
Random Features Model with General Convex Regularization: A Fine Grained Analysis with Precise Asymptotic Learning Curves

David Bosch
Chalmers University
of Technology

Ashkan Panahi
Chalmers University
of Technology

Ayca Özcelikkale
Uppsala University

Devdatt Dubhashi
Chalmers University
of Technology

Abstract

We compute precise asymptotic expressions for the learning curves of least squares random feature (RF) models with either a separable strongly convex regularization or the ℓ_1 regularization. We propose a novel multi-level application of the convex Gaussian min max theorem (CGMT) to overcome the traditional difficulty of finding computable expressions for random features models with correlated data. Our result takes the form of a computable 4-dimensional scalar optimization. In contrast to previous results, our approach does not require solving an often intractable proximal operator, which scales with the number of model parameters. Furthermore, we extend the universality results for the training and generalization errors for RF models to ℓ_1 regularization. In particular, we demonstrate that under mild conditions, random feature models with elastic net or ℓ_1 regularization are asymptotically equivalent to a surrogate Gaussian model with the same first and second moments. We numerically demonstrate the predictive capacity of our results, and show experimentally that the predicted test error is accurate even in the non-asymptotic regime.

1 INTRODUCTION

It has been recently understood that classical statistical theory requires revisiting to describe the behavior of overparameterized models (Zhang et al., 2021; Belkin et al., 2019). Since then, studying the asymptotic regime of a machine learning (ML) model, in which the number of data points and model parameters grow infinite at a constant ra-

tio, has become a popular method of analysis (Belkin et al., 2020; Hastie et al., 2019; Bartlett et al., 2020a; Tsigler and Bartlett, 2020). The asymptotic analysis of regularized Random Feature (RF) models (Rahimi and Recht, 2007) has been of particular interest as they can capture a large range of other interesting models (Mei and Montanari, 2019; Goldt et al., 2020a; d’Ascoli et al., 2020; Dhifallah and Lu, 2020). Despite remarkable progress in the analysis of RF models, existing asymptotic results are not directly computable for the majority of regularization functions, and in this generic scenario, precise asymptotic learning curves are still lacking. In this paper, we address this limitation and provide a novel technique that provides computable, exact asymptotic learning curves under a large family of separable, strongly convex regularization, as well as the ℓ_1 regularization (also known as LASSO).

Similar to many recent papers, we make use of the convex Gaussian Min Max theorem (CGMT) (Thrampoulidis et al., 2014; Hastie et al., 2019; Montanari et al., 2019; Dhifallah and Lu, 2020; Goldt et al., 2020b), where there are generally two steps. The RFs are non-Gaussian due to non-linear activation functions, but it is shown that they can be equivalently replaced by a surrogate Gaussian model with matching first two statistical moments (Panahi and Hassibi, 2017; Oymak and Tropp, 2018; Hu and Lu, 2020). Establishing this equivalence between the RF model and surrogate Gaussian model is generally referred to as universality (Panahi and Hassibi, 2017; Oymak and Tropp, 2018; Hu and Lu, 2020). Next, the CGMT is applied, which provides an alternative optimization problem whose analysis is provably tied to the original problem. This alternative optimization formulation has been a great tool for computing precise asymptotic learning curves in the case of uncorrelated features. However, for the general RF formulation, the surrogate features are inevitably heavily correlated. As a result, the alternative optimization has been generally as difficult to analyze as the original RF model. More precisely, solution of this alternative optimization typically involves solving a proximal operator of a non separable m -dimensional vector that scales with the number of model parameters (Loureiro et al., 2021), even if the regulariza-

tion function is separable. Only in the case of ridge (ℓ_2^2) regularization, where a specific rotational symmetry holds true, can this difficulty be overcome (Chang et al., 2020; Montanari et al., 2019; d’Ascoli et al., 2020).

Contributions: The first main contribution of this paper is a novel multilevel application of the CGMT to the correlated surrogate model that overcomes the difficulties with the analysis of the alternative optimization and substantially simplifies the final results. With this method, we provide a computable technique for obtaining learning curves of surrogate Gaussian model with arbitrary separable, strongly convex; or ℓ_1 regularization. Our next contribution is to establish universality, i.e. to show that our analysis also applies to the original, non-Gaussian random features. This result has been previously established for regularization functions that are thrice differentiable and strongly convex (Hu and Lu, 2020). We extend this result in two steps. First we show that a wider variety of potentially nondifferentiable, strongly convex functions satisfy universality. In particular, we show that a combination of ℓ_1 and ℓ_2^2 , known as elastic net (Zou and Hastie, 2005), is universal. Furthermore, under the assumptions that the activation function is continuous and Lipschitz, and solution vector that is sufficiently sparse, we show that the ℓ_2^2 part of the elastic net regularization can be removed and the universality of pure ℓ_1 (which is not strongly convex) is established.

