work in this broader context. Universality results are pervasive in the literature on random matrix theory. Over the past decade, the exploration of universality concerning the choice of the design matrix distribution has gained traction in the context of high-dimensional linear models. In the context of the Lasso, [57] established universality of the cost optimum, while for the elastic net, [40] established universality of the estimation risk. Similar results exist for random linear inverse problems ([78], [79]). Recently, [47] studied a broad spectrum of penalized regression problems and established universality of the optimum value of the empirical risk and the risk of regularized estimators under i.i.d. designs. [45] established universality of the optimizer as well as their risk under a broader class of structured and/or semi-random covariate distributions.

Parallel universality results exist for approximate message passing algorithms (cf. [39], [41], [46], [61]). Further, [42] demonstrated equivalence between regression problems where the rows of the design matrix arise from nonlinear transforms of Gaussian features and design matrices with linearized Gaussian features. Subsequently, [48] developed an extension

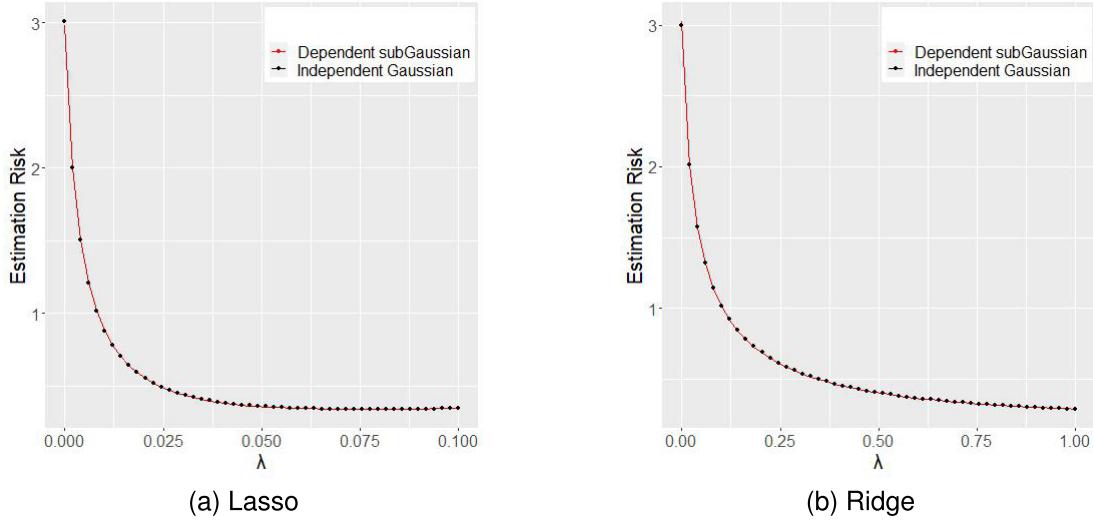


Fig. 3. Estimation risk in the penalized regression with genetic covariates averaged over 50 iterations with $n = 400$, $p_0 = 30$, and $d = 10$. The red curves corresponds to the risk $\|\beta_{\mathcal{X}}^L - \beta_0\|^2$ and $\|\beta_{\mathcal{X}}^R - \beta_0\|^2$ for Lasso and risk respectively, while the corresponding risks for the Gaussian $\|\beta_{\mathcal{G}}^L - \beta_0\|^2$ and $\|\beta_{\mathcal{G}}^R - \beta_0\|^2$ are overlaid as black dots.

of this result to the neural tangent model and a linear transformation of the i.i.d. model.

We note that in terms of dependent covariates under the linear model framework, [48] is the most relevant to our setting. Although the two-layer neural tangent model and the random feature model presented there are intriguing examples of dependence from a machine learning standpoint, for this discussion, we focus on the setup outlined in [48, Section 3.3], i.e., linear functions of independent entries. Notably, this example involves rows of the form $\Sigma^{1/2}Z_i$, where the entries of Z_i are independent sub-Gaussian. In this article, we broaden this understanding to encompass a different form of dependence with block structure (see Section II) that prior work fails to capture. As noted earlier, this block dependence arises in various widely used statistical models, as we demonstrate in Section II. We also note another limitation of [48] in the context of estimation risk. The paper proves the universality of the training error, which is the same as our optimum but instead of estimation risk, discusses the universality of test error. Although estimation risk can be obtained in their setting by a similar argument, it requires some strong convexity conditions for the penalty function. In our work, we overcome this barrier. Leveraging methods from [47] in conjunction with new proof ideas, we are able to address the case of the Lasso where strong convexity does not hold.

Our paper also builds a bridge between the proportional asymptotics literature and the nonlinear regression literature, which has a long history. In particular, we study additive models (see Section II) that were previously analyzed in the high-dimensional setting with $\log(p) = o(n)$ (where p quantifies the the number of component functions in the additive model) along with sparsity assumptions (see [67]) on the number of active component functions. [67] (see [80], [81], [82], [83] for similar models) considers both smoothness and sparsity based penalties and also allows for weak sparsity. We explore the additive models but under a high-dimensional setting where $p/n \rightarrow \kappa$ (κ can be greater than 1), without

imposing sparsity assumptions. Furthermore, while the prior literature was concerned with the rates of the estimation risk, we characterize the exact behavior of the risk or rather the exact constant that it converges to in our setting. A similar setup arises while dealing with the functional regression model, where we obtain results without any sparsity assumption whatsoever on the underlying function class. Finally, our paper considers a model from genomics where the covariates may be dependent but uncorrelated (see [74]). We have described a toy model (see [73], [76] for other models) in Section II-C, which showcases how such structures might arise and then showed that the model is amenable to our universality analysis.

VII. PROOFS

Our universality proof relies on the following key observation: for showing (i) universality of the optimum value of the empirical risk (Theorem 4.1) as well as ii) universality of the estimation risk of regularized estimators (Corollary 4.2), one requires to first establish that the estimators lie in some “nice” set with high probability. After this is established, points (i) and (ii) can be proved using separate techniques. For the subsequent discussion, we will use the term universality of the optimizer to mean (IV.2) in the statement of Theorem 4.2. Before we delve into our proof, we provide a brief description of our proof path below.

As a first step for points (i) and (ii), we require to prove that the optimizer lies in a compact set whose ℓ_2 norm is bounded by a constant and ℓ_∞ norm approaches 0 in the large n limit. We prove these key results in Lemmas 7.1 and 7.2 respectively. Subsequently for part (i), we establish a version of normal approximation for linear combinations of dependent random variables. For part (ii), we draw inspiration from the techniques in [47]. That said, we emphasize that the techniques used in [47] are heavily dependent on independence of the entries of the design, a property that does not hold in our setting. This raises additional technical challenges in our setting that

we handle as we proceed. For part (ii), it suffices to prove Theorem 4.2, once we have established the aforementioned compactness property. For Theorem 4.2 in turn, we require the compactness property, part (i), and a notion of separation in the Gaussian model, which we name the Gordon gap.

As will be explained in the next section, the compactness property is non-trivial to establish in the dependent sub-Gaussian setup, and we use a novel leave-d-out argument to solve the problem. In addition, while Stein's method has been used to show universality in several contexts, to the best of our knowledge, the dependency neighbourhood technique that we employ in this paper to show our normal approximation as part of point (i) has never seen prior usage in the context of regression models with dependent designs. Finally, the Gordon gap is a technical tool that is only relevant to the independent Gaussian design problem (see Section VII-C) and has been used before in [47] and [84]. Nonetheless, the scaling factor Λ , which is critical for covering our example in Section II-B, poses additional challenges involving convergence of fixed points of certain equations that we overcome via an equicontinuity argument.

A. Compactness Property

As a first step, we prove that the optimizers $\hat{\varphi}_{\mathcal{X}}^K := \hat{\beta}_{\mathcal{X}}^K - \beta_0$ belong to the following set with probability at least $1 - n^{-c}$

$$\mathcal{T}_n := \{\varphi \in \mathbb{R}^p : \|\varphi\|_{\infty} \leq C \sqrt{\frac{d \log n}{n}}, \|\varphi\| \leq C\}. \quad (\text{VII.1})$$

Since the ridge regression objective is strongly convex, the ℓ_2 bound $\|\hat{\varphi}_{\mathcal{X}}^R\| \leq C$ follows trivially. For the Lasso the analogous condition follows due to a restricted strong convexity type property that we discuss next. In the traditional theory for Lasso, in the setting of ultra-high dimensions ($p \gg n$), it is customary to assume sparsity on the underlying signal. This allows one to define the familiar and famous restricted strong convexity condition (see [85] and references cited therein). Using this, one can subsequently control the ℓ_2 norm of the Lasso solution. Since we refrain from explicitly assuming sparsity assumptions on the signal, this strategy does not directly translate to our setting. However, in our high-dimensional regime, even when $p > n$, when the tuning parameter is sufficiently large (Assumption 3, see also [47, Lemma 6.3]), one can establish that the Lasso solution is sparse (with sparsity proportional to the sample size) with high probability. This implies that having control on the smallest singular value of \mathcal{X}_S (matrix formed by restricting \mathcal{X} to a sparse subset) is sufficient for ensuring a restricted strong convexity type property of the Lasso solution. This allows one to bound the ℓ_2 norm of the Lasso solution.

Establishing the ℓ_{∞} bound poses challenges since the leave-one-out method as outlined in [47] fails in the presence of dependence in the data. Instead we use a novel leave-d-out method which essentially considers the following optimization problem

$$\hat{\varphi}^{(S)} = \arg \min_{\varphi \in \mathbb{R}^p, \varphi_S=0} \frac{1}{2n} \|X\varphi - \xi\|^2 + f(\varphi),$$

where S is a subset of $1, \dots, p$ of cardinality d . A detailed analysis together with our block dependence assumption then yields a suitable upper bound on $\|\hat{\varphi}_{\mathcal{X},S}^K\|$, which in turn provides the desired upper bound on $\|\hat{\varphi}_{\mathcal{X}}^K\|_{\infty}$. Such leave-d-out arguments have been used in prior literature in [9], [20], and [21] in other contexts, however, these works do not use this for universality arguments. For convenience we will show a scaled version of $\hat{\varphi}_{\mathcal{X}}^K$ (which we call $\hat{\theta}_{\mathcal{X}}^K$ and is defined below) belongs to \mathcal{T}_n (see (VII.1)), from which it easily follows that $\hat{\varphi}_{\mathcal{X}}^K \in \mathcal{T}_n$. First we introduce some notations.

Recall the empirical risk minimization problem

$$\hat{R}_n(Y, \mathcal{X}, \beta) = \frac{1}{2n} \sum_{i=1}^n \|Y_i - \mathcal{X}_i^T \beta\|^2 + \lambda f(\beta).$$

Since \mathcal{X} is of the form $X\Lambda^{1/2}$, we let $\vartheta = \Lambda^{1/2}\beta$. Then the risk minimization problem can be rewritten as

$$\hat{R}_n(Y, X, \vartheta) = \frac{1}{2n} \sum_{i=1}^n \|Y_i - X_i^T \vartheta\|^2 + \lambda f(\Lambda^{-1/2}\vartheta).$$

Since $\varphi = \beta - \beta_0$, we define $\theta = \vartheta - \vartheta_0$ so that $\theta = \Lambda^{1/2}\varphi$. The new parametrization is concisely presented in the following display

$$\begin{aligned} \vartheta &= \Lambda^{1/2}\beta, & \vartheta_0 &= \Lambda^{1/2}\beta_0 \\ \theta &= \vartheta - \vartheta_0, & \Lambda^{-1/2} &= \tilde{\Omega}. \end{aligned} \quad (\text{VII.2})$$

For brevity, we will also denote the value of $\hat{R}_n(Y, X, \vartheta)$ on using the variable θ as $R^K(X, \theta)$ (with K=L or R depending on whether ℓ_1 or ℓ_2 penalty is used), where we have suppressed the dependence on ϑ_0 (and hence Y) and n . Finally, we define the following optimizers

$$\begin{aligned} \hat{\theta}_X^R &= \arg \min_{\theta \in \mathbb{R}^p} \frac{1}{2n} \|X\theta - \xi\|^2 + \frac{\lambda}{2} \|\tilde{\Omega}(\theta + \vartheta_0)\|^2, \\ \hat{\theta}_X^L &= \arg \min_{\theta \in \mathbb{R}^p} \frac{1}{2n} \|X\theta - \xi\|^2 + \frac{\lambda}{\sqrt{n}} \|\tilde{\Omega}(\theta + \vartheta_0)\|_1. \end{aligned}$$

The following lemmas proved in the appendix show that the optimizers lie in the set \mathcal{T}_n with high probability.

Lemma 7.1: Under Assumptions 1-3 we have

$$\|\hat{\theta}_X^K\|_2 \leq C$$

with probability at least $1 - n^{-c}$ for positive constants c and C .

Lemma 7.2: Under Assumptions 1-3 we have

$$\|\hat{\theta}_X^K\|_{\infty} \leq C \sqrt{\frac{d \log n}{n}}$$

with probability at least $1 - n^{-c}$ for positive constants c and C .

B. Universality of the Optimum

We next describe our proof for the universality of the optimum that corresponds to point (i) discussed in the beginning of Section VII. We may divide this proof into three steps. The first step replaces the minimum over the compact set with a minimum over a discrete set, and shows that the approximation error goes to zero. The second step involves approximating the minimization over the discrete set by a smoothed minimum,

and showing once again that the approximation error is negligible. The final step proves that this smoothed minimum is universal, that is, its value is asymptotically the same if the dependent design is replaced by an i.i.d. Gaussian design with matching moments.

This proof strategy has been used in prior work (see [47, Theorem 2.3] or [48, Theorem 1]). The universality of the smoothed minimum is the crucial step and depends heavily on the dependence/independence structure in the underlying design matrix. It involves showing for any given set \mathcal{S}_n we have

$$\sup_{\theta \in \mathcal{T}_n \cap \mathcal{S}_n} |E[\phi(X_i^\top \theta)] - E[\phi(G_i^\top \theta)]| \rightarrow 0 \quad (\text{VII.3})$$

for a suitable choice of test function ϕ , where X_i is a row vector from the block dependent design and G_i is a row vector from the independent Gaussian design. The arguments showing this step differs in our case compared to prior work. To establish such a pointwise normality condition, we use a version of Stein's method that involves dependency graphs together with the fact that the optimizer lies in the compact set \mathcal{T}_n .

Heuristically the universality result holds because $X_i^\top \hat{\beta}^K$ is a weighted sum of dependent sub-Gaussian random variables that converges to a Gaussian random variable if the dependence is “low” in a suitable sense. It can be easily observed that the weighted sum may fail to converge if the weights are chosen arbitrarily, for example if $\hat{\beta}^K = (1, 0, \dots, 0)$, convergence to a Gaussian cannot be expected. For this reason we require the coordinates of $\hat{\beta}^K$ to be small enough, which justifies the ℓ_∞ bound proved in Lemma 7.2.

Proof of Theorem 4.1:

For this proof, we will assume that \mathcal{S}_n is a subset of \mathcal{T}_n , which we can assume without loss of generality.

Step 1: Discretization

Pick a δ net of the set \mathcal{S}_n and call it $\mathcal{S}_{n,\delta}$. Let $\hat{\theta} = \arg \min_{\theta \in \mathcal{S}_n} R(\theta, X)$ and $\tilde{\theta}$ be a point in $\mathcal{S}_{n,\delta}$ such that $\|\hat{\theta} - \tilde{\theta}\| \leq \delta$. For the purposes of this proof, we write $\hat{\theta}_X^K$ as $\hat{\theta}$ and $R^K(\theta, X)$ as $R(\theta, X)$ since the same proof works for both Lasso and ridge regression. We observe the following chain of inequalities

$$\begin{aligned} & |\min_{\theta \in \mathcal{S}_{n,\delta}} R(\theta, X) - \min_{\theta \in \mathcal{S}_n} R(\theta, X)| \\ &= \min_{\theta \in \mathcal{S}_{n,\delta}} R(\theta, X) - \min_{\theta \in \mathcal{S}_n} R(\theta, X) \\ &\leq R(\tilde{\theta}, X) - R(\hat{\theta}, X) \\ &\leq |R(\tilde{\theta}, X) - R(\hat{\theta}, X)| \\ &\leq \frac{1}{2n} (\|X\tilde{\theta} - \xi\|^2 - \|X\hat{\theta} - \xi\|^2) + |f(\tilde{\theta}) - f(\hat{\theta})| \\ &\leq \frac{1}{2n} \|X(\tilde{\theta} - \hat{\theta})\|^2 + \frac{1}{n} |(X\tilde{\theta} - \xi)^\top X(\tilde{\theta} - \hat{\theta})| \\ &\quad + C_1 \|\tilde{\theta} - \hat{\theta}\| \\ &\leq \frac{1}{2n} \|X(\tilde{\theta} - \hat{\theta})\|^2 + \frac{1}{n} (\|X\|\|\tilde{\theta}\| + \|\xi\|) \|X\| \|\tilde{\theta} - \hat{\theta}\| \\ &\quad + C \|\tilde{\theta} - \hat{\theta}\|. \end{aligned}$$

Noting that $\|X\| \leq C\sqrt{n}$ and $\|\xi\| \leq C\sqrt{n}$ with probability at least $1 - n^{-c}$ and $\max\{\|\theta\|, \|\tilde{\theta}\|\} \leq C$ (since \mathcal{S}_n and $\mathcal{S}_{n,\delta}$

are subsets of \mathcal{T}_n) we conclude that

$$|\min_{\theta \in \mathcal{S}_{n,\delta}} R(\theta, X) - \min_{\theta \in \mathcal{S}_n} R(\theta, X)| \leq C\delta$$

with probability at least $1 - n^{-c}$. For a bounded differentiable function with bounded Lipschitz derivative we have

$$\begin{aligned} & |E[\psi(\min_{\theta \in \mathcal{S}_{n,\delta}} R(\theta, X))] - E[\psi(\min_{\theta \in \mathcal{S}_n} R(\theta, X))]| \\ &\leq C\|\psi'\|_\infty \delta. \end{aligned} \quad (\text{VII.4})$$

Step 2: Approximation by smoothed minimum

The smoothed minimum is defined as follows

$$f_\delta(\alpha, X) = -\frac{1}{n\alpha} \log \sum_{\theta \in \mathcal{S}_{n,\delta}} \exp(-n\alpha R(\theta, X))$$

with the following approximation error (cf. [48])

$$|f_\delta(\alpha, X) - \min_{\theta \in \mathcal{S}_{n,\delta}} R(\theta, X)| \leq C(\delta)\alpha^{-1}.$$

Again using a Lipschitz function as in (VII.4), we obtain

$$\begin{aligned} & |E[\psi(f_\delta(\alpha, X))] - E[\psi(\min_{\theta \in \mathcal{S}_{n,\delta}} R(\theta, X))]| \\ &\leq C\|\psi'\|_\infty \alpha^{-1}. \end{aligned} \quad (\text{VII.5})$$

Step 3: Universality of the smoothed minimum

The universality of the smoothed minimum follows via the interpolation path as shown in [48]. In particular one defines matrices U_t where the rows are given by

$$U_{t,i} = \sin(t)X_i + \cos(t)G_i$$

and uses the relation

$$\begin{aligned} & |E(\psi(f_\delta(\alpha, X)) - \psi(f_\delta(\alpha, G)))| \\ &= \frac{2}{\pi} \left| E \int_0^{\pi/2} \frac{\partial}{\partial t} (\psi(f_\delta(\alpha, U_t))) \right|. \end{aligned} \quad (\text{VII.6})$$

It can be shown that $|E \frac{\partial}{\partial t} (\psi(f_\delta(\alpha, U_t)))| \rightarrow 0$ once the pointwise normality condition (see (VII.3)) is satisfied (see [48, Lemma 4]), which we prove via Stein's method. Thus from (VII.4), (VII.5), and (VII.6), we conclude that

$$\lim_{n \rightarrow \infty} |E[\psi(\min_{\theta \in \mathcal{S}_n} R(\theta, X))] - E[\psi(\min_{\theta \in \mathcal{S}_n} R(\theta, G))]| \rightarrow 0$$

by suitably choosing $\delta \rightarrow 0$, $\alpha \rightarrow \infty$, and $n \rightarrow \infty$, once we verify (VII.3). To show the latter, recall that \mathcal{S}_n can be assumed to a subset of \mathcal{T}_n without loss of generality and hence we aim to show

$$\sup_{\theta \in \mathcal{S}_n} |E[\phi(X_i^\top \theta)] - E[\phi(G_i^\top \theta)]| \rightarrow 0.$$

We use the following relation

$$\begin{aligned} & |E[\phi(X_i^\top \theta)] - E[\phi(G_i^\top \theta)]| \\ &= |E[\phi(\nu X_i^\top \theta / \nu)] - E[\phi(\nu Z)]|, \end{aligned}$$

where $\nu^2 = \|\theta\|^2$ and $Z \sim N(0, 1)$. Define

$$\psi(x) = e^{x^2/2} \int_{-\infty}^x e^{-y^2/2} (\phi(\nu y) - E(\phi(\nu Z))) dy,$$

which satisfies the Stein's identity

$$\psi'(x) - x\psi(x) = \phi(\nu x) - E(\nu Z),$$

and if $\|\phi'(x)\|_\infty \leq B$, then

$$\max\{\psi(x), \psi'(x), \psi''(x)\} \leq 2\nu B \leq C, \quad (\text{VII.7})$$

since $\|\theta\| \leq C_1$ (for some constant C_1) in our setting. Thus it suffices to upper bound

$$E \left[\psi' \left(\frac{X_i^\top \theta}{\nu} \right) - \frac{X_i^\top \theta}{\nu} \psi \left(\frac{X_i^\top \theta}{\nu} \right) \right].$$

By using (VII.7), the block dependent structure (Definition 2.1), and [86, Theorem 3.6] we have

$$\begin{aligned} & E \left[\psi' \left(\frac{X_i^\top \theta}{\nu} \right) - \frac{X_i^\top \theta}{\nu} \psi \left(\frac{X_i^\top \theta}{\nu} \right) \right] \\ & \leq Cd^2\varsigma_3 + Cd^{3/2}\sqrt{\varsigma_4}, \end{aligned}$$

with

$$\begin{aligned} \varsigma_3 &= \frac{\max_j E|X_{ij}|^3 \sum_{j=1}^p |\theta_j|^3}{\|\theta\|^3}, \\ \varsigma_4 &= \frac{\max_j E|X_{ij}|^4 \sum_{j=1}^p |\theta_j|^4}{\|\theta\|^4}. \end{aligned}$$

Now suppose $\max_j |\theta_j| \leq \alpha_n$ and $\|\theta\| > \epsilon$, then with $\varsigma = \max\{\max_j E|x_{ij}|^3, \max_j E|x_{ij}|^4\}$ we have

$$\begin{aligned} & E \left[\psi' \left(\frac{X_i^\top \theta}{\nu} \right) - \frac{X_i^\top \theta}{\nu} \psi \left(\frac{X_i^\top \theta}{\nu} \right) \right] \\ & \leq \frac{d^2\varsigma\alpha_n}{\epsilon} + \frac{d^{3/2}\sqrt{\varsigma}\alpha_n}{\epsilon}. \end{aligned} \quad (\text{VII.8})$$

Since $\theta \in \mathcal{S}_n \subset \mathcal{T}_n$, $\alpha_n \rightarrow 0$ and hence the RHS of (VII.8) goes to 0 as $n \rightarrow \infty$. On the other hand, when $\|\theta\| < \epsilon$, we have

$$\begin{aligned} & |E[\phi(X_i^\top \theta)] - E[\phi(G_i^\top \theta)]| \\ & \leq \|\phi'\|_\infty \left[(E(X_i^\top \theta)^2)^{1/2} + (E(G_i^\top \theta)^2)^{1/2} \right] \leq C_1\epsilon, \end{aligned}$$

for some C_1 since $\|\theta\| \leq C$ for $\theta \in \mathcal{T}_n$. Thus we have

$$\begin{aligned} & \sup_{\theta \in \mathcal{S}_n} |E[\phi(X_i^\top \theta)] - E[\phi(G_i^\top \theta)]| \\ & \leq \sup_{\{\theta \in \mathcal{S}_n\} \cap \{\|\theta\| \geq \epsilon\}} |E[\phi(X_i^\top \theta)] - E[\phi(G_i^\top \theta)]| \\ & \quad + \sup_{\{\theta \in \mathcal{S}_n\} \cap \{\|\theta\| \leq \epsilon\}} |E[\phi(X_i^\top \theta)] - E[\phi(G_i^\top \theta)]| \\ & \leq \frac{d^2\varsigma\alpha_n}{\epsilon} + \frac{d^{3/2}\sqrt{\varsigma}\alpha_n}{\epsilon} + C_1\epsilon. \end{aligned}$$

Now choose

$$\epsilon = \sqrt{\frac{Cd^2\varsigma\alpha_n + Cd^{3/2}\sqrt{\varsigma}\alpha_n}{C_1}}$$

so that

$$\begin{aligned} & \sup_{\theta \in \mathcal{S}_n} |E[\phi(X_i^\top \theta)] - E[\phi(G_i^\top \theta)]| \\ & \leq 2\sqrt{C_1(Cd^2\varsigma\alpha_n + Cd^{3/2}\sqrt{\varsigma}\alpha_n)} \rightarrow 0 \end{aligned}$$

as desired.

Remark: Note that in the proof we only require $d^2\alpha_n$ converges to 0. So the result holds for growing d . In particular, α_n can be chosen as $(d \log n/n)^{1/2}$ and the proof works as long as $d^5 \log n/n \rightarrow 0$. However, we note that since the tightness of the upper bounds given in Lemma 7.2 and (VII.8) cannot be guaranteed in general, the maximum d beyond which universality fails remains an interesting open question.

C. Universality of the Optimizer

We use the terms universality of the optimizer and universality of the estimation risk interchangeably. This is because the proof of Theorem 4.2 shows the universality of the optimizer and an application of Theorem 4.2 gives Corollary 4.2, which deals with the universality of the estimation risk.

Our proof of this step builds upon [47]. The universality of the optimizer is established in [47, Theorem 2.4] under two conditions (called (O1) and (O2) in [47]) with the regression setup considered in [47, Theorem 3.1]. We note that condition (O1) mentioned in there is the same as the bound on the ℓ_∞ norm of the optimizer, which we prove in Lemma 7.2 by our leave- d -out method. Condition (O2) is the Gordon gap that states the following: if there is a difference in the unrestricted optimum and the optimum restricted to a particular set with high probability, then the optimizer does not belong to that set with high probability. In other words, for any pair of positive constants ψ^K and z , the following holds

$$\begin{aligned} & \{\inf_{\theta \in R^p} R^K(\theta, G) \leq \psi^K + z\} \\ & \cap \{\inf_{\theta \in \mathcal{S}_n} R^K(\theta, G) \geq \psi^K + 2z\} \subset \{\hat{\theta}_G^K \in S_n^c\}. \end{aligned} \quad (\text{VII.9})$$

We observe that the statement above concerns independent Gaussian design only i.e. we do not need to establish such a condition for the dependent sub-Gaussian design matrix. Although it may seem nontrivial that the universality of the optimizer follows from checking this condition for the Gaussian case, we show that (see proof of Theorem 4.2) the Gordon gap coupled with Theorem 4.1 (which shows universality of the optimum value of the empirical risk for all subsets of the nice compact set) readily yields the desired universality of the optimizer. This proof strategy was also utilized in [47] and [48]. Although our proof for the Gordon gap builds upon [47, Theorems 3.4, 3.8], the scaling matrix Λ calls for a more nuanced analysis. In particular, Lemma 7.3 below plays a critical role in the proof of Theorem 4.2. The scaling factor introduces challenges in establishing this lemma. We overcome this by employing a nontrivial equicontinuity argument (see Appendix A-C), a step not necessitated in the case of [47].

Let $\omega_i = \tilde{\Omega}_{ii}$ and $\mu_0 = \sqrt{n}\vartheta_0$ where $\tilde{\Omega}$ and ϑ are defined in (VII.2). Recall the fixed point solution β_*^K, γ_*^K of the equations (III.4) and define $\theta_*^K = w_*^K/\sqrt{n}$ where

$$w_{*,i}^K = \eta_k \left(\mu_{0,i} + \gamma_*^K g_i, \frac{\gamma_*^K \lambda \omega_i^k}{\beta_*^K} \right) - \mu_{0,i}$$

with $k = 1$ for $K = R$ and $k = 2$ for $K = L$. Let us define the set S_n^ϵ as follows

$$S_n^\epsilon := \{\theta : |\varrho(\theta) - E(\varrho(\theta_*^K))| \geq \epsilon\}. \quad (\text{VII.10})$$

Henceforth we will omit the ϵ and call it \mathcal{S}_n since Theorem 4.2 is proved for a fixed ϵ . We have the following lemma.

Lemma 7.3: There exists positive constants ψ^R and ψ^L (denoted by ψ^K with $K = R$ for ridge and $K = L$ for Lasso) and z such that the following holds

$$P\left(\inf_{\theta \in R^p} R^K(\hat{\theta}, G) \geq \psi^K + z\right) \leq \delta_n$$

$$P\left(\inf_{\theta \in \mathcal{S}_n} R^K(\hat{\theta}, G) \leq \psi^K + 2z\right) \leq \delta_n$$

for some $\delta_n \rightarrow 0$.

Remark 7.1: In the above lemma $\psi^K = \psi^K(\beta_*^K, \gamma_*^K)$ where the function $\psi^K(\cdot, \cdot)$ is given in (III.3) and the constants β_*^K and γ_*^K is given by the (III.4). Further z is constant which depends only on ϵ - the same ϵ used to define the set S_n^ϵ in (VII.10).

Now we are ready to prove Theorem 4.2.