2 RELATED WORKS

The asymptotic analysis of RF models is recently culminated in the study of the so-called double descent phenomenon, where increasing the model size beyond the interpolation threshold, surprisingly improves the learning performance, leading to a learning curve with two descent regions. The double descent phenomenon has a long history (Loog et al., 2020), but was first discussed in its modern form by (Belkin et al., 2019) (see also (Geiger et al., 2020)). Overparameterized systems have since been studied extensively, for an incomplete list see (Tsigler and Bartlett, 2020; Hastie et al., 2019; Mei and Montanari, 2019; Bartlett et al., 2020b; Belkin et al., 2020; Muthukumar et al., 2019; Kobak et al., 2020; Deng et al., 2019; Taheri et al., 2021; Lolas, 2020; Mignacco et al., 2020; Kini and Thrampoulidis, 2020; Liang and Sur, 2020; Montanari et al., 2019; Taheri et al., 2020; Salehi et al., 2019)

Gaussian comparison theorems have played a central role in obtaining exact learning curves, which go back to (Gordon, 1985, 1988). They show an asymptotic equivalence between certain optimization problems over Gaussian random variables. (Thrampoulidis et al., 2015; Thrampoulidis et al., 2014) showed that in the presence of convexity, the bounds provided by Gordon could be refined. The applications of comparison theorems to the study of the asymptotic

regime are numerous (Bosch et al., 2021; Loureiro et al., 2021; Dhifallah and Lu, 2020; Thrampoulidis et al., 2016; Chang et al., 2020). A principal difficulty with the CGMT is in the case of correlated covariates, as in the RF model. This results in the alternative optimization problem of the CGMT to be no more tractable than the original problem. In the case of ℓ_2^2 regularization, rotational symmetry may be applied to study correlated models. In the papers such as (Chang et al., 2020; Mei and Montanari, 2019; Dhifallah and Lu, 2020) this symmetry is exploited to derive analytic expressions. We are not aware of any analytic expressions derived by means of the CGMT considering RFs with more generic regularization. As a contribution of this paper, we resolved the issue of correlated covariates with a novel approach involving multiple applications of CGMT and extend the analysis of regularized least squares into RF features with a larger set of regularization functions.

The Gaussian Equivalence Principle (GEP) expresses that there exists an asymptotic equivalence between RF models and Gaussian models with identical first and second moments. This universality was shown for (regularized) least squares by Panahi and Hassibi (2017), extended to generic convex regularization by Hu and Lu (2020) and for generative models by Goldt et al. (2022). More recent results by Montanari and Saeed (2022) extends universality to empirical risk minimization with regularization. Ba et al. (2022) has also extended universality results to RF models after a single step of gradient descent with small step sizes. The results of Hu and Lu (2020) and Montanari and Saeed (2022) however do not hold in the case of ℓ_1 regularization, while those of Panahi and Hassibi (2017) do not apply to the random features case. We extend the universality results of Hu and Lu (2020) to the case of ℓ_1 and elastic net regularization. Liang and Sur (2020) also demonstrate the universality of ℓ_1 regularization but for a different setup of max-margin classifiers. Their results cannot simply be translated to that of ours. Firstly, they only consider universality of the objective value, while we additionally demonstrate universality for strongly convex functions of the solution vector. Secondly they require that the activation function is restricted to a compact set, which we do not require here.

3 RANDOM FEATURES MODEL

We consider a dataset $\{(\mathbf{z}_i, y_i) \in \mathbb{R}^d \times \mathbb{R}\}_{i=1}^n$ and wish to determine the relationship between the data vector \mathbf{z}_i and the labels y_i by means of a function of the following form:

$$f(\mathbf{z}_i; \boldsymbol{\theta}, \boldsymbol{\varphi}) = \frac{1}{\sqrt{m}} \boldsymbol{\theta}^T \boldsymbol{\varphi}(\mathbf{z}_i) \quad \boldsymbol{\theta} \in \mathbb{R}^m. \quad (1)$$

Here $\boldsymbol{\varphi} : \mathbb{R}^d \rightarrow \mathbb{R}^m$ is a fixed nonlinear feature map, whose relation to the labels y_i is characterized by a variable weight vector $\boldsymbol{\theta}$. We determine $\boldsymbol{\theta}$ by the following optimization

problem:

$$\hat{\boldsymbol{\theta}} = \arg \min_{\boldsymbol{\theta}} \sum_{i=1}^n l(f(\mathbf{z}_i; \boldsymbol{\theta}, \boldsymbol{\varphi}), y_i) + r(\boldsymbol{\theta}), \quad (2)$$

where $l(x, y) = \frac{1}{2}(x - y)^2$ is the square-loss function and $r(x)$ is a regularization function. We consider a wide-range of regularization functions which are explained in Section 4. We restrict ourselves to the feature map

$$\boldsymbol{\varphi}(\mathbf{z}_i) = \sigma \left(\frac{1}{\sqrt{d}} \mathbf{W} \mathbf{z}_i \right), \quad (3)$$