Proof of Theorem 4.2:

We start with formalizing the statement in (VII.9). We have the following set of inequalities for suitable choices of γ_n , δ_n , $\epsilon_n \rightarrow 0$.

$$\begin{aligned} & P(\hat{\theta}_X^K \in \mathcal{S}_n) \\ & \leq P(\hat{\theta}_X^K \in \mathcal{S}_n \cap \mathcal{T}_n) + \epsilon_n \\ & \leq P\left(\inf_{\theta \in \mathcal{T}_n} R^K(\theta, X) \geq \psi^K + 3z\right) \\ & \quad + P\left(\inf_{\theta \in \mathcal{S}_n \cap \mathcal{T}_n} R^K(\theta, X) \leq \psi^K + 6z\right) + \epsilon_n \\ & \leq P\left(\inf_{\theta \in \mathcal{T}_n} R^K(\theta, G) \geq \psi^K + z\right) \\ & \quad + P\left(\inf_{\theta \in \mathcal{S}_n \cap \mathcal{T}_n} R^K(\theta, G) \leq \psi^K + 2z\right) + \epsilon_n + 2\gamma_n \\ & \leq P\left(\inf_{\theta} R^K(\theta, G) \geq \psi^K + z\right) \\ & \quad + P\left(\inf_{\theta \in \mathcal{S}_n} R^K(\theta, G) \leq \psi^K + 2z\right) + 3\epsilon_n + 2\gamma_n \\ & \leq 2\delta_n + 3\epsilon_n + 2\gamma_n. \end{aligned}$$

We note that the first and fourth inequality follows from Lemmas 7.1 and 7.2. The second inequality follows from the following inclusion

$$\begin{aligned} & \left\{ \inf_{\theta \in \mathcal{T}_n} R^K(\theta, X) \leq \psi^K + 3z \right\} \\ & \cap \left\{ \inf_{\theta \in \mathcal{S}_n \cap \mathcal{T}_n} R^K(\theta, X) \geq \psi^K + 6z \right\} \\ & \subset \left\{ \hat{\theta}_X^K \in (\mathcal{S}_n \cap \mathcal{T}_n)^c \right\}. \end{aligned}$$

The third inequality follows from Theorem 4.1 with a suitable choice of test function, while the last inequality follows from Lemma 7.3. Next we use (VII.2) to obtain our statements with

the old parameters. For φ , we use the fact that $\varphi = \tilde{\Omega}\theta$ and $\theta = w/\sqrt{n}$ to obtain

$$\varphi_*^K = \tilde{\Omega} \frac{w_*^K}{\sqrt{n}}$$

so that

$$\varphi_{*,i}^R = \eta_2 \left(\beta_{0,i} + \frac{\gamma_*^R \omega_i g_i}{\sqrt{n}}, \frac{\gamma_*^R \lambda \omega_i^2}{\beta_*^R} \right) - \beta_{0,i} \quad (\text{VII.11})$$

$$\varphi_{*,i}^L = \eta_1 \left(\beta_{0,i} + \frac{\gamma_*^L \omega_i g_i}{\sqrt{n}}, \frac{\gamma_*^L \lambda \omega_i^2}{\sqrt{n} \beta_*^L} \right) - \beta_{0,i}, \quad (\text{VII.12})$$

where we have used the fact that

$$\eta_1(\alpha x, \alpha \lambda) = \alpha \eta_1(x, \lambda), \quad \eta_2(\alpha x, \lambda) = \alpha \eta_2(x, \lambda).$$

Since $\varphi_X^K = \tilde{\Omega} \hat{\theta}_X^K$, where the entries of $\tilde{\Omega}$ are bounded, we can easily redefine ϱ so that

$$P(|\varrho(\hat{\varphi}_X^K) - E(\varrho(\varphi_*^K))| \geq \epsilon) \leq 2\delta_n + 2\epsilon_n + 3\gamma_n. \quad (\text{VII.13})$$

Now we observe that although the above result is proved for dependent sub-Gaussian designs, the latter subsumes the case of independent Gaussian design and hence it trivially follows that

$$\begin{aligned} & P(|\varrho(\hat{\varphi}_G^K) - E(\varrho(\varphi_*^K))| \geq \epsilon) \leq 2\delta_n + 3\epsilon_n + 2\gamma_n. \quad (\text{VII.14}) \\ & \text{Combining (VII.13) and (VII.14), we obtain that} \end{aligned}$$

$$P(|\varrho(\hat{\varphi}_X^K) - \varrho(\hat{\varphi}_G^K)| \geq 2\epsilon) \leq 4\delta_n + 6\epsilon_n + 4\gamma_n.$$

Defining $4\delta_n + 6\epsilon_n + 4\gamma_n$ as ε_n and replacing 2ϵ by ϵ we observe that the desired claim is proved. \square

We now have all ingredients in place for the proof of Corollary 4.2.

Proof of Corollary 4.2: Choose ϱ as the function $\|\cdot\|^2$. Although this function is not Lipschitz in the entire domain \mathbb{R}^p it is Lipschitz in the bounded domain $\{\varphi : \|\varphi\| \leq C\}$. Since our optimizers have bounded ℓ_2 norm with probability at least $1 - n^{-c}$, we work in the domain $\{\varphi : \|\varphi\| \leq C\}$ for our purposes and use Theorem 4.2 to conclude

$$\begin{aligned} & \|\hat{\varphi}_X^R\|^2 - E_{\hat{Q}_n} \left[\eta_2 \left(M + \frac{\gamma_*^R \Omega Z}{\sqrt{n}}, \frac{\gamma_*^R \lambda \Omega^2}{\beta_*^R} \right) - M \right]^2 \\ & \xrightarrow{P} 0, \\ & \|\hat{\varphi}_X^L\|^2 - E_{\hat{Q}_n} \left[\eta_1 \left(M + \frac{\gamma_*^L \Omega Z}{\sqrt{n}}, \frac{\gamma_*^L \lambda \Omega^2}{\sqrt{n} \beta_*^L} \right) - M \right]^2 \\ & \xrightarrow{P} 0. \end{aligned}$$

Substituting $\hat{\varphi}_X^K = \hat{\beta}_X^K - \beta_0$, we conclude our proof.

VIII. DISCUSSION

In this paper, we established universality results within the framework of empirical risk minimization, demonstrating that the optimal value of the regularized empirical risk and estimation risk of regularized estimators converge to the same value for block dependent sub-Gaussian designs, as in the case of i.i.d. Gaussian designs. In sum, to the best of the authors' knowledge, our paper is the first in the literature that handles dependence structures beyond correlation-based

dependence or those allowed by right-rotationally-invariant designs. Our examples from Section II demonstrate the importance of studying such general dependencies, providing the first segue into nonlinear problems. Notably, this allows us to characterize the (asymptotically) exact risk of estimators for popular nonlinear regression models under high-dimensional covariates. We conclude our manuscript with a couple of directions for future work.

First, our results do not account for dependence captured by an explicit covariance matrix. A logical extension of our work would involve proving dependence when feature vectors take the form $\Sigma^{1/2}Z_i$. Here, instead of Z_i being a random vector with independent sub-Gaussian entries, as considered in [48], Z_i is isotropic, and block dependent sub-Gaussian vectors. We conjecture that under some regularity assumptions on the covariance matrix Σ , the optimal value of the risk and the risk of regularized estimators should be asymptotically the same as that of the design with rows of the form $\Sigma^{1/2}G_i$ where the entries of G_i are i.i.d. Gaussian. While proving universality of the optimum in this scenario aligns with the approach in [48], generalizing our leave-d-out method in presence of a general Σ and subsequently proving universality of the optimizer presents nontrivial challenges, which we defer to future work.

Second, an intriguing avenue involves considering a more general dependence pattern, where each entry of the feature vector depends on only a small number of other entries, without a block structure. Currently, the block structure is pivotal in our proofs of ℓ_∞ bounds via our leave-d-out method. We view our current contribution as a stepping stone towards investigating these other notions of dependence. The absence of the block structure implies a stronger form of dependence, necessitating new tools. We defer this to future investigation.

APPENDIX A

PROOF OF THE LEMMAS

First we state a concentration result which will be used in proving the theorems.

A. Concentration Result

Lemma 1.1: Let X be an $n \times p$ matrix whose rows X_i are independent sub-Gaussian isotropic random vectors in \mathbb{R}^n . Then for every $t \geq 0$, with probability at least $1 - 2\exp(-c_K t^2)$ one has

$$\begin{aligned} \sqrt{n} - C_K \sqrt{p} - t &\leq s_{\min}(X) \leq s_{\max}(X) \\ &\leq \sqrt{n} + C_K \sqrt{p} + t, \end{aligned} \quad (\text{A.1})$$

where $s_{\min}(X)$ and $s_{\max}(X)$ are the smallest and largest singular values of the matrix X . The constant s, C_K, c_K depend only on the sub-Gaussian norm $K = \max_i \|X_i\|_{\psi_2}$ of the rows.

Proof: See [87]. □

B. Bounds on the Optimizers

Throughout this section we will work with Assumptions 1 and 3. Further we use the following modified assumption instead of 2 to account for change in scale (see (VII.2))

Assumption 4: The signal satisfies the bound $\|\vartheta_0\| \leq C$. Let $v_i = \sqrt{n}\vartheta_{0,i}$ for $1 \leq i \leq p$. Then we assume that

$$\frac{1}{p} \sum_{i=1}^p \delta_{\lambda_i, v_i} \xrightarrow{W_2} \mu.$$

Further we assume that $\|v\|_\infty \leq C$.

Proof of lemma 7.1: Recall that we work under the proportional asymptotic setup i.e. $\lim_{n \rightarrow \infty} p/n = \kappa$. For the sake of clarity we will first deal with the setup with $\kappa^{-1} > C_0 > C_K$ for some C_0 where C_K is the constant appearing in Lemma 1.1. We will later show that in the general setup similar bounds can be achieved by suitably modifying the tuning parameter λ .

Case 1: $n/p \rightarrow \kappa^{-1} > C_0 > C_K^2$

Ridge regression

Recall that $\tilde{\Omega} = \Lambda^{-1/2}$. By the optimality of $\hat{\theta}_X^R$, it easily follows that

$$\begin{aligned} &\frac{1}{2} \|X\hat{\theta}_X^R - \xi\|^2 + \frac{n\lambda}{2} \|\tilde{\Omega}(\hat{\theta}_X^R + \vartheta_0)\|^2 \\ &\leq \frac{1}{2} \|\xi\|^2 + \frac{n\lambda}{2} \|\tilde{\Omega}\vartheta_0\|^2, \end{aligned}$$

whence $\|X\hat{\theta}_X^R\|^2 \leq C(\frac{1}{2}\|\xi\|^2 + \frac{n\lambda}{2}\|\tilde{\Omega}\vartheta_0\|^2)$.

Using a lower bound on the singular value of X from Lemma 1.1 we observe that the left-hand side of the last inequality is lower bounded by $cn\|\hat{\theta}_X^R\|^2$ with probability at least $1 - n^{-c}$ by choosing $t = c_1\sqrt{\log n}$ (for suitable choice of c_1). Similarly using $E[\|\xi\|^2] \leq Cn$ and sub-Gaussian concentration inequality for the random vector ξ we easily conclude that $\|\hat{\theta}_X^R\| \leq C$ with probability at least $1 - n^{-c}$. Note that we have used $\|\vartheta_0\|^2 \leq C$ and that $\tilde{\Omega}$ is a diagonal matrix with bounded entries.

Lasso

We have an analogous bound for Lasso penalty. Indeed, similar to the above argument we obtain

$$\|X\hat{\theta}_X^L\|^2 \leq C \left(\frac{1}{2} \|\xi\|^2 + \frac{n\lambda}{\sqrt{n}} \|\tilde{\Omega}\vartheta_0\|_1 \right).$$

Since entries of $\tilde{\Omega}$ are bounded by some constant C and $\|\vartheta_0\| \leq C$, we get $\|\tilde{\Omega}\vartheta_0\|_1 \leq C\sqrt{p}\|\vartheta_0\| \leq C\sqrt{p}$. Using the fact that $\|\xi\| \leq C\sqrt{n}$ with probability at least $1 - n^{-c}$ we obtain $\|X\hat{\theta}_X^L\|^2 \leq Cn$ (we have used that $(p/n)^{1/2} \leq C$) with probability at least $1 - n^{-c}$. A lower bound on the singular value of X from Lemma 1.1 then yields $\|\hat{\theta}_X^L\| \leq C$.

Case 2: $n/p \rightarrow \kappa^{-1} \leq C_K^2$

Ridge regression

This is a relatively simpler case since we can directly analyze the closed form expression of the ridge estimator and use strong convexity. Indeed for the ridge estimator $\hat{\theta}_X^R$, we have

$$\begin{aligned} \|\hat{\theta}_X^R\| &= \|(X^T X + n\lambda\tilde{\Omega}^2)^{-1}(X^T X\vartheta_0 + X^T \xi)\| \\ &\leq (n\lambda c)^{-1} (\|X^T X\| \|\vartheta_0\| + \|X^T\| \|\xi\|). \end{aligned}$$

By lemma 1.1 we can upper bound $\|X^T X\|$ and $\|X\|$ by n and \sqrt{n} respectively. Since $\|\vartheta_0\|$ is bounded by a constant and $\|\xi\| \leq C\sqrt{n}$ with probability at least $1 - n^{-c}$ we obtain $\|\hat{\theta}_X^R\| \leq C$ with probability at least $1 - n^{-c}$ which in turn implies $\|\hat{\theta}_X^R\| \leq C$ with probability at least $1 - n^{-c}$.

Lasso

This case is more difficult because when the condition $n/p \rightarrow \kappa^{-1} > C_K^2$ is not satisfied, the bound on the smallest singular value (see Lemma 1.1) does not work. In fact, for $p > n$ the smallest singular value is zero. However, we can follow the idea outlined in [47] - ensuring that λ is larger than a certain value one can ensure that the optimizer is sparse with the number of nonzero components a small fraction of n . Restricting ourselves to such a set S , it is then enough to give a lower bound for the smallest singular value of X_S .

The following lemma shows that for large enough λ the optimizer is sparse.

Lemma 1.2: For any $c_0 < 1$ There exist a λ_0 which depends on c_0 , such that for $\lambda \geq \lambda_0$, the solution $\hat{\theta}_X^L$ is $(c_0 n)$ sparse with probability at least $1 - n^{c(\lambda, c_0)}$ where $c(\lambda, c_0)$ is a constant depending on c_0 and the tuning parameter λ .