where $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ is a non linear, odd activation function applied element wise (eg. $\tanh(x)$), and $\mathbf{W} \in \mathbb{R}^{m \times d}$ is a random weight matrix whose elements are i.i.d standard Gaussians, independent of \mathbf{z}_i . We note that this choice of the random feature map can be interpreted as a Neural Network (NN) with one hidden layer. We let the matrix \mathbf{X} be given such that $\mathbf{X}_{ij} = \varphi_j(\mathbf{z}_i) = \sigma \left(\frac{1}{\sqrt{d}} \mathbf{w}_j^T \mathbf{z}_i \right)$, where \mathbf{w}_j^T is the j^{th} row of \mathbf{W} . We consider two metrics of the performance of the solution $\hat{\boldsymbol{\theta}}$ of (2), the training error, expressed in matrix notation as

$$\mathcal{E}_{train}(\boldsymbol{\theta}) = \frac{1}{2n} \|\mathbf{y} - \frac{1}{\sqrt{m}} \mathbf{X} \boldsymbol{\theta}\|_2^2 + \frac{1}{m} r(\boldsymbol{\theta}) \quad (4)$$

and the generalization error

$$\mathcal{E}_{gen}(\boldsymbol{\theta}) = \mathbb{E} \left[\frac{1}{2} (y_{new} - f(\mathbf{z}_{new}; \boldsymbol{\theta}, \boldsymbol{\varphi}))^2 \right], \quad (5)$$

where $(\mathbf{z}_{new}, y_{new})$ is a new sample pair independent of, but identically distributed to the training data.

Analysis of this problem requires making assumptions on the distribution of the the dataset. We assume that $\mathbf{z}_i \stackrel{i.i.d.}{\sim} \mathcal{N}(0, \mathbf{I}_d)$ and that the labels y_i are generated according to

$$y_i = \frac{1}{\sqrt{m}} \boldsymbol{\theta}^{*T} \boldsymbol{\varphi}(\mathbf{z}_i) + \epsilon_i, \quad (6)$$

where $\boldsymbol{\theta}^*$ is a fixed weight vector that may be deterministic or random and ϵ_i is i.i.d. noise with $\mathbb{E}[\epsilon_i] = 0$, $\mathbb{E}[\epsilon_i^2] = \sigma_\epsilon^2$ and $\mathbb{E}[\epsilon_i^4] < \infty$, and $\boldsymbol{\varphi}$ is given in (3). We note that this method of label generation is different that that of Hu and Lu (2020), we note that their results still apply in this context. For a discussion of this fact see remark 2 in the appendix.

Under these assumptions, the main goal of this paper is to predict the values of $\mathcal{E}_{gen}(\hat{\boldsymbol{\theta}})$, $\mathcal{E}_{train}(\hat{\boldsymbol{\theta}})$, where $\hat{\boldsymbol{\theta}}$ is given by (2). Further, we provide the asymptotic value of $h(\hat{\boldsymbol{\theta}})$ where h is an arbitrary test function from a wide range of choices, as we elaborate.

4 MAIN RESULTS

4.1 Overview of Main Results

Before delving into details, we provide an overview of our main results. A more detailed and rigorous treatment is provided in the subsequent sections.

The key optimization problem in (2) can be written as

$$P_1 = \min_{\boldsymbol{\theta}} \frac{1}{2n} \|\mathbf{y} - \frac{1}{\sqrt{m}} \mathbf{X} \boldsymbol{\theta}\|_2^2 + \frac{1}{m} r(\boldsymbol{\theta}). \quad (7)$$

Hence, the optimal solution of P_1 is given by (2). However we consider a slightly more general problem of the following form:

$$\begin{aligned} \tilde{P}_1(\tau_1, \tau_2) = \min_{\boldsymbol{\theta}} & \frac{1}{2n} \|\mathbf{y} - \frac{1}{\sqrt{m}} \mathbf{X} \boldsymbol{\theta}\|_2^2 + \frac{1}{m} r(\boldsymbol{\theta}) \\ & + \frac{\tau_1}{m} (\boldsymbol{\theta} - \boldsymbol{\theta}^*)^T \mathbf{R} (\boldsymbol{\theta} - \boldsymbol{\theta}^*) + \frac{\tau_2}{m} h(\boldsymbol{\theta}), \end{aligned} \quad (8)$$

where τ_1, τ_2 are real numbers and $h(\mathbf{e})$ is a test function such that $r + \tau_2 h$ is convex. Moreover, \mathbf{R} is the feature covariance matrix $\mathbb{E}_{\mathbf{z}}[\boldsymbol{\varphi}(\mathbf{z}) \boldsymbol{\varphi}(\mathbf{z})^T]$. We refer to the solution of (8) as $\tilde{\boldsymbol{\theta}}_1(\tau_1, \tau_2)$.

We note that setting $\tau_1 = \tau_2 = 0$, we obtain the original problem (7), i.e. $P_1 = \tilde{P}_1(0, 0)$ and $\hat{\boldsymbol{\theta}}_1 = \tilde{\boldsymbol{\theta}}_1(0, 0)$. These additional “ τ ” are added to the problem definition to prove the universality of generalization error and of generic strongly convex functions. We note that the τ_1 term corresponds to a component of the generalization function and τ_2 is attached to the generic function $h(\boldsymbol{\theta})$. Taking the derivative with respect to τ_1 , or τ_2 allows these terms to be recovered, this property is made use of in the proof of the universality, see proof of theorem 4.