We defer the proof of this lemma for the moment. Now we can easily derive the ℓ_2 bound for the Lasso case. By the optimality of $\hat{\theta}_X^L$, it easily follows that

$$\begin{aligned} & \frac{1}{2} \|X\hat{\theta}_X^L - \xi\|^2 + \frac{\sqrt{n}\lambda}{2} \|\tilde{\Omega}(\hat{\theta}_X^L + \vartheta_0)\|_1 \\ & \leq \frac{1}{2} \|\xi\|^2 + \frac{\sqrt{n}\lambda}{2} \|\tilde{\Omega}\vartheta_0\|_1. \end{aligned}$$

We can rewrite the inequality using ϑ instead of θ as follows:

$$\begin{aligned} & \frac{1}{2} \|X\hat{\theta}_X^L - Y\|^2 + \frac{\sqrt{n}\lambda}{2} \|\tilde{\Omega}\hat{\theta}_X^L\|_1 \\ & \leq \frac{1}{2} \|\xi\|^2 + \frac{\sqrt{n}\lambda}{2} \|\tilde{\Omega}\vartheta_0\|_1. \end{aligned}$$

Rearranging the terms it is not difficult to observe that

$$\frac{1}{2} \|X\hat{\theta}_X^L\|^2 \leq C \left(\frac{1}{2} \|\xi\|^2 + \frac{\sqrt{n}\lambda}{2} \|\tilde{\Omega}\vartheta_0\|_1 + \|Y\|^2 \right).$$

The right hand side is less than equal to n due to sub-Gaussianity assumptions. Indeed $\|\xi\|^2 \leq Cn$ with high probability and the same holds for the term $\|X\vartheta_0\|^2$ using Lemma 1.1 and boundedness of $\|\vartheta_0\|$. These two bounds can be used to show that $\|Y\|^2 \leq Cn$ with high probability. Also, using the boundedness of the entries of $\tilde{\Omega}$ and Cauchy-Schwarz inequality we have $\frac{\sqrt{n}\lambda}{2} \|\tilde{\Omega}\vartheta_0\|_1 \leq C\sqrt{np}\|\vartheta_0\| \leq Cn$. From these bounds the claimed bound on the RHS follows. On the other hand the LHS is greater than $cn\|\hat{\theta}_X^L\|^2$ with high probability by Lemma 1.2. Thus we obtain $\|\hat{\theta}_X^L\| \leq C$ for some constant c with high probability. Since $\|\vartheta_0\| \leq C$, we obtain that $\|\hat{\theta}_X^L\| \leq C$ with high probability. \square

Proof of Lemma 7.2:

Case 1: $n/p \rightarrow \psi^{-1} > C_0 > C_K^2$

Ridge Regression

Let $A = X/\sqrt{n}$, $w = \sqrt{n}\theta$, $\hat{w}_A^R = \sqrt{n}\hat{\theta}_X^R$. We also denote $\mu_0 = \sqrt{n}\vartheta_0$ and write \hat{w}_A^R as \hat{w} for the sake of brevity. Thus it is enough to show that $\|\hat{w}\|_\infty \leq \sqrt{d \log n}$ with high probability. We consider the column-leave- d -out version of the problem. Let S be a given subset of $\{1, \dots, p\}$.

$$\hat{w}^{(S)} = \arg \min_{w \in \mathbb{R}^p, w_S=0} \frac{1}{2} \|Aw - \xi\|^2 + \frac{\lambda}{2} \|\tilde{\Omega}(\mu_0 + w)\|^2.$$

Since \hat{w} is the optimal solution we obtain the following set of inequalities

$$\begin{aligned} 0 & \leq \frac{1}{2} \|A\hat{w}^{(S)} - \xi\|^2 + \frac{\lambda}{2} \|\tilde{\Omega}(\mu_0 + \hat{w}^{(S)})\|^2 \\ & \quad - \frac{1}{2} \|A\hat{w} - \xi\|^2 - \frac{\lambda}{2} \|\tilde{\Omega}(\mu_0 + \hat{w})\|^2 \\ & = -\frac{1}{2} \|A(\hat{w}^{(S)} - \hat{w})\|^2 \\ & \quad + \langle A(\hat{w}^{(S)} - \hat{w}), A\hat{w}^{(S)} - \xi \rangle \\ & \quad + \frac{\lambda}{2} \|\tilde{\Omega}(\mu_0 + \hat{w}^{(S)})\|^2 - \frac{\lambda}{2} \|\tilde{\Omega}(\mu_0 + \hat{w})\|^2. \end{aligned}$$

We have the following decomposition

$$\begin{aligned} & \langle A(\hat{w}^{(S)} - \hat{w}), A\hat{w}^{(S)} - \xi \rangle \\ & = -\langle A_S \hat{w}_S, (A_{-S} \hat{w}_{-S}^{(S)} - \xi) \rangle \\ & \quad + \langle A_{-S}(\hat{w}_{-S}^{(S)} - \hat{w}_{-S}), (A_{-S} \hat{w}_{-S}^{(S)} - \xi) \rangle \\ & = -\langle A_S \hat{w}_S, (A_{-S} \hat{w}_{-S}^{(S)} - \xi) \rangle \\ & \quad - \lambda \langle \hat{w}_{-S}^{(S)} - \hat{w}_{-S}, \tilde{\Omega}_{-S}^2(\mu_{0,-S} + \hat{w}_{-S}^{(S)}) \rangle, \end{aligned}$$

where the last equality uses the KKT condition $A_{-S}^T (A_{-S} \hat{w}_{-S}^{(S)} - \xi) = -\lambda \tilde{\Omega}_{-S}^2(\mu_{0,-S} + \hat{w}_{-S}^{(S)})$. Thus we have

$$\begin{aligned} 0 & \leq -\frac{1}{2} \|A(\hat{w}^{(S)} - \hat{w})\|^2 - \langle A_S \hat{w}_S, (A_{-S} \hat{w}_{-S}^{(S)} - \xi) \rangle \\ & \quad - \lambda \langle \hat{w}_{-S}^{(S)} - \hat{w}_{-S}, \tilde{\Omega}_{-S}^2(\mu_{0,-S} + \hat{w}_{-S}^{(S)}) \rangle \\ & \quad + \frac{\lambda}{2} \|\tilde{\Omega}(\mu_0 + \hat{w}^{(S)})\|^2 - \frac{\lambda}{2} \|\tilde{\Omega}(\mu_0 + \hat{w})\|^2 \\ & \leq -\frac{1}{2} \|A(\hat{w}^{(S)} - \hat{w})\|^2 - \langle A_S \hat{w}_S, (A_{-S} \hat{w}_{-S}^{(S)} - \xi) \rangle \\ & \quad + \frac{\lambda}{2} \|\tilde{\Omega}_S \mu_{0,S}\|^2 - \frac{\lambda}{2} \|\tilde{\Omega}_S(\mu_{0,S} + \hat{w}_S)\|^2, \end{aligned}$$

where the last inequality follows from the convexity of the function $\|\cdot\|^2$. Thus we obtain

$$\begin{aligned} & \frac{1}{2} \|A(\hat{w}^{(S)} - \hat{w})\|^2 + \frac{\lambda}{2} \|\tilde{\Omega}_S \hat{w}_S\|^2 \\ & \leq \|\hat{w}_S\| (\|A_S^T (A_{-S} \hat{w}_{-S}^{(S)} - \xi)\| + \lambda \|\tilde{\Omega}_S^2 \mu_{0,S}\|). \end{aligned}$$

From this relation, it easily follows that for some constant c

$$\begin{aligned} c \|\hat{w}_S\| & \leq 2\lambda^{-1} (\|A_S^T (A_{-S} \hat{w}_{-S}^{(S)} - \xi)\| + \lambda \|\tilde{\Omega}_S^2 \mu_{0,S}\|) \\ & \leq 2\lambda^{-1} (\|A_S^T A_{-S} \hat{w}_{-S}^{(S)}\| + \|A_S^T \xi\| \\ & \quad + \lambda \|\tilde{\Omega}_S^2 \mu_{0,S}\|). \end{aligned} \tag{A.2}$$

By the optimality of $\hat{w}^{(S)}$, we have

$$\begin{aligned} & \frac{1}{2} \|A_{-S} \hat{w}_{-S}^{(S)} - \xi\|^2 + \frac{\lambda}{2} \|\tilde{\Omega}(\mu_0 + \hat{w}_{-S}^{(S)})\|^2 \\ & \leq \frac{1}{2} \|\xi\|^2 + \frac{\lambda}{2} \|\tilde{\Omega} \mu_0\|^2, \end{aligned}$$

whence $\|A_{-S} \hat{w}_{-S}^{(S)}\|^2 \leq C \left(\frac{1}{2} \|\xi\|^2 + \frac{\lambda}{2} \|\tilde{\Omega} \mu_0\|^2 \right)$. Since $\|\xi\|^2 \leq Cn$ with probability at least $1 - n^{-c}$ and $\|\tilde{\Omega} \mu_0\|^2 \leq Cn$, we obtain $\|A_{-S} \hat{w}_{-S}^{(S)}\|^2 \leq Cn$ with probability at least $1 - n^{-c}$.

Now for a single coordinate s there exist a set S (with $|S| \leq d$) such that $s \in S$ and the columns of A_S and the

columns of A_{-S} are independent. Let $i \in S$. Using the facts that the elements of A_i are independent sub-Gaussian and that A_i and $A_{-S}\hat{w}_{-S}^{(S)}$ are independent, we have

$$\|A_S^T A_{-S} \hat{w}_{-S}^{(S)}\| \leq \sqrt{d} \max_{i \in S} |A_i^T A_{-S} \hat{w}_{-S}^{(S)}| \leq C \sqrt{d \log n} \quad (\text{A.3})$$

holds with probability at least $1 - n^{-c}$ conditioned on the event $\|A_{-S}\hat{w}_{-S}^{(S)}\|^2 \leq Cn$ which itself holds with probability $1 - n^{-c}$.

Since the elements of A_i are i.i.d. and A_i and ξ are independent, $|A_i^T \xi| \leq C \sqrt{\log n}$ holds with probability at least $1 - n^{-c}$. Indeed conditioning on the event $\|\xi\| \leq C \sqrt{n}$ the result follows from the sub-Gaussian concentration and the event itself holds with probability at least $1 - n^{-c}$. Similar to the bounds obtained for $\|A_S^T A_{-S} \hat{w}_{-S}^{(S)}\|$ we obtain $\|A_S^T \xi\| \leq C \sqrt{d \log n}$ with probability at least $1 - n^{-c}$. Finally since the entries of the signal μ_0 are bounded (recall that the entries of ϑ_0 are of the order $O(p^{-1/2})$ by Assumption 4), using (A.2) we have the following inequality

$$|\hat{w}_s| \leq \|\hat{w}_S\| \leq C \sqrt{d \log n},$$

with probability at least $1 - n^{-c}$.

Also this event holds uniformly in s since we can choose $c > 1$ and use the union bound. Hence

$$\|\hat{w}^R\|_\infty \leq C \sqrt{d \log n},$$

with probability at least $1 - n^{-c}$.

Lasso

Once again we start with the leave-d-out estimator. Let S be a given subset of $\{1, \dots, p\}$ and define:

$$\hat{w}^{(S)} = \arg \min_{w \in \mathbb{R}^p, w_S=0} \frac{1}{2} \|Aw - \xi\|^2 + \lambda \|\tilde{\Omega}(\mu_0 + w)\|_1.$$

The KKT condition is given by

$$A_{-S}^T (A_{-S} \hat{w}_{-S}^{(S)} - \xi) + \lambda \tilde{\Omega}_{-S} v_{-S} = 0,$$

where v_{-S} is the subgradient of $\|\cdot\|_1$ at $\tilde{\Omega}(\mu_0 + \hat{w}^{(S)})$. We write the risk $R(Y, A, \vartheta)$ as a function of w i.e. $R(w)$ suppressing the other variables. Denoting \hat{w} as the Lasso optimizer we have

$$\begin{aligned} & R(0, \hat{w}_{-S}) \\ &= \frac{1}{2} \|A_{-S} \hat{w}_{-S} - \xi\|^2 + \lambda \|\tilde{\Omega}(\hat{w}_{-S} + \mu_0)\|_1 \\ &= \frac{1}{2} \|A_{-S}(\hat{w}_{-S} - \hat{w}_{-S}^{(S)})\|^2 + \frac{1}{2} \|A_{-S} \hat{w}_{-S}^{(S)} - \xi\|^2 \\ & \quad + \langle A_{-S}(\hat{w}_{-S} - \hat{w}_{-S}^{(S)}), A_{-S} \hat{w}_{-S}^{(S)} - \xi \rangle \\ & \quad + \lambda \|\tilde{\Omega}(\hat{w}_{-S}^{(S)} + \mu_0)\|_1 + \lambda \|\tilde{\Omega}(\hat{w}_{-S} + \mu_0)\|_1 \\ & \quad - \lambda \|\tilde{\Omega}(\hat{w}_{-S}^{(S)} + \mu_0)\|_1 \\ &= R(0, \hat{w}_{-S}^{(S)}) + \frac{1}{2} \|A_{-S}(\hat{w}_{-S} - \hat{w}_{-S}^{(S)})\|^2 \\ & \quad + \langle A_{-S}(\hat{w}_{-S} - \hat{w}_{-S}^{(S)}), A_{-S} \hat{w}_{-S}^{(S)} - \xi \rangle \\ & \quad + \lambda \|\tilde{\Omega}(\hat{w}_{-S} + \mu_0)\|_1 - \lambda \|\tilde{\Omega}(\hat{w}_{-S}^{(S)} + \mu_0)\|_1 \\ &= R(0, \hat{w}_{-S}^{(S)}) + \frac{1}{2} \|A_{-S}(\hat{w}_{-S} - \hat{w}_{-S}^{(S)})\|^2 \end{aligned}$$

$$\begin{aligned} & - \langle \hat{w}_{-S} - \hat{w}_{-S}^{(S)}, \lambda \tilde{\Omega}_{-S} v_{-S} \rangle \\ & + \lambda \|\tilde{\Omega}(\hat{w}_{-S} + \mu_0)\|_1 - \lambda \|\tilde{\Omega}(\hat{w}_{-S}^{(S)} + \mu_0)\|_1 \\ & \geq R(0, \hat{w}_{-S}^{(S)}) + \frac{1}{2} \|A_{-S}(\hat{w}_{-S} - \hat{w}_{-S}^{(S)})\|^2. \end{aligned}$$

The fourth equality follows from the KKT condition and the last inequality from the convexity of the function $\|\cdot\|_1$. Define $f_1(x) := R((x, \hat{w}_{-S}))$ and $f_2(x) := R((x, \hat{w}_{-S}^{(S)}))$. Then we have shown

$$\begin{aligned} & \frac{1}{2} \|A_{-S}(\hat{w}_{-S} - \hat{w}_{-S}^{(S)})\|^2 \\ & \leq f_1(0) - f_2(0) \\ & \leq (f_1(0) - \min_x f_1(x)) - (f_2(0) - \min_x f_2(x)). \end{aligned} \quad (\text{A.4})$$

The last inequality follows from the fact

$$\min_x f_1(x) = R(\hat{w}) \leq \min_x R((x, \hat{w}_{-S}^{(S)})) = \min_x f_2(x).$$

We expand $f_1(x)$ as follows:

$$\begin{aligned} & f_1(x) \\ &= \frac{1}{2} \|A_S x + A_{-S} \hat{w}_{-S} - \xi\|^2 + \lambda \|\tilde{\Omega}_S(x + \mu_{0,S})\|_1 \\ & \quad + \lambda \|\tilde{\Omega}_{-S}(\hat{w}_{-S} + \mu_{0,-S})\|_1 \\ &= \frac{1}{2} \|A_{-S} \hat{w}_{-S} - \xi\|^2 + \frac{1}{2} \|A_S x\|^2 \\ & \quad + \langle A_S x, A_{-S} \hat{w}_{-S} - \xi \rangle \\ & \quad + \lambda \|\tilde{\Omega}_S(x + \mu_{0,S})\|_1 + \lambda \|\tilde{\Omega}_{-S}(\hat{w}_{-S} + \mu_{0,-S})\|_1 \\ & \quad + \lambda \|\tilde{\Omega}_S \mu_{0,S}\|_1 - \lambda \|\tilde{\Omega}_S \mu_{0,S}\|_1 \\ &= f_1(0) + \frac{1}{2} \|A_S x\|^2 + \langle A_S x, A_{-S} \hat{w}_{-S} - \xi \rangle \\ & \quad + \lambda \|\tilde{\Omega}_S(x + \mu_{0,S})\|_1 - \lambda \|\tilde{\Omega}_S \mu_{0,S}\|_1. \end{aligned}$$

Let $y_1 = A_{-S} \hat{w}_{-S} - \xi$ and $y_2 = A_{-S} \hat{w}_{-S}^{(S)} - \xi$. Also define

$$g(y) = \min_x \frac{1}{2} \|A_S x\|^2 + \langle A_S x, y \rangle + \lambda \|\tilde{\Omega}_S(x + \mu_{0,S})\|_1.$$

Then we observe that

$$\begin{aligned} & f_1(0) - f_2(0) \\ & \leq (f_1(0) - \min_x f_1(x)) - (f_2(0) - \min_x f_2(x)) \\ & = g(y_2) - g(y_1). \end{aligned}$$

Further we define the following functions:

$$\begin{aligned} h(x) &= \frac{1}{2} \|A_S x\|^2 + \langle A_S x, y \rangle \\ b(x) &= \lambda \|\tilde{\Omega}_S(x + \mu_{0,S})\|_1. \end{aligned}$$

Thus we have,

$$\begin{aligned} g(y) &= \min_x h(x) + b(x) \\ &= \sup_x -h^*(x) - b^*(-x), \end{aligned}$$

where $h^*(x)$ and $g^*(x)$ are Fenchel duals of $h(x)$ and $g(x)$ respectively and the last equality is obtained using Fenchel duality theorem. We note that with probability at least $1 - n^{-c}$ $s_{\min}(A_S) \geq c$ and hence $A_S^T A_S$ is invertible and we have the following lemma:

Lemma 1.3: For the functions $h(x)$ and $b(x)$ defined above, we have the following Fenchel duals:

$$\begin{aligned} h^*(x) &= \frac{1}{2}(x - A_S^T y)^T (A_S^T A_S)^{-1} (x - A_S^T y), \\ b^*(x) &= \begin{cases} -\langle \mu_{0,S}, x \rangle & \text{if } \|\tilde{\Omega}_S^{-1} x\|_\infty \leq \lambda \\ \infty & \text{otherwise.} \end{cases} \end{aligned} \quad (\text{A.5})$$

It follows that

$$\begin{aligned} g(y) &= \sup_{\|\tilde{\Omega}_S^{-1} x\|_\infty \leq \lambda} -\frac{1}{2}(x - A_S^T y)^T (A_S^T A_S)^{-1} (x - A_S^T y) \\ &\quad - \langle \mu_{0,S}, x \rangle, \end{aligned}$$

and consequently defining $P_S = A_S (A_S^T A_S)^{-1} A_S^T$ and using (A.4) we obtain the following set of inequalities,

$$\begin{aligned} &\frac{1}{2} \|A_{-S}(\hat{w}_{-S} - \hat{w}_{-S}^{(S)})\|^2 \\ &\leq |g(y_2) - g(y_1)| \\ &\leq \sup_{\|\tilde{\Omega}_S^{-1} x\|_\infty \leq \lambda} |x^T (A_S^T A_S)^{-1} A_S^T (y_1 - y_2)| \\ &\quad + \frac{1}{2} |y_1^T P_S y_1 - y_2^T P_S y_2| \\ &\leq \sup_{\|\tilde{\Omega}_S^{-1} x\|_\infty \leq \lambda} \|x\| \|(A_S^T A_S)^{-1} A_S^T (y_1 - y_2)\| \\ &\quad + \frac{1}{2} (y_1 - y_2)^T P_S (y_1 - y_2) + |y_2^T P_S (y_1 - y_2)|. \end{aligned}$$

Recall that $y_1 - y_2 = A_{-S}(\hat{w}_{-S} - \hat{w}_{-S}^{(S)})$; hence we obtain

$(y_1 - y_2)^T P_S (y_1 - y_2) = \|P_S A_{-S}(\hat{w}_{-S} - \hat{w}_{-S}^{(S)})\|^2$. For $\|\tilde{\Omega}_S^{-1} x\|_\infty \leq \lambda$ we have the following inequality

$$\|x\| \leq \|x\|_\infty \sqrt{d} \leq C \|\tilde{\Omega}_S^{-1} x\|_\infty \sqrt{d} \leq C \lambda \sqrt{d}$$

and consequently we obtain,

$$\begin{aligned} &\frac{1}{2} \|P_S^\perp A_{-S}(\hat{w}_{-S} - \hat{w}_{-S}^{(S)})\|^2 \\ &\leq C \lambda \sqrt{d} \|(A_S^T A_S)^{-1} A_S^T A_{-S}(\hat{w}_{-S} - \hat{w}_{-S}^{(S)})\| \\ &\quad + \|(A_{-S} \hat{w}_{-S}^{(S)} - \xi)^T P_S A_{-S}(\hat{w}_{-S} - \hat{w}_{-S}^{(S)})\| \\ &\leq C \lambda \sqrt{d} \|(A_S^T A_S)^{-1} A_S^T A_{-S}(\hat{w}_{-S} - \hat{w}_{-S}^{(S)})\| \\ &\quad + \|(A_{-S} \hat{w}_{-S}^{(S)} - \xi)^T A_S (A_S^T A_S)^{-1} A_S^T A_{-S}(\hat{w}_{-S} - \hat{w}_{-S}^{(S)})\| \\ &\leq (C \lambda \sqrt{d} \|(A_S^T A_S)^{-1}\| \|A_S^T\| \|A_{-S}\| \\ &\quad + \|A_{-S}^T\| \|A_S\| \|(A_S^T A_S)^{-1}\| \|A_S^T (A_{-S} \hat{w}_{-S}^{(S)} - \xi)\|) \|\hat{w}_{-S} - \hat{w}_{-S}^{(S)}\|. \end{aligned}$$

By the same argument as equation (A.3) we have $\|A_S^T A_{-S} \hat{w}_{-S}^{(S)}\| \leq C \sqrt{d \log n}$ with probability at least $1 - n^{-c}$ and by independence of A_S and ξ , $\|A_S^T \xi\| \leq C \sqrt{d \log n}$ with probability at least $1 - n^{-c}$. Hence $\|A_S^T (A_{-S} \hat{w}_{-S}^{(S)} - \xi)\| \leq C \sqrt{d \log n}$ with probability at least $1 - n^{-c}$. By Lemma 1.1 the terms $\|A_S\|, \|A_{-S}\|, \|(A_S^T A_S)^{-1}\|$ are bounded by some

constant C with probability at least $1 - n^{-c}$. Hence the following holds

$$\frac{1}{2} \|P_S^\perp A_{-S}(\hat{w}_{-S} - \hat{w}_{-S}^{(S)})\|^2 \leq C \sqrt{d \log n} \|\hat{w}_{-S} - \hat{w}_{-S}^{(S)}\|, \quad (\text{A.6})$$

with probability at least $1 - n^{-c}$. Now

$$\begin{aligned} &P_S^\perp (A_{-S}(\hat{w}_{-S} - \hat{w}_{-S}^{(S)})) \\ &= A_S (-(A_S^T A_S)^{-1} A_S^T A_{-S}(\hat{w}_{-S} - \hat{w}_{-S}^{(S)})) \\ &\quad + A_{-S}(\hat{w}_{-S} - \hat{w}_{-S}^{(S)}) = A\bar{w}, \end{aligned}$$

where

$$\begin{aligned} \bar{w}^T &= ((-(A_S^T A_S)^{-1} A_S^T A_{-S}(\hat{w}_{-S} - \hat{w}_{-S}^{(S)}))^T, \\ &\quad (\hat{w}_{-S} - \hat{w}_{-S}^{(S)})^T). \end{aligned}$$

By Lemma 1.1 and the condition $n/p > C_K^2$ we have $\|A\bar{w}\|^2 \geq c\|\bar{w}\|^2$ (for some $c > 0$). We also note that $\|\bar{w}\|^2 \geq \|\hat{w}_{-S} - \hat{w}_{-S}^{(S)}\|^2$ and hence,

$$\begin{aligned} c\|\hat{w}_{-S} - \hat{w}_{-S}^{(S)}\|^2 &\leq C \sqrt{d \log n} \|\hat{w}_{-S} - \hat{w}_{-S}^{(S)}\|, \\ \|\hat{w}_{-S} - \hat{w}_{-S}^{(S)}\| &\leq c^{-1} C \sqrt{d \log n}. \end{aligned} \quad (\text{A.7})$$

By the KKT condition

$$A^T A w = A^T \xi - \lambda \tilde{\Omega} v,$$

for subgradient v . Arguing that $\|A^T \xi\|_\infty$ and $\|\tilde{\Omega} v\|_\infty$ are bounded by $C \sqrt{\log n}$ with probability at least $1 - n^{-c}$ we obtain

$$\begin{aligned} \|(A^T A \hat{w})_S\| &= \|A_S^T A_S \hat{w}_S + A_S^T A_{-S} \hat{w}_{-S}\| \\ &\leq C \sqrt{d \log n}. \end{aligned}$$

Rearranging terms and using triangle inequality then yields,

$$\|A_S^T A_S \hat{w}_S\| \leq C(\sqrt{d \log p} + \|A_S^T A_{-S} \hat{w}_{-S}\|).$$

Noting that $\|A_S^T A_S\| \geq c$ holds with probability at least $1 - n^{-c}$ we have the following decomposition:

$$\begin{aligned} \|\hat{w}_S\| &\leq C(\|A_S^T A_{-S} \hat{w}_{-S}\| + \sqrt{d \log n}) \\ &\leq C(\|A_S^T A_{-S} \hat{w}_{-S}^{(S)}\| \\ &\quad + \|A_S\| \|A_{-S}\| \|\hat{w}_{-S} - \hat{w}_{-S}^{(S)}\| \\ &\quad + \sqrt{d \log n}). \end{aligned}$$

The term $\|A_S^T A_{-S} \hat{w}_{-S}^{(S)}\|$ can be bounded by $C \sqrt{d \log n}$ by the same argument as equation (A.3). Also $\|A_S\|, \|A_{-S}\|$ are bounded by some constant C with probability at least $1 - n^{-c}$ by applying Lemma 1.1. These facts together with (A.7) yields

$$|\hat{w}|_s \leq \|\hat{w}_S\| \leq C \sqrt{d \log n},$$

with probability at least $1 - n^{-c}$ and using a union bound we obtain

$$\|\hat{w}^L\|_\infty \leq C \sqrt{d \log n},$$

with probability at least $1 - n^{-c}$.

Case 2: $n/p \rightarrow \kappa^{-1} \leq C_K^2$

We note that for the ℓ_∞ bound in the ridge regression the lower bound of the singular value from Lemma 1.1 was not

used, hence in this case the ℓ_∞ bound proof is identical to the one used in Case 1.

For the ℓ_∞ in Lasso we note the lower bound on the singular from Lemma 1.1 was only used to give lower bound to $A\bar{w}$ (see the equation after A.6). However in this case one can use easily show that for $\lambda \geq \lambda_0$ $\hat{w}_{-S} - \hat{w}_S^{(S)}$ and hence \bar{w} is $2c_0n$ sparse (see Lemma 1.2). The rest of the proof is same as the Lasso case of Lemma 7.2. \square

Proof of Lemma 1.2: This follows from the proof of Lemma 6.3 of [47] with certain modification due to the dependence. For the sake of completeness we give the entire proof.

From the KKT condition we can write

$$A_i^T(A\hat{w} - \xi) + \lambda(\tilde{\Omega}v)_i = 0, \quad (\text{A.8})$$

with v be the subgradient of $\|\cdot\|_1$ at $\tilde{\Omega}(\hat{w} + \mu_0)$. Let S_+ be the set such that for $i \in S_+$, $\hat{\mu}_i = (\hat{w} + \mu_0)_i > 0$ (and similarly define S_-). We will show that both S_+ and S_- are bounded by c_0n with probability at least $1 - n^{-c}$ which ensures that the Lasso solution is $2c_0n$ sparse with probability at least $1 - n^{-c}$.

Note that the KKT condition (A.8) can be used to show

$$\begin{aligned} & \sum_{i \in S_+} z_i^2 + \lambda^2 \sum_{i \in S_+} \omega_i^2 - 2\lambda \sum_{i \in S_+} z_i \omega_i \\ & \leq \|A\hat{w}\|^2 \|A_{S_+} A_{S_+}^T\|, \\ & \quad c^2 \lambda^2 |S_+| - 2\lambda \sum_{i \in S_+} z_i \omega_i \\ & \leq \|A\hat{w}\|^2 \|A_{S_+} A_{S_+}^T\|, \end{aligned}$$

where $z = A^T \xi$ and we have used the lower bound on ω_i . Thus with probability at least $1 - n^{-c}$ the following event holds

$$E := \{c^2 \lambda^2 |S_+| \leq 2\lambda \sum_{i \in S_+} z_i \omega_i + Cn\}.$$

We have used the fact that $\|A\hat{w}\| \leq C\sqrt{n}$ with probability at least $1 - n^{-c}$ and $\|A_{S_+} A_{S_+}^T\| \leq C$ with the same probability. If $|S_+| \geq c_0n$, we have $c^2 \lambda^2 |S_+| \leq 2\lambda \sum_{i \in S_+} z_i \omega_i + (C/c_0) |S_+|$ so that if λ is chosen greater than $\lambda_0 = (2C/c^2 c_0)^{1/2}$, we have

$$c^2 \lambda^2 |S_+|/2 \leq 2\lambda \sum_{i \in S_+} z_i \omega_i.$$

To show that this event hold with vanishing probability we first write $\sum_{i \in S_+} z_i \omega_i = \xi^T A e_s$ where e_s is the vector whose j^{th} entry is $\omega_j \mathbb{1}\{j \in S_+\}$. We observe that $(A e_s)_i$ is sub-Gaussian with sub-Gaussian factor $Cc^2 s/n$ where $s = |S_+|$, since $A = X/\sqrt{n}$ and X has isotropic sub-Gaussian rows. Hence conditioned on ξ the random variable $\xi^T A e_s$ has sub-Gaussian factor $Cc^2 \|\xi\|^2 s/n$ since ξ is independent of A . We define the event $F := \{\|\xi\| \leq C\sqrt{n}\}$. Since $\|\xi\| \leq C\sqrt{n}$ with probability $1 - n^{-c}$, for a given subset S of size s we have

$$\begin{aligned} P(\{c^2 \lambda^2 s/2 \leq 2\lambda \sum_{i \in S_+} z_i \omega_i\} \cap F) \\ \leq \exp(-c^4 \lambda^2 (s/4)^2 / 2Cs) \end{aligned}$$

where we have simplified the constants. Finally we obtain,

$$P(\{|S_+| \geq c_0n\} \cap E \cap F)$$

$$\begin{aligned} & \leq P(\{\lambda c^2 |S_+|/4 \leq |\sum_{i \in S_+} z_i \omega_i|\} \cap F) \\ & \leq \sum_{s=c_0n}^p \sum_{S:|S|=s} \exp(-c^4 \lambda^2 (s/4)^2 / 2Cs) \\ & \leq C \exp(-Cc_0 \lambda^2 n), \end{aligned}$$

where we have simplified all the constants except c_0 . We can remove the conditioning on E and F since both holds with probability at least $1 - n^{-c}$.