We analyze the problem in (8) by considering two alternative problem formulations, and demonstrating that they are asymptotically equivalent to one another.

Consider the linear feature map

$$\tilde{\boldsymbol{\varphi}}(\mathbf{z}) = \frac{\rho_1}{\sqrt{d}} \mathbf{W} \mathbf{z} + \rho_* \mathbf{g}, \quad (9)$$

where $\rho_1 = \mathbb{E}_a[a\sigma(a)]$ and $\rho_*^2 = \mathbb{E}_a[\sigma^2(a)] - \rho_1^2$, with $a \sim \mathcal{N}(0, 1)$, and $\mathbf{g} \sim \mathcal{N}(0, \mathbf{I}_m)$. This feature map is obtained by means of a truncated Hermite polynomial expansion of the original feature map (3), as discussed in Mei and Montanari (2019), and unlike the original feature maps in (3) these feature are Gaussian (for fixed weights \mathbf{W}). Let $(\tilde{\mathbf{X}})_{ij} = \tilde{\varphi}_j(\mathbf{z}_i) = \frac{\rho_1}{\sqrt{d}} \mathbf{w}_j^T \mathbf{z}_i + \rho_* g_{ij}$, where g_{ij} are i.i.d Gaussian, and consider the problem

$$\begin{aligned} \tilde{P}_2(\tau_1, \tau_2) = \min_{\boldsymbol{\theta}} & \frac{1}{2n} \|\mathbf{y} - \frac{1}{\sqrt{m}} \tilde{\mathbf{X}} \boldsymbol{\theta}\|_2^2 + \frac{1}{m} r(\boldsymbol{\theta}) \\ & + \frac{\tau_1}{m} (\boldsymbol{\theta} - \boldsymbol{\theta}^*)^T \tilde{\mathbf{R}} (\boldsymbol{\theta} - \boldsymbol{\theta}^*) + \frac{\tau_2}{m} h(\boldsymbol{\theta}), \end{aligned} \quad (10)$$

where $y_i = \boldsymbol{\theta}^{*T} \tilde{\boldsymbol{\varphi}}(\mathbf{z}_i) + \epsilon_i$ and $\tilde{\mathbf{R}} = \mathbb{E}_{\mathbf{z}}[\tilde{\boldsymbol{\varphi}}(\mathbf{z})\tilde{\boldsymbol{\varphi}}(\mathbf{z})^T]$. The optimal solution of $\tilde{P}_2(\tau_1, \tau_2)$ is referred to as $\hat{\boldsymbol{\theta}}_2(\tau_1, \tau_2)$. In particular, we denote $\hat{\boldsymbol{\theta}}_2 = \tilde{\boldsymbol{\theta}}_2(0, 0)$.

Now, we define $\psi(\beta, q, \xi, t, \tau_1, \tau_2)$ as follows

$$\begin{aligned} \psi(\beta, q, \xi, t, \tau_1, \tau_2) = & \frac{1}{m} \mathbb{E} \left[\mathcal{M}_{\frac{1}{2c_1}}(r + \tau_2 h) \left(\boldsymbol{\theta}^* - \frac{c_2 \sqrt{\gamma}}{2c_1} \boldsymbol{\phi} \right) \right] \\ & - \frac{c_2^2 \gamma}{4c_1} + \frac{\xi t}{2} + \frac{\beta q}{2} + \frac{\beta \sigma_\epsilon^2}{2q} + \frac{\xi \beta^2}{2t\eta} \\ & - \frac{(\beta + 2\tau_1 q) \xi^2}{2q} - \frac{q \beta^2}{2(\beta + 2q\tau_1)\eta} - \frac{\beta^2}{2}, \end{aligned} \quad (11)$$

where $\mathcal{M}_{\frac{1}{2c_1}}(r + \tau_2 h)$ is the Moreau envelope of $r + \tau_2 h$ with the step size $\frac{1}{2c_1}$ (see supplement definition 1), $\boldsymbol{\phi}$ is a standard Gaussian vector, c_1 and c_2 are functions of $\beta, q, \xi, t, \tau_1, \tau_2$ given by

$$c_1 = \frac{(\beta + 2\tau_1 q)^2 \rho_1^2 \xi}{2q^2 t} + \frac{(\beta + 2q\tau_1) \rho_*^2}{2q} \quad (12)$$

$$c_2 = \sqrt{\frac{(\beta + 2\tau_1 q)^2 \rho_1^2 \xi^2 \eta}{q^2} + \beta^2 \rho_*^2}. \quad (13)$$

The expectation is taken with respect to $\boldsymbol{\phi}$ and hence the function ψ is not random. Accordingly, we define the key alternative optimization problem, i.e. a four-dimensional scalar optimization problem, in our development:

$$\tilde{P}_3(\tau_1, \tau_2) = \max_{\beta > 0} \min_{q > 0} \max_{\xi > 0} \min_{t > 0} \psi(\beta, q, \xi, t, \tau_1, \tau_2). \quad (14)$$

Let $\tilde{\beta}, \tilde{q}, \tilde{\xi}, \tilde{t}$ be the optimal point of \tilde{P}_3 and let $\tilde{c}_1 = c_1(\tilde{\beta}, \tilde{q}, \tilde{\xi}, \tilde{t})$ and $\tilde{c}_2 = c_2(\tilde{\beta}, \tilde{q}, \tilde{\xi})$. Accordingly, we define $\tilde{\boldsymbol{\theta}}_3(\tau_1, \tau_2)$ as follows

$$\tilde{\boldsymbol{\theta}}_3(\tau_1, \tau_2) := \text{prox}_{\frac{1}{2\tilde{c}_1}}(r + \tau_2 h) \left(\boldsymbol{\theta}^* - \frac{\tilde{c}_2 \sqrt{\gamma}}{2\tilde{c}_1} \boldsymbol{\phi} \right), \quad (15)$$

where $\text{prox}_{\frac{1}{2\tilde{c}_1}}(r + \tau_2 h)$ denotes the proximal operator of $r + \tau_2 h$ with the step size $\frac{1}{2\tilde{c}_1}$. Similar to the two previous cases, we define $\hat{\boldsymbol{\theta}}_3 = \tilde{\boldsymbol{\theta}}_3(0, 0)$. The training and generalization error corresponding to problem \tilde{P}_3 are not given by (4) and (5), instead we have that

$$\tilde{\mathcal{E}}_{train} = \tilde{P}_3(0, 0) \quad \tilde{\mathcal{E}}_{gen} = \sigma_\epsilon^2 + \left. \frac{\partial \tilde{P}_3(\tau_1, 0)}{\partial \tau_1} \right|_{\tau_1=0}. \quad (16)$$

Now, we provide a summary of our main results:

Theorem 1. Informal statement of the main results

There exist symmetric intervals $\tau_1 \in [-\tau_1^*, \tau_1^*]$ and $\tau_2 \in [-\tau_2^*, \tau_2^*]$ with sufficiently small universal constants τ_1^*, τ_2^* , a wide family of strongly convex, separable functions r and potentially non-convex, separable test functions h , for which in the asymptotic limit,

$$\tilde{P}_1(\tau_1, \tau_2) \approx \tilde{P}_2(\tau_1, \tau_2), \quad \tilde{P}_2(\tau_1, \tau_2) \approx \tilde{P}_3(\tau_1, \tau_2) \quad (17)$$

and hence

$$\tilde{P}_1(\tau_1, \tau_2) \approx \tilde{P}_3(\tau_1, \tau_2). \quad (18)$$

By the above result, we may conclude for such scenarios that

$$\mathcal{E}_{train}(\hat{\boldsymbol{\theta}}_1) \approx \mathcal{E}_{train}(\hat{\boldsymbol{\theta}}_2) \approx \tilde{\mathcal{E}}_{train}, \quad (19)$$

$$\mathcal{E}_{gen}(\hat{\boldsymbol{\theta}}_1) \approx \mathcal{E}_{gen}(\hat{\boldsymbol{\theta}}_2) \approx \tilde{\mathcal{E}}_{gen}, \quad (20)$$

and

$$h(\hat{\boldsymbol{\theta}}_1) \approx h(\hat{\boldsymbol{\theta}}_2) \approx h(\hat{\boldsymbol{\theta}}_3). \quad (21)$$

The above result also holds for ℓ_1 regularization under some considerations about the true model $\boldsymbol{\theta}^*$ and the activation function.

Discussion of Main Result: By Theorem 1, the generalization/training error and other properties of the original problem P_1 , represented by a test function h , can be found using the solution of $P_3 = \tilde{P}_3(\tau_1 = 0, \tau_2 = 0)$. See Theorem 5 for a precise statement. Note that P_3 is scalar and since r is separable, calculating $\mathbb{E} \left[\mathcal{M}_{\frac{1}{2c_1}}(r) \left(\boldsymbol{\theta}^* - \frac{c_2 \sqrt{\gamma}}{2c_1} \boldsymbol{\phi} \right) \right]$ is straightforward (τ_1, τ_2 are set to zero). Hence, P_3 is simple to evaluate using standard computation techniques.

We note that Theorem 1 is, at first sight, similar to Theorem 1 in Loureiro et al. (2021), which is also based on the Moreau envelope and the proximal operator of the regularization function. However, we note that the argument of the Moreau envelope in their expression is more complex and cannot be generally evaluated even if r is separable. Hence, our result is novel and not the same as (Loureiro et al., 2021, Thm. 1) and allows significantly easier calculation of the generalization error compared to other existing methods in the literature for the correlated RF model.

Our proof has two building blocks: Using a novel multi-level application of CGMT, we show in Theorem 2, the convergence of \tilde{P}_2 to the scalar optimization problem \tilde{P}_3 in the left hand side of (17). The universality result, i.e. the asymptotic convergence between \tilde{P}_1 and \tilde{P}_2 in the right hand side of (17) is presented in Section 4.4. The other claims i.e (19),(20) and (21) are subsequently obtained by an individual argument.