Similar bounds hold for S_- and we have shown (redefining c_0) that for $\lambda \geq \lambda_0$, the Lasso solution is c_0n sparse with probability at least $1 - n^{-c}$. \square

Proof of Lemma 1.3: We start with the Fenchel dual for the function h ,

$$h^*(x) = \sup_z \langle z, x \rangle - \frac{1}{2} \|A_S z\|^2 - \langle A_S z, y \rangle.$$

It can be easily verified that the maximum is attained at $z = (A_S^T A_S)^{-1}(x - A_S^T y)$ and plugging in the value we obtain

$$h^*(x) = \frac{1}{2} (x - A_S^T y)^T (A_S^T A_S)^{-1} (x - A_S^T y).$$

To compute $b^*(x)$ we start with defining the function $\bar{b}(x) = \lambda \omega |x + y|$ for scalars x and y . Thus we have,

$$\bar{b}^*(x) = \sup_z xz - \lambda \omega |z + y|.$$

It can be easily observed that $\bar{b}^*(x) = \infty$ unless $|x| \leq \lambda \omega$.

Further define the function

$$\alpha(x, z) = xz - \lambda \omega |z + y|.$$

If $z \geq -y$, the function $\alpha(x, z) = z(x - \lambda \omega) - \lambda y$ attains maximum value at $z = -y$ as $x - \lambda \omega \leq 0$. Similarly if $z \leq -y$, $\alpha(x, z) = z(x + \lambda \omega) + \lambda y$ attains maximum value at $z = -y$ as $x + \lambda \omega \geq 0$. Hence

$$\bar{b}^*(x) = \begin{cases} \sup_z \alpha(x, z) = -xy, & \text{if } |x| \leq \lambda \omega \\ \infty, \text{ otherwise} . \end{cases}$$

Now we consider the function

$$b^*(x) = \sup_z \langle z, x \rangle - \lambda \|\tilde{\Omega}_S(z + \mu_{0,S})\|_1$$

and note that the optimization problem in the RHS is separable in its arguments. Hence we have

$$b^*(x) = \begin{cases} -\langle \mu_{0,S}, x \rangle, & \text{if } |x_i/\omega_i| \leq \lambda \text{ for all } i \in S \\ \infty, \text{ otherwise} , \end{cases}$$

which was to be shown. \square

C. Universality of the Optimizer

Proof of Lemma 7.3:

We only prove the result for the ridge case, since the Lasso case follows by suitably modifying (to account for the term $\tilde{\Omega}$) the proof of Theorem 3.8 of [47] together with the statements from [84]. Recall that G is the design matrix with i.i.d. Gaussian entries. Denoting $\bar{G} = G/\sqrt{n}$ problem, we consider the following functions:

$$r(w, u) = \frac{1}{n} u^T \bar{G} w - \frac{1}{n} u^T \xi - \frac{1}{2n} \|u\|^2$$

$$\ell(w, u) = -\frac{1}{n^{3/2}} \|u\| g^T w + \frac{\|w\| h^T u}{n^{3/2}} - \frac{1}{n} u^T \xi - \frac{1}{2n} \|u\|^2 + \frac{\lambda}{2n} (\|\tilde{\Omega}(w + \mu_0)\|^2 - \|\tilde{\Omega}\mu_0\|^2),$$

where $u \in \mathbb{R}^n$ and h and g are independent random vectors with distributions $h \sim N(0, I_n)$, $g \sim N(0, I_p)$ respectively. We denote by $R(w) = \max_u h(w, u)$ and $L(w) = \max_u \ell(w, u)$ the ridge cost and the associated Gordon cost respectively.

Note that using $w = \sqrt{n}\theta$, $R(w) = R^K(\theta, G)$ for $K = R$ and we work with the notation $R(w)$ since the Gaussian design and the ridge setup is clear from the context. Our goal is to demonstrate that the optimal value of $R(w)$ i.e. $\min_w R(w)$ is close to ψ^R . We will use CGMT to show that $\min_w R(w)$ is well approximated by $\min_w L(w)$ and the latter is close to the desired constant ψ^R (a constant that we will define shortly). The proof proceeds by showing that $\min_w L(w)$ is well approximated by the value at saddle point of a function $\psi_n(\beta, \gamma)$, which is further approximated by the value at saddle point of its non-stochastic version $\psi(\beta, \gamma)$. Denoting the saddle point of the latter by (β_*, γ_*) , ψ^R will be defined as $\psi(\beta_*, \gamma_*)$. We start by defining ψ_n and ψ .

Define the function $\psi_n(\beta, \gamma)$ as follows:

$$\begin{aligned} \psi_n(\beta, \gamma) &= \left(\frac{\sigma^2}{\gamma} + \gamma \right) \frac{\beta}{2} - \frac{\beta^2}{2} \\ &\quad - \frac{1}{n/p} \frac{1}{p} \sum_{i=1}^p \frac{\gamma \lambda^2 \omega_i^4}{2(\beta + \lambda \omega_i^2 \gamma)} \left(\mu_{0,i} - \frac{\beta g_i}{\lambda \omega_i^2} \right)^2. \end{aligned} \quad (\text{A.9})$$

The above equation can be written compactly as

$$\begin{aligned} \psi_n(\beta, \gamma) &= \left(\frac{\sigma^2}{\gamma} + \gamma \right) \frac{\beta}{2} - \frac{\beta^2}{2} \\ &\quad - \frac{1}{n/p} E_{P_n} \left[\frac{\gamma \lambda^2 \Omega^4}{2(\beta + \lambda \Omega^2 \gamma)} \left(M - \frac{\beta \Gamma}{\lambda \Omega^2} \right)^2 \right], \end{aligned} \quad (\text{A.10})$$

where the expectation is taken with respect to the law P_n of the 3-tuple of random variables $(M, \Omega, \Gamma) \sim \frac{1}{p} \sum_{i=1}^p \delta_{\mu_{0,i}, \omega_i, g_i}$.

Also define the function $\psi(\beta, \gamma)$ as follows:

$$\begin{aligned} \psi(\beta, \gamma) &= \left(\frac{\sigma^2}{\gamma} + \gamma \right) \frac{\beta}{2} - \frac{\beta^2}{2} \\ &\quad - \frac{1}{n/p} E_{Q_n} \left[\frac{\gamma \lambda^2 \Omega^4}{2(\beta + \lambda \Omega^2 \gamma)} \left(M - \frac{\beta \Gamma}{\lambda \Omega^2} \right)^2 \right], \end{aligned} \quad (\text{A.11})$$

where the expectation is taken with respect to the law Q_n of the 3-tuple of random variables $(M, \Omega, \Gamma) \sim (\frac{1}{p} \sum_{i=1}^p \delta_{\mu_{0,i}, \omega_i}) \otimes N(0, 1)$. Taking the derivates of the function ψ_n with respect to β and γ one obtains the following fixed point equations for the saddle points:

$$\begin{aligned} \gamma_{*,n}^2 &= \sigma^2 + \frac{1}{n/p} E_{P_n} \left[\eta_2 \left(M + \gamma_{*,n} \Gamma; \frac{\gamma_{*,n} \lambda \Omega^2}{\beta_{*,n}} \right) \right. \\ &\quad \left. - M \right]^2, \end{aligned} \quad (\text{A.12})$$

$$\beta_{*,n} = \gamma_{*,n} - \frac{1}{n/p} E_{P_n} \left[\Gamma \cdot \left(\eta_2 \left(M + \gamma_{*,n} \Gamma; \frac{\gamma_{*,n} \lambda \Omega^2}{\beta_{*,n}} \right) \right. \right. \\ &\quad \left. \left. - M \right) \right]. \quad (\text{A.13})$$

We can compare (A.12) and (A.13) to (III.4) and see that the latter easily follows from the previous two by using Gaussian integration by parts. Similarly from equation (A.11) we get the fixed point equations for the limiting case as follows:

$$\gamma_*^2 = \sigma^2 + \frac{1}{n/p} E_{Q_n} \left[\eta_2 \left(M + \gamma_* \Gamma; \frac{\gamma_* \lambda \Omega^2}{\beta_*} \right) - M \right]^2, \quad (\text{A.14})$$

$$\beta_* = \gamma_* - \frac{1}{n/p} E_{Q_n} \left[\Gamma \cdot \left(\eta_2 \left(M + \gamma_* \Gamma; \frac{\gamma_* \lambda \Omega^2}{\beta_*} \right) \right. \right. \\ &\quad \left. \left. - M \right) \right]. \quad (\text{A.15})$$

It can be shown following the steps of Proposition 5.2 of [47] that there exists a unique solution to the system formed by (A.12) and (A.13). Further the solutions $\beta_{*,n}$ and $\gamma_{*,n}$ lie in some compact subset of \mathbb{R}^+ bounded away from 0 with probability at least $1 - n^{-c}$. Similar results on the unicity and compactness of the solutions β_* and γ_* can be shown as well. Next we show that $|\beta_{*,n} - \beta_*|$ and $|\gamma_{*,n} - \gamma_*|$ converges to 0 with probability at least $1 - n^{-c}$. However because of the matrix Ω this convergence does not follow easily from [47]. We use different techniques and prove the statement in a separate lemma whose proof is technical and hence we defer for the moment

Lemma 1.4: With probability at least $1 - n^{-c}$, we have

$$|\beta_{*,n} - \beta_*| = O(\epsilon_n) \quad |\gamma_{*,n} - \gamma_*| = O(\epsilon_n).$$

for some $\epsilon_n \rightarrow 0$.

The rest of the proof follows in the lines of [47]. Since the modification involving the matrix Ω is routine we only sketch the outline and omit detailed proof. One shows (see [47, Lemma 5.2]) that with probability at least $1 - n^{-c}$ the following statements hold:

- 1) $|\psi_n(\beta_{*,n}, \gamma_{*,n}) - \psi(\beta_*, \gamma_*)| = O(\epsilon_n)$,
- 2) $|\min_w L(w) - \max_{\beta > 0} \min_{\gamma > 0} \psi_n(\beta, \gamma)| = O(\epsilon_n)$,
- 3) $|\min_w H(w) - \psi(\beta_*, \gamma_*)| = O(\epsilon_n)$.

Now for any Lipschitz function ϱ define the following set

$$S_{n,\epsilon}(\varrho) = \{w \in \mathbb{R}^p : |\varrho(w/\sqrt{n}) - E\varrho(w_*/\sqrt{n})| \geq \epsilon^{-1/2}\}.$$

Then it can be shown (see [47, Theorem 3.4] and its proof) that the following statement holds:

$$\begin{aligned} P \left(\min_w R(w) \geq \psi(\beta_*, \gamma_*) + \epsilon \right) &\leq C n^{-c}, \\ P \left(\min_{w \in S_{n,\epsilon}(\varrho)} R(w) \leq \psi(\beta_*, \gamma_*) + K^{-1} \epsilon \right) &\leq C n^{-c} \end{aligned}$$

for some $K > 0$. Recall that $\theta = w/\sqrt{n}$, $\psi^R = \psi(\beta_*, \gamma_*)$. Then letting $\delta_n = C n^{-c}$ and choosing z suitably the above inequalities imply that the statement of the Lemma 7.3 is true if we can prove Lemma 1.4. \square

Proof of Lemma 1.4:

First we adopt the method used in [58] to show that with probability at least $1 - n^{-c}$

$$\|(\beta_{*,n}, \gamma_{*,n}) - (\beta^\dagger, \gamma^\dagger)\| = O(\epsilon_n) \quad (\text{A.16})$$

holds for some $\epsilon_n \rightarrow 0$ and for some constants $\beta^\dagger, \gamma^\dagger$. Let us denote the joint law of the random variables $(\bar{M}, \bar{\Omega}, Z)$ as P , where $(\bar{M}, \bar{\Omega})$ has law μ (see Assumption 4) and Z is independently distributed as $N(0, 1)$. Note that by Assumption 4 $P_n \xrightarrow{W_2} P$ and $Q_n \xrightarrow{W_2} P$. For notational convenience we will define functions of a general 3-tuple of random variables (M, Ω, Γ) and depending on the context they will be distributed as P_n, Q_n or P .

Step 1 We write the optimizers as solutions of some functions i.e. define

$$F_n(\beta, \gamma) = \frac{1}{n/p} \left[\eta_2 \left(M + \gamma \Gamma; \frac{\gamma \lambda \Omega^2}{\beta} \right) - M \right]^2 + \sigma^2 - \gamma^2, \quad (\text{A.17})$$

$$G_n(\beta, \gamma) = \frac{1}{n/p} \left[\Gamma \cdot \left(\eta_2 \left(M + \gamma \Gamma; \frac{\gamma \lambda \Omega^2}{\beta} \right) - M \right) \right] + \beta - \gamma, \quad (\text{A.18})$$

$$F(\beta, \gamma) = \kappa \left[\eta_2 \left(M + \gamma \Gamma; \frac{\gamma \lambda \Omega^2}{\beta} \right) - M \right]^2 + \sigma^2 - \gamma^2, \quad (\text{A.19})$$

$$G(\beta, \gamma) = \kappa \left[\Gamma \cdot \left(\eta_2 \left(M + \gamma \Gamma; \frac{\gamma \lambda \Omega^2}{\beta} \right) - M \right) \right] + \beta - \gamma. \quad (\text{A.20})$$

Note that $(\beta_{*,n}, \gamma_{*,n})$ are solutions of the equations $E_{P_n} F_n(\beta, \gamma)$ and $E_{P_n} G_n(\beta, \gamma)$. We define $(\beta^\dagger, \gamma^\dagger)$ as solutions of the equation $E_P F(\beta, \gamma)$ and $E_P G(\beta, \gamma)$. It can be shown along the lines of Lemma 5.2 of [47] that both solutions are unique.

Step 2: We will show that the following uniform convergence holds in some compact set \mathcal{C} of $\mathbb{R}^+ \times \mathbb{R}^+$ bounded away from $(0, 0)$:

$$\lim_{n \rightarrow \infty} \sup_{(\gamma, \beta) \in \mathcal{C}} |E_{P_n} F_n(\beta, \gamma) - E_P F(\beta, \gamma)| = 0, \quad (\text{A.21})$$

$$\lim_{n \rightarrow \infty} \sup_{(\gamma, \beta) \in \mathcal{C}} |E_{P_n} G_n(\beta, \gamma) - E_P G(\beta, \gamma)| = 0. \quad (\text{A.22})$$

We defer the proof for the moment.

Step 3: Since the optimizers in both the finite sample case $((\beta_{*,n}, \gamma_{*,n}))$ and the limiting case $((\beta^\dagger, \gamma^\dagger))$ lie in compact sets with probability at least $1 - n^{-c}$ (see Lemma 5.2 of [47]), we can choose \mathcal{C} , so that the optimizers belong to \mathcal{C} with probability at least $1 - n^{-c}$. Now let $(\tilde{\beta}, \tilde{\gamma})$ be a limit point of the sequence $(\beta_{*,n}, \gamma_{*,n})$ then we prove that $(\tilde{\beta}, \tilde{\gamma}) = (\beta^\dagger, \gamma^\dagger)$.