For strongly convex and thrice differentiable regularization functions, the universality relation in the right hand side of (17) has already been demonstrated in Hu and Lu (2020). Here, we extend these results to the case of a sequence of strongly convex, thrice differentiable functions with bounded third derivatives that converge uniformly to the regularization function (Theorem 4). Such functions may not be even differentiable. Moreover, while Hu and Lu (2020) also shows the universality of the generalization/test errors, we extend this result and show that the entire discussion holds true for an arbitrary test function h obtained as the uniform limit of a sequence of thrice differentiable functions with bounded third derivatives. Exact

assumptions will be shortly presented. The above approach also allows us to extend the universality results to elastic net (Corollary 1) and ℓ_1 regularization (Theorem 6), which have not been provided in the literature before.

4.2 Assumptions

Below, we provide a list of all assumptions considered in our study. The specific assumptions that are used for each result is provided under the statement of the associated result.

A1 The regularization function satisfies one of the below:

- Case A: For positive constants $\mu, L > 0$, there exists a sequence of functions $r^{(k)}$ that are separable, μ -strongly convex and thrice differentiable with L -uniformly bounded third derivatives¹. The sequence $r^{(k)}$ converges uniformly in the limit of $k \rightarrow \infty$ to the regularization function r .
- Case B: The regularization function is $r(\boldsymbol{\theta}) = \lambda \|\boldsymbol{\theta}\|_1$.

Note that it is sufficient that one of these assumptions, either Case A or Case B, holds true.

A2 For positive constants $l, L > 0$, there exists a sequence of thrice differentiable functions $h^{(k)}$ with l -uniformly bounded second derivatives and L -uniformly bounded third derivatives. The sequence $h^{(k)}$ converges uniformly in the limit of $k \rightarrow \infty$ to the test function h .

A3 The noise vector $\boldsymbol{\epsilon}$ has elements ϵ_i which are i.i.d with $\mathbb{E}[\epsilon_i] = 0$, $\mathbb{E}[\epsilon_i^2] = \sigma_\epsilon^2 < \infty$ and $\mathbb{E}[\epsilon_i^4] < \infty$.

A4 The dimensions n, m, d remain at constant ratio when they are increased to infinity. These ratios are given by $\gamma = \frac{n}{m}$, $\eta = \frac{n}{d}$ and $\delta = \gamma\eta = \frac{m}{d}$

A5 The true model $\boldsymbol{\theta}^*$ is independent of \mathbf{X} . We assume that for some constants $c, c', C > 0$, $\mathbb{P}\left(\frac{1}{\sqrt{m}} \max(\|\nabla r(\boldsymbol{\theta}^*)\|_2, \|\nabla h(\boldsymbol{\theta}^*)\|_2) > c\right) \rightarrow 0$, and $\mathbb{P}(\max_i |(\nabla h(\boldsymbol{\theta}^*))_i| \geq c \log m) \leq C e^{-c'(\log m)^2}$.

A6 The activation function $\sigma(\cdot)$ is odd, with bounded first, second, and third derivatives.

Given the assumptions, we state the values of the bounds on τ_1 and τ_2 :

$$|\tau_1| \leq \tau_1^* = \frac{\mu/8}{\rho_1^2(1 + 2\sqrt{\delta})^2 + \rho_*^2} \quad |\tau_2| \leq \tau_2^* = \frac{\mu}{4l}, \quad (22)$$

¹Note that for a generic multi-variable function, the derivatives are tensors and we refer to their operator norm for bounds. However, as the functions are separable, i.e. a scalar function is applied element-wise, the bounds are simply on the derivatives of the scalar function.

where the values of μ and l are given in A1 and A2, respectively. Both of these bounds are chosen to ensure that the sum of the regularization function and the two “ τ terms” remains strongly convex with high probability.

4.3 Asymptotic Gaussian Results

In this section, we state our main result connecting P_2 in (10) and P_3 in (14).

Theorem 2. *Let Assumptions A3-A5 hold and $r + \tau_2 h$ is $\frac{\mu}{2}$ -strongly convex for $\tau_2 \in [-\tau_2^*, \tau_2^*]$. Then for all $\tau_1 \in [-\tau_1^*, \tau_1^*]$ and $\tau_2 \in [-\tau_2^*, \tau_2^*]$,*

$$\left| \tilde{P}_2(\tau_1, \tau_2) - \tilde{P}_3(\tau_1, \tau_2) \right| \xrightarrow[n, m, d \rightarrow \infty]{P} \mathbf{0} \quad (23)$$

Moreover,

$$\left| \left(\mathcal{E}_{train}(\hat{\boldsymbol{\theta}}_2), \mathcal{E}_{train}(\hat{\boldsymbol{\theta}}_2), \frac{1}{m} h(\hat{\boldsymbol{\theta}}_2) \right) - \left(\tilde{\mathcal{E}}_{train}, \tilde{\mathcal{E}}_{gen}, \frac{1}{m} h(\hat{\boldsymbol{\theta}}_3) \right) \right| \xrightarrow[n, m, d \rightarrow \infty]{P} \mathbf{0} \quad (24)$$

where $\hat{\boldsymbol{\theta}}_2$ is the solution to problem (10) and $\hat{\boldsymbol{\theta}}_3$ is the solution presented in (15) associated with P_3 in (14).

This result makes the statement in the second equation of (17) precise. Note that we do not need A6 and the assumption for r, h is weaker than the combination of A1-Case A and A2. A6, A1-Case A and A2 are required for the next step concerning P_1 . The results for A1-Case B will be obtained from the study of A1-Case A, in a suitable limit. Note that for this result, ρ_1 and ρ^* in (9) can be arbitrary, but we will set them to the values discussed in text following (9) for the subsequent results.

4.3.1 Proof Sketch of Theorem 2

The proof of this statement makes use of the Convex Gaussian Min Max Theorem (CGMT), which establishes an asymptotic equivalence between a primary (P) and an alternative (A) optimization problem of the following form:

$$P(\mathbf{A}) = \min_{\mathbf{x} \in S_x} \max_{\mathbf{y} \in S_y} \mathbf{x}^T \mathbf{A} \mathbf{y} + \psi(\mathbf{x}, \mathbf{y}) \quad (25)$$

$$A(\mathbf{g}, \mathbf{h}) = \min_{\mathbf{x} \in S_x} \max_{\mathbf{y} \in S_y} \|\mathbf{y}\|_2 \mathbf{x}^T \mathbf{g} + \|\mathbf{x}\|_2 \mathbf{y}^T \mathbf{h} + \psi(\mathbf{x}, \mathbf{y}) \quad (26)$$

Here, $\mathbf{A} \in \mathbb{R}^{m \times n}$, $\mathbf{g} \in \mathbb{R}^m$, $\mathbf{h} \in \mathbb{R}^n$ have i.i.d standard Gaussian elements, $\psi(\mathbf{x}, \mathbf{y})$ is an arbitrary convex-concave function, and $S_x \subset \mathbb{R}^m$, $S_y \subset \mathbb{R}^n$ are compact and convex sets. For more details, see supplement A. To prove Theorem 2, we first fix \mathbf{W} and $\boldsymbol{\theta}^*$ and translate the original minimization problem into a min-max problem of the form in (25) by suitable transformations and change of variables. Then, we invoke the CGMT which eliminates the

randomness (in \mathbf{X}) due to the data set \mathbf{z}^2 and re-express the problem in terms of (26). The resulting expression is given in (Dhifallah and Lu, 2020; Loureiro et al., 2021), but it is well-known to be intractable as it depends on the covariance matrix of the Gaussian feature map. Here, we introduce a key novel step. We show that assuming random weights \mathbf{W} , under further suitable, non-trivial transformations, the resulting equivalent form in (26) itself can be transformed into the form of (25) with a new random matrix \mathbf{A} representing the randomness of the weights. This allows us to apply the CGMT again, resulting in the elimination of the random matrix \mathbf{W} . Finally, we simplify the expressions obtained by the second CGMT application, which leads to the results in Theorem 2. The full proof is given in the Appendix B.

4.4 Universality

Next, we demonstrate universality. Here we show that the solution vectors problems P_1 given in (7) and problem P_2 given in (10) result in asymptotically equivalent values, not only in the training and generalization error, but also in a wide family of other test functions h . We provide two novel theorems, in this sections, that extend the existing results for the universality of random feature models. For completeness we first state the existing results by Hu and Lu (2020).

Theorem 3 (Hu and Lu, 2020) Theorem 1, Proposition 1). *Let assumptions A3-A5 hold. Set $\tau_2 = 0$ and let $r(\boldsymbol{\theta})$ be a regularization function that is strongly convex and thrice differentiable with uniformly bounded third derivatives. Let $\hat{\boldsymbol{\theta}}_1, \hat{\boldsymbol{\theta}}_2$ be the optimal solution to the problems given in (7) and (10), respectively. Then for all $\tau_1 \in [-\tau_1^*, \tau_1^*]$,*

$$\tilde{P}_1(\tau_1, 0) \rightarrow \tilde{P}_2(\tau_1, 0) \quad (27)$$

As a result,

$$\left| \left(\mathcal{E}_{train}(\hat{\boldsymbol{\theta}}_1), \mathcal{E}_{gen}(\hat{\boldsymbol{\theta}}_1) \right) - \left(\mathcal{E}_{train}(\hat{\boldsymbol{\theta}}_2), \mathcal{E}_{gen}(\hat{\boldsymbol{\theta}}_2) \right) \right| \xrightarrow[n, m, d \rightarrow \infty]{P} \mathbf{0} \quad (28)$$

Remark 1. The statement of the Theorem 3 is adapted to the particular setup that we consider here. For completeness the original theorem is given in appendix C as Theorem 9.