Indeed,

$$\begin{aligned} & |E_P(F(\tilde{\beta}, \tilde{\gamma}))| \\ & \leq |E_P(F(\tilde{\beta}, \tilde{\gamma})) - E_P(F(\beta_{*,n}, \gamma_{*,n}))| \\ & \quad + |E_P(F(\beta_{*,n}, \gamma_{*,n})) - E_P(F(\beta_{*,n}, \gamma_{*,n}))|, \end{aligned}$$

since $E_{P_n}(F_n(\beta_{*,n}, \gamma_{*,n})) = 0$. Now from the continuity of $E_P(F(\cdot, \cdot))$ we have

$$|E_P(F(\tilde{\beta}, \tilde{\gamma})) - E_P(F(\beta_{*,n}, \gamma_{*,n}))| \rightarrow 0,$$

while from equation (A.21) we conclude that

$$|E_{P_n}(F(\beta_{*,n}, \gamma_{*,n})) - E_P(F(\beta_{*,n}, \gamma_{*,n}))| \rightarrow 0.$$

Thus $E_P(F(\tilde{\beta}, \tilde{\gamma})) = 0$ and similarly we can prove $E_P(G(\tilde{\beta}, \tilde{\gamma})) = 0$. Since $(\beta^\dagger, \gamma^\dagger)$ is the unique solution of the set of equations $E_P(F(\beta, \gamma)) = 0$ and $E_P(G(\beta, \gamma)) = 0$, we obtain $(\tilde{\beta}, \tilde{\gamma}) = (\beta^\dagger, \gamma^\dagger)$. Noting that the optimizers $(\beta_{*,n}, \gamma_{*,n})$ and $(\beta^\dagger, \gamma^\dagger)$ lie in the compact set \mathcal{C} with probability at least $1 - n^{-c}$ this proves that for some ϵ_n , we have $\|(\beta_{*,n}, \gamma_{*,n}) - (\beta^\dagger, \gamma^\dagger)\| = O(\epsilon_n)$ with probability at least $1 - n^{-c}$.

Thus, to prove the lemma it is sufficient to prove (A.21) and (A.22). Since $p/n \rightarrow \kappa$ it is enough to prove

$$\lim_{n \rightarrow \infty} \sup_{(\gamma, \beta) \in \mathcal{C}} |E_{P_n} F(\beta, \gamma) - E_P F(\beta, \gamma)| = 0.$$

By Arzela-Ascoli theorem the proof follows if we can show the following statements:

- 1) Pointwise convergence: For a given (γ, β) in \mathcal{C} , $\lim_{n \rightarrow \infty} |E_{P_n} F(\beta, \gamma) - E_P F(\beta, \gamma)| = 0$.
- 2) Equicontinuity: $\forall \epsilon, \exists \delta$ such that if $\|(\beta, \gamma) - (\beta', \gamma')\| < \delta$,

$$\sup_n |E_{P_n} F(\beta, \gamma) - E_{P_n} F(\beta', \gamma')| < \epsilon.$$

Writing $F(\beta, \gamma)$ as $f(M, \Gamma, \Omega)$, a function of M, Γ, Ω , one observes that to show (1), the following result can be used: $P_n \xrightarrow{W_2} P$ implies $|E_{P_n} F(\beta, \gamma) - E_P F(\beta, \gamma)| = 0$, if one can show $\sup_{(M, \Gamma, \Omega)} \frac{f(M, \Gamma, \Omega)}{\|M, \Gamma, \Omega\|^2} < \infty$. Noting the form of $F(\beta, \gamma)$ it is enough to show that,

$$\begin{aligned} & \sup_{(M, \Gamma, \Omega)} \frac{\left[\eta_2 \left(M + \gamma \Gamma; \frac{\gamma \lambda \Omega^2}{\beta} \right) - M \right]^2}{\|M, \Gamma, \Omega\|^2} \\ & \leq \sup_{(M, \Gamma, \Omega)} \frac{\left[(-\lambda \Omega^2 M + \beta \Gamma) / (\beta \gamma^{-1} + \lambda \Omega^2) \right]^2}{\|M, \Gamma\|^2} < \infty. \end{aligned}$$

Since Ω is bounded, the numerator is bounded by a quadratic function of M and G and hence the claim is proved.

For the function G , we similarly have to show

$$\begin{aligned} & \sup_{(M, \Gamma, \Omega)} \frac{\Gamma \left(\eta_2 \left(M + \gamma \Gamma; \frac{\gamma \lambda \Omega^2}{\beta} \right) - M \right)}{\|M, \Gamma, \Omega\|^2} \\ & \leq \sup_{(M, \Gamma, \Omega)} \frac{\Gamma(-\lambda \Omega^2 M + \beta \Gamma) / (\beta \gamma^{-1} + \lambda \Omega^2)}{\|M, \Gamma\|^2} < \infty. \end{aligned}$$

Again, the numerator is upper bounded by a quadratic function of Γ and M is bounded; hence, the claim is proved.

To show (2) we observe that,

$$\begin{aligned} & F(\beta, \gamma) - F(\beta', \gamma') \\ & = \left[(-\lambda \Omega^2 M + \beta \Gamma) / (\beta \gamma^{-1} + \lambda \Omega^2) \right]^2 \\ & \quad - \left[(-\lambda \Omega^2 M + \beta' \Gamma) / (\beta' \gamma'^{-1} + \lambda \Omega^2) \right]^2 + \gamma'^2 - \gamma^2 \\ & = (\beta \gamma^{-1} + \lambda \Omega^2)^{-2} (\beta' \gamma'^{-1} + \lambda \Omega^2)^{-2} \left[(\beta' \gamma'^{-1} \right. \\ & \quad \left. + \lambda \Omega^2)^2 (-\lambda \Omega^2 M + \beta \Gamma)^2 - (\beta \gamma^{-1} \right. \\ & \quad \left. + \lambda \Omega^2)^2 (-\lambda \Omega^2 M + \beta' \Gamma)^2 \right] \\ & \quad + \gamma'^2 - \gamma^2. \end{aligned}$$

Noting that $\beta, \beta', \gamma, \gamma'$ lies in compact sets bounded away from 0, we need to show that $\|(\beta, \gamma) - (\beta', \gamma')\| < \delta$ implies

$$|E_{P_n}[(\beta'\gamma'^{-1} + \lambda\Omega^2)^2(-\lambda\Omega^2M + \beta\Gamma)^2 - (\beta\gamma^{-1} + \lambda\Omega^2)^2(-\lambda\Omega^2M + \beta'\Gamma)^2]| < \epsilon. \quad (\text{A.23})$$

First, we show that

$$|E_{P_n}[(\beta'\gamma'^{-1} + \lambda\Omega^2)^2((-\lambda\Omega^2M + \beta\Gamma)^2 - (-\lambda\Omega^2M + \beta'\Gamma)^2)]| < \epsilon/2. \quad (\text{A.24})$$

Indeed since $(\beta'\gamma'^{-1} + \lambda\Omega^2)^2$ is bounded by some constant C , we consider the expression the following set of inequalities:

$$\begin{aligned} & |E_{P_n}[((-\lambda\Omega^2M + \beta\Gamma)^2 - (-\lambda\Omega^2M + \beta'\Gamma)^2)]| \\ & \leq E_{P_n}[|\beta'^2 - \beta^2|\Gamma^2] + E_{P_n}[2|\beta - \beta'|\lambda\Omega^2M\Gamma] \\ & \leq E_{P_n}[|\beta'^2 - \beta^2|\Gamma^2] + CE_{P_n}[|\beta - \beta'|\lambda(M^2 + \Gamma^2)]. \end{aligned}$$

Since both x^2 and x are uniformly continuous functions in the compact domain and $P_n \xrightarrow{W_2} P$ implies that $E_{P_n}\Gamma^2$ and $E_{P_n}(M^2 + \Gamma^2)$ is uniformly bounded we can choose δ small enough so that equation (A.24) holds.

Next we show that

$$|E_{P_n}[((\beta'\gamma'^{-1} + \lambda\Omega^2)^2 - (\beta\gamma^{-1} + \lambda\Omega^2)^2)(-\lambda\Omega^2M + \beta'\Gamma)^2]| < \epsilon/2. \quad (\text{A.25})$$

Using the facts that Ω is bounded, $\beta\gamma^{-1}$ and $\beta^2\gamma^{-2}$ are uniformly continuous functions in the compact domain (the domain also being bounded away from 0), we have $((\beta'\gamma'^{-1} + \lambda\Omega^2)^2 - (\beta\gamma^{-1} + \lambda\Omega^2)^2) < \epsilon$ for sufficiently small δ . Also, $P_n \xrightarrow{W_2} P$ implies that $E_{P_n}(-\lambda\Omega^2M + \beta'\Gamma)^2$ is uniformly bounded in n , so that choosing ϵ small enough, equation (A.25) holds. Finally we use triangle inequality to show (A.23) from (A.24) and (A.25). The equicontinuity of G can be shown along similar lines; hence we omit the proof.

Step 4: A calculation analogous to the one shown above proves that the following statements hold:

$$\lim_n \sup_{(\gamma, \beta) \in \mathcal{C}} |E_{Q_n}F_n(\beta, \gamma) - E_P F(\beta, \gamma)| = 0, \quad (\text{A.26})$$

$$\lim_n \sup_{(\gamma, \beta) \in \mathcal{C}} |E_{Q_n}G_n(\beta, \gamma) - E_P G(\beta, \gamma)| = 0. \quad (\text{A.27})$$

An argument similar to the one used in Step 3 can be used to conclude that with probability at least $1 - n^{-c}$ the following holds

$$\|(\beta_*, \gamma_*) - (\beta^\dagger, \gamma^\dagger)\| = O(\epsilon_n) \quad (\text{A.28})$$

Using equations (A.16) and (A.28) we conclude that the desired result holds. \square

ACKNOWLEDGMENT

The authors would like to thank Nabarun Deb, Rounak Dey, Rajarshi Mukherjee, and Subhabrata Sen for helpful discussions at various stages of the manuscript.

REFERENCES

- [1] V. A. Marchenko and L. A. Pastur, "Distribution of eigenvalues for some sets of random matrices," *Matematicheskii Sbornik*, vol. 114, no. 4, pp. 507–536, 1967.
- [2] D. L. Donoho, A. Maleki, and A. Montanari, "Message-passing algorithms for compressed sensing," *Proc. Nat. Acad. Sci. USA*, vol. 106, no. 45, pp. 18914–18919, Jul. 2009. [Online]. Available: <https://www.pnas.org/doi/abs/10.1073/pnas.0909892106>
- [3] M. Bayati and A. Montanari, "The dynamics of message passing on dense graphs, with applications to compressed sensing," *IEEE Trans. Inf. Theory*, vol. 57, no. 2, pp. 764–785, Feb. 2011.
- [4] M. Bayati and A. Montanari, "The LASSO risk for Gaussian matrices," *IEEE Trans. Inf. Theory*, vol. 58, no. 4, pp. 1997–2017, Apr. 2012.
- [5] S. Rangan, "Generalized approximate message passing for estimation with random linear mixing," in *Proc. IEEE Int. Symp. Inf. Theory*, Jul. 2011, pp. 2168–2172.
- [6] A. Javanmard and A. Montanari, "State evolution for general approximate message passing algorithms, with applications to spatial coupling," *Inf. Inference J. IMA*, vol. 2, no. 2, pp. 115–144, 2013.
- [7] L. Zdeborová and F. Krzakala, "Statistical physics of inference: Thresholds and algorithms," *Adv. Phys.*, vol. 65, no. 5, pp. 453–552, 2016.
- [8] D. Donoho and A. Montanari, "High dimensional robust M-estimation: Asymptotic variance via approximate message passing," *Probab. Theory Rel. Fields*, vol. 166, nos. 3–4, pp. 935–969, Dec. 2016.
- [9] P. Sur and E. J. Candès, "A modern maximum-likelihood theory for high-dimensional logistic regression," *Proc. Nat. Acad. Sci. USA*, vol. 116, no. 29, pp. 14516–14525, Jul. 2019. [Online]. Available: <https://www.pnas.org/doi/abs/10.1073/pnas.1810420116>
- [10] J. Barbier, F. Krzakala, N. Macris, L. Miolane, and L. Zdeborová, "Optimal errors and phase transitions in high-dimensional generalized linear models," *Proc. Nat. Acad. Sci. USA*, vol. 116, no. 12, pp. 5451–5460, Mar. 2019.
- [11] S. Wang, H. Weng, and A. Maleki, "Which bridge estimator is the best for variable selection?" *Ann. Statist.*, vol. 48, no. 5, pp. 2791–2823, Oct. 2020.
- [12] M. Mondelli and R. Venkataraman, "Approximate message passing with spectral initialization for generalized linear models," in *Proc. Int. Conf. Artif. Intell. Statist.*, 2021, pp. 397–405.
- [13] O. Y. Feng, R. Venkataraman, C. Rush, and R. J. Samworth, "A unifying tutorial on approximate message passing," *Found. Trends Mach. Learn.*, vol. 15, no. 4, pp. 335–536, 2022. [Online]. Available: <http://ieeexplore.ieee.org/document/9785928>
- [14] G. Li, W. Fan, and Y. Wei, "Approximate message passing from random initialization with applications to \mathbb{Z}_2 synchronization," *Proc. Nat. Acad. Sci. USA*, vol. 120, no. 31, Aug. 2023, Art. no. e2302930120. [Online]. Available: <https://www.pnas.org/doi/abs/10.1073/pnas.2302930120>
- [15] M. Mézard, G. Parisi, and M. A. Virasoro, *Spin Glass Theory and Beyond: An Introduction To the Replica Method and Its Applications*, vol. 9. Singapore: World Scientific, 1987.
- [16] M. Talagrand, *Spin Glasses: A Challenge for Mathematicians: Cavity and Mean Field Models*, vol. 46. Springer, 2003.
- [17] M. Mezard and A. Montanari, *Information, Physics, and Computation*. London, U.K.: Oxford Univ. Press, 2009.
- [18] N. El Karoui, D. Bean, P. J. Bickel, C. Lim, and B. Yu, "On robust regression with high-dimensional predictors," *Proc. Nat. Acad. Sci. USA*, vol. 110, no. 36, pp. 14557–14562, Sep. 2013. [Online]. Available: <https://www.pnas.org/doi/abs/10.1073/pnas.1307842110>
- [19] D. Bean, P. J. Bickel, N. El Karoui, and B. Yu, "Optimal M-estimation in high-dimensional regression," *Proc. Nat. Acad. Sci. USA*, vol. 110, no. 36, pp. 14563–14568, Sep. 2013. [Online]. Available: <https://www.pnas.org/doi/abs/10.1073/pnas.1307845110>
- [20] N. El Karoui, "On the impact of predictor geometry on the performance on high-dimensional ridge-regularized generalized robust regression estimators," *Probab. Theory Rel. Fields*, vol. 170, pp. 95–175, Feb. 2018.
- [21] P. Sur, Y. Chen, and E. J. Candès, "The likelihood ratio test in high-dimensional logistic regression is asymptotically a rescaled chi-square," *Probab. Theory Rel. Fields*, vol. 175, nos. 1–2, pp. 487–558, Oct. 2019.
- [22] Y. Chen, Y. Chi, J. Fan, and C. Ma, "Spectral methods for data science: A statistical perspective," *Found. Trends Mach. Learn.*, vol. 14, no. 5, pp. 566–806, 2021.
- [23] M. Stojnic, "A framework to characterize performance of LASSO algorithms," 2013, *arXiv:1303.7291*.
- [24] M. Stojnic, "Meshes that trap random subspaces," 2013, *arXiv:1304.0003*.

[25] M. Stojnic, "Upper-bounding ℓ_1 -optimization weak thresholds," 2013, *arXiv:1303.7289*.

[26] S. Oymak, C. Thrampoulidis, and B. Hassibi, "The squared-error of generalized LASSO: A precise analysis," in *Proc. 51st Annu. Allerton Conf. Commun., Control, Comput. (Allerton)*, Oct. 2013, pp. 1002–1009.

[27] C. Thrampoulidis, S. Oymak, and B. Hassibi, "Regularized linear regression: A precise analysis of the estimation error," in *Proc. 28th Conf. Learn. Theory*, vol. 40, Paris, France, P. Grünwald, E. Hazan, and S. Kale, Eds., Jul. 2015, pp. 1683–1709. [Online]. Available: <https://proceedings.mlr.press/v40/Thrampoulidis15.html>

[28] C. Thrampoulidis, E. Abbasi, and B. Hassibi, "Precise error analysis of regularized M -estimators in high dimensions," *IEEE Trans. Inf. Theory*, vol. 64, no. 8, pp. 5592–5628, Aug. 2018.

[29] E. J. Candès and P. Sur, "The phase transition for the existence of the maximum likelihood estimate in high-dimensional logistic regression," *Ann. Statist.*, vol. 48, no. 1, pp. 27–42, Feb. 2020.

[30] Q. Zhao, P. Sur, and E. J. Candès, "The asymptotic distribution of the MLE in high-dimensional logistic models: Arbitrary covariance," *Bernoulli*, vol. 28, no. 3, pp. 1835–1861, Aug. 2022.

[31] P. Sur and E. J. Candès, "A modern maximum-likelihood theory for high-dimensional logistic regression," *Proc. Nat. Acad. Sci. USA*, vol. 116, no. 29, pp. 14516–14525, Jul. 2019. [Online]. Available: <https://www.pnas.org/content/pnas/suppl/2019/06/29/1810420116.DCSupplemental/pnas.1810420116.sapp.pdf>

[32] A. Montanari and S. Sen, "A friendly tutorial on mean-field spin glass techniques for non-physicists," 2022, *arXiv:2204.02909*.

[33] T. Liang and P. Sur, "A precise high-dimensional asymptotic theory for boosting and minimum- ℓ_1 -norm interpolated classifiers," *Ann. Statist.*, vol. 50, no. 3, pp. 1669–1695, 2022.

[34] K. Jiang, R. Mukherjee, S. Sen, and P. Sur, "A new central limit theorem for the augmented IPW estimator: Variance inflation, cross-fit covariance and beyond," 2022, *arXiv:2205.10198*.

[35] S. Chatterjee, "A generalization of the Lindeberg principle," *Ann. Probab.*, vol. 34, no. 6, pp. 2061–2076, Nov. 2006.

[36] L. Erdős, B. Schlein, and H.-T. Yau, "Local semicircle law and complete delocalization for Wigner random matrices," *Commun. Math. Phys.*, vol. 287, no. 2, pp. 641–655, Apr. 2009.

[37] T. Tao and V. Vu, "Random matrices: The universality phenomenon for Wigner ensembles," in *Modern Aspects of Random Matrix Theory*, vol. 72, 2014, pp. 121–172.

[38] L. Erdős, "Universality of Wigner random matrices: A survey of recent results," *Russian Math. Surv.*, vol. 66, no. 3, pp. 507–626, Jun. 2011.

[39] M. Bayati, M. Lelarge, and A. Montanari, "Universality in polytope phase transitions and message passing algorithms," *Ann. Appl. Probab.*, vol. 25, no. 2, pp. 753–822, Apr. 2015.

[40] A. Montanari and P.-M. Nguyen, "Universality of the elastic net error," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Jun. 2017, pp. 2338–2342.

[41] W.-K. Chen and W.-K. Lam, "Universality of approximate message passing algorithms," *Electron. J. Probab.*, vol. 26, pp. 1–44, Jan. 2021.

[42] H. Hu and Y. M. Lu, "Universality laws for high-dimensional learning with random features," *IEEE Trans. Inf. Theory*, vol. 69, no. 3, pp. 1932–1964, Mar. 2023.

[43] A. Panahi and B. Hassibi, "A universal analysis of large-scale regularized least squares solutions," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 30, 2017.

[44] F. Gerace, F. Krzakala, B. Loureiro, L. Stephan, and L. Zdeborová, "Gaussian universality of perceptrons with random labels," 2022, *arXiv:2205.13303*.

[45] R. Dudeja, S. Sen, and Y. M. Lu, "Spectral universality in regularized linear regression with nearly deterministic sensing matrices," *IEEE Trans. Inf. Theory*, early access, Sep. 12, 2024, doi: [10.1109/TIT.2024.3458953](https://doi.org/10.1109/TIT.2024.3458953).

[46] T. Wang, X. Zhong, and Z. Fan, "Universality of approximate message passing algorithms and tensor networks," 2022, *arXiv:2206.13037*.

[47] Q. Han and Y. Shen, "Universality of regularized regression estimators in high dimensions," *Ann. Statist.*, vol. 51, no. 4, pp. 1799–1823, Aug. 2023.

[48] A. Montanari and B. N. Saeed, "Universality of empirical risk minimization," in *Proc. 355th Conf. Learn. Theory*, vol. 178, Jul. 2022, pp. 4310–4312.

[49] A. Montanari, F. Ruan, B. Saeed, and Y. Sohn, "Universality of max-margin classifiers," 2023, *arXiv:2310.00176*.

[50] Q. Han and X. Xu, "The distribution of ridgeless least squares interpolators," 2023, *arXiv:2307.02044*.

[51] A. I. Mundo, J. R. Tipton, and T. J. Muldoon, "Generalized additive models to analyze nonlinear trends in biomedical longitudinal data using R: Beyond repeated measures ANOVA and linear mixed models," *Statist. Med.*, vol. 41, no. 21, pp. 4266–4283, Sep. 2022. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/sim.9505>

[52] G. Z. Heller, K. P. Robledo, and I. C. Marschner, "Distributional regression in clinical trials: Treatment effects on parameters other than the mean," *BMC Med. Res. Methodol.*, vol. 22, no. 1, Dec. 2022, Art. no. 56.

[53] G. S. Marcillo et al., "Implementation of a generalized additive model (GAM) for soybean maturity prediction in African environments," *Agronomy*, vol. 11, no. 6, p. 1043, May 2021. [Online]. Available: <https://www.mdpi.com/2073-4395/11/6/1043>

[54] K. Ravindra, P. Rattan, S. Mor, and A. N. Aggarwal, "Generalized additive models: Building evidence of air pollution, climate change and human health," *Environ. Int.*, vol. 132, Nov. 2019, Art. no. 104987. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0160412019309341>

[55] J. S. Morris, "Functional regression," *Annu. Rev. Statist. Appl.*, vol. 2, pp. 321–359, Apr. 2015.

[56] R. C. Lewontin, "On measures of gametic disequilibrium," *Genetics*, vol. 120, no. 3, pp. 849–852, Nov. 1988.

[57] S. B. Korada and A. Montanari, "Applications of the Lindeberg principle in communications and statistical learning," *IEEE Trans. Inf. Theory*, vol. 57, no. 4, pp. 2440–2450, Apr. 2011.

[58] A. Montanari, F. Ruan, Y. Sohn, and J. Yan, "The generalization error of max-margin linear classifiers: Benign overfitting and high dimensional asymptotics in the overparametrized regime," 2019, *arXiv:1911.01544*.

[59] Z. Deng, A. Kammoun, and C. Thrampoulidis, "A model of double descent for high-dimensional binary linear classification," *Inf. Inference, J. IMA*, vol. 11, no. 2, pp. 435–495, Jun. 2022.

[60] T. Liang, S. Sen, and P. Sur, "High-dimensional asymptotics of Langevin dynamics in spiked matrix models," *Inf. Inference, J. IMA*, vol. 12, no. 4, pp. 2720–2752, Sep. 2023.

[61] R. Dudeja, Y. M. Lu, and S. Sen, "Universality of approximate message passing with semirandom matrices," *Ann. Probab.*, vol. 51, no. 5, pp. 1616–1683, Sep. 2023.

[62] D. Donoho and J. Tanner, "Observed universality of phase transitions in high-dimensional geometry, with implications for modern data analysis and signal processing," *Phil. Trans. Roy. Soc. A, Math., Phys. Eng. Sci.*, vol. 367, no. 1906, pp. 4273–4293, Nov. 2009.

[63] H. Monajemi et al., "Deterministic matrices matching the compressed sensing phase transitions of Gaussian random matrices," *Proc. Nat. Acad. Sci. USA*, vol. 110, no. 4, pp. 1181–1186, 2013.

[64] A. Abbara, A. Baker, F. Krzakala, and L. Zdeborová, "On the universality of noiseless linear estimation with respect to the measurement matrix," *J. Phys. A, Math. Theor.*, vol. 53, no. 16, Apr. 2020, Art. no. 164001.

[65] Y. Li and P. Sur, "Spectrum-aware debiasing: A modern inference framework with applications to principal components regression," 2023, *arXiv:2309.07810*.

[66] C. J. Stone, "Additive regression and other nonparametric models," *Ann. Statist.*, vol. 13, no. 2, pp. 689–705, Jun. 1985.

[67] Z. Tan and C.-H. Zhang, "Doubly penalized estimation in additive regression with high-dimensional data," *Ann. Statist.*, vol. 47, no. 5, pp. 2567–2600, Oct. 2019.

[68] A. B. Tsybakov, *Introduction to Nonparametric Estimation* (Springer Series in Statistics). New York, NY, USA: Springer, 2009.

[69] P. J. Bickel, Y. Ritov, and A. B. Tsybakov, "Simultaneous analysis of lasso and Dantzig selector," *Ann. Statist.*, vol. 37, no. 4, pp. 1705–1732, Aug. 2009.

[70] P. T. Reiss, J. Goldsmith, H. L. Shang, and R. T. Ogden, "Methods for scalar-on-function regression," *Int. Stat. Rev.*, vol. 85, no. 2, pp. 228–249, Aug. 2017.

[71] J. Goldsmith, J. Bobb, C. M. Crainiceanu, B. Caffo, and D. Reich, "Penalized functional regression," *J. Comput. Graph. Statist.*, vol. 20, no. 4, pp. 830–851, Dec. 2011.

[72] V. Koltchinskii and S. Minsker, " L_1 -penalization in functional linear regression with subgaussian design," *J. de l'École Polytechnique, Mathématiques*, vol. 1, pp. 269–330, Oct. 2014. [Online]. Available: <https://jep.centre-mersenne.org/articles/10.5802/jep.11/>

[73] M. Sesia, S. Bates, E. Candès, J. Marchini, and C. Sabatti, "False discovery rate control in genome-wide association studies with population structure," *Proc. Nat. Acad. Sci. USA*, vol. 118, no. 40, Oct. 2021, Art. no. e2105841118.

- [74] D. Gianola, E. Manfredi, and H. Simianer, "On measures of association among genetic variables," *Animal Genet.*, vol. 43, no. s1, pp. 19–35, Jul. 2012.
- [75] C. Pattaro, I. Ruczinski, D. M. Fallin, and G. Parmigiani, "Haplotype block partitioning as a tool for dimensionality reduction in SNP association studies," *BMC Genomics*, vol. 9, no. 1, pp. 1–15, Dec. 2008.
- [76] M. Sesia, C. Sabatti, and E. J. Candès, "Gene hunting with hidden Markov model knockoffs," *Biometrika*, vol. 106, no. 1, pp. 1–18, Mar. 2019.
- [77] P. Scheet and M. Stephens, "A fast and flexible statistical model for large-scale population genotype data: Applications to inferring missing genotypes and haplotypic phase," *Amer. J. Hum. Genet.*, vol. 78, no. 4, pp. 629–644, Apr. 2006.
- [78] E. Abbasi, F. Salehi, and B. Hassibi, "Universality in learning from linear measurements," in *Advances in Neural Information Processing Systems*, vol. 32, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, Eds., Red Hook, NY, USA: Curran Associates, 2019. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2019/file/dffbb6efd376d8dbb22cdf491e481edc-Paper.pdf
- [79] S. Oymak and J. A. Tropp, "Universality laws for randomized dimension reduction, with applications," *Inf. Inference, J. IMA*, vol. 7, no. 3, pp. 337–446, Sep. 2018.
- [80] L. Meier, S. van de Geer, and P. Bühlmann, "High-dimensional additive modeling," *Ann. Statist.*, vol. 37, no. 6B, pp. 3779–3821, Dec. 2009.
- [81] V. Koltchinskii and M. Yuan, "Sparsity in multiple kernel learning," *Ann. Statist.*, vol. 38, no. 6, pp. 3660–3695, Dec. 2010.
- [82] M. Yuan and D.-X. Zhou, "Minimax optimal rates of estimation in high dimensional additive models," *Ann. Statist.*, vol. 44, no. 6, pp. 2564–2593, Dec. 2016.
- [83] G. Raskutti, M. J. Wainwright, and B. Yu, "Minimax-optimal rates for sparse additive models over kernel classes via convex programming," *J. Mach. Learn. Res.*, vol. 13, no. 2, pp. 389–427, 2012.
- [84] L. Miolane and A. Montanari, "The distribution of the lasso: Uniform control over sparse balls and adaptive parameter tuning," *Ann. Statist.*, vol. 49, no. 4, pp. 2313–2335, Aug. 2021.
- [85] S. N. Negahban, P. Ravikumar, M. J. Wainwright, and B. Yu, "A unified framework for high-dimensional analysis of M -estimators with decomposable regularizers," *Stat. Sci.*, vol. 27, no. 4, pp. 538–557, Nov. 2012.
- [86] N. Ross, "Fundamentals of Stein's method," *Probab. Surv.*, vol. 8, pp. 210–293, Oct. 2011.
- [87] R. Vershynin, *Introduction to the Non-asymptotic Analysis of Random Matrices*. Cambridge, U.K.: Cambridge Univ. Press, 2012, pp. 210–268.

Samriddha Lahiry received the Bachelor of Statistics (B.Stat.) and Master of Statistics (M.Stat.) degrees from Indian Statistical Institute, Kolkata, India, in 2015 and 2017, respectively, and the Ph.D. degree in statistics from Cornell University in 2022. He is currently a Visiting Fellow with the Department of Statistics and Data Science, National University of Singapore, Singapore. His research interests include high-dimensional statistics, exact asymptotics, nonparametric estimation, and quantum statistical inference.

Pragya Sur received the Bachelor of Statistics (B.Stat.) and Master of Statistics (M.Stat.) from Indian Statistical Institute, Kolkata, India, in 2012 and 2014, respectively, and the Ph.D. degree in statistics from Stanford University in 2019. She is currently an Assistant Professor with the Department of Statistics, Harvard University, Cambridge, MA, USA. Her research interests include high-dimensional statistics, machine learning theory, learning under heterogeneity, and distribution shifts. She was a recipient of several awards, including NSF DMS 2113426, a William F. Milton Fund Award, a Dean's Competitive Fund for Promising Scholarship, and the Theodore W. Anderson Theory of Statistics Dissertation Award.