We are now ready to present our contribution. Firstly, we demonstrate the following theorem, relaxing the condition on the regularizer in Theorem 3, to A1-Case A and extending the result to an arbitrary test function h :

Theorem 4. *Let A2-A6 hold and the regularization function r satisfies A1-Case A. Then for all $\tau_1 \in [-\tau_1^*, \tau_1^*]$ and*

²This means that the terms including the random matrix \mathbf{A} in $P(\mathbf{A})$ will be removed and replaced by terms including random vectors \mathbf{g}, \mathbf{h} in $A(\mathbf{g}, \mathbf{h})$.

$$\tau_2 \in [-\tau_2^*, \tau_2^*],$$

$$\left| \tilde{P}_1(\tau_1, \tau_2) - \tilde{P}_2(\tau_1, \tau_2) \right| \xrightarrow[n, m, d \rightarrow \infty]{P} 0 \quad (29)$$

As a result,

$$\left| \left(\mathcal{E}_{train}(\hat{\boldsymbol{\theta}}_1), \mathcal{E}_{gen}(\hat{\boldsymbol{\theta}}_1), \frac{1}{m} h(\hat{\boldsymbol{\theta}}_1) \right) \left(\mathcal{E}_{train}(\hat{\boldsymbol{\theta}}_2), \mathcal{E}_{gen}(\hat{\boldsymbol{\theta}}_2), \frac{1}{m} h(\hat{\boldsymbol{\theta}}_2) \right) \right| \xrightarrow[n, m, d \rightarrow \infty]{P} \mathbf{0} \quad (30)$$

The next result illustrates that universality can also be applied to elastic net regularization

Corollary 1. *Let A2-A5 hold. Let $r(\boldsymbol{\theta}) = \lambda \|\boldsymbol{\theta}\|_1 + \frac{\mu}{2} \|\boldsymbol{\theta}\|_2^2$. Then, the claims of Theorem 4 hold true.*

4.4.1 Proof Sketch of Theorem 4

The original proof given by Hu and Lu (2020) is valid only for regularization functions that are strongly convex and thrice differentiable with uniformly bounded third derivatives. We first extend these results to sequence of regularization functions $r^{(k)}$ that converge uniformly to a function r . Noting that this theorem holds for all $r^{(k)}$ with $k < \infty$ the proof consists of demonstrating that the relations hold in the limit. Second, Hu and Lu (2020) does not consider the term $\tau_2 h(\boldsymbol{\theta})$. We adopt the original proof of Hu and Lu (2020) and modify it to demonstrate that the results similarly hold with a more generic test function $h(\boldsymbol{\theta})$. The proof of these results are given in the Appendix C.3.

For the specific case of elastic net, we construct a valid sequence $r^{(k)}(\boldsymbol{\theta})$ that uniformly converges to the elastic net regularization function, see Appendix C.4.

4.5 Random Features and Scalar Optimization Problem

We now connect the original problem P_1 to the scalar optimization problem P_3 by combining the results in Section 4.3 and Section 4.4. This leads to the following precise statement of the main result in Theorem 1:

Theorem 5. *Let Assumptions A2 - A6 and A1. Case A hold. Then for all $\tau_1 \in [-\tau_1^*, \tau_1^*]$ and $\tau_2 \in [-\tau_2^*, \tau_2^*]$,*

$$\left| \tilde{P}_1(\tau_1, \tau_2) - \tilde{P}_3(\tau_1, \tau_2) \right| \xrightarrow[n, m, d \rightarrow \infty]{P} 0 \quad (31)$$

Moreover,

$$\left| \left(\mathcal{E}_{train}(\hat{\boldsymbol{\theta}}_1), \mathcal{E}_{train}(\hat{\boldsymbol{\theta}}_1), \frac{1}{m} h(\hat{\boldsymbol{\theta}}_1) \right) \left(\tilde{\mathcal{E}}_{train}, \tilde{\mathcal{E}}_{gen}, \frac{1}{m} h(\hat{\boldsymbol{\theta}}_3) \right) \right| \xrightarrow[n, m, d \rightarrow \infty]{P} \mathbf{0} \quad (32)$$

4.6 Results for ℓ_1 regularization

We further extend these results to the case of ℓ_1 regularization. For this case, additional assumptions are needed. In particular, we may only consider scenarios, where problem (2) is sufficiently sparse. This is defined by the following:

$$M_0 = \frac{1}{m} \sum_{i=1}^m \Pr(\hat{\theta}_{i,3} \neq 0) \quad (33)$$

where $\hat{\theta}_{i,3}$ denotes the i^{th} element of $\hat{\boldsymbol{\theta}}_3$ for the regularization function $r(\boldsymbol{\theta}) = \lambda \|\boldsymbol{\theta}\|_1$. We prove the following theorem:

Theorem 6. *Let Assumptions A2 - A6 hold and $r(\boldsymbol{\theta}) = \lambda \|\boldsymbol{\theta}\|_1$. There exists a constant ρ only depending on the activation function σ and the parameters of the problem $(\lambda, \sigma_\epsilon^2, \gamma, \eta)$ such that for $M_0 < \rho$, the results of Theorem 5 holds for $r(\boldsymbol{\theta}) = \lambda \|\boldsymbol{\theta}\|_1$.*

4.6.1 Proof Sketch of Theorem 6

We adopt the proof in Panahi and Hassibi (2017), which performs this procedure for i.i.d. sub-Gaussian features, and modify it for the random feature model. Extending the results for the ℓ_1 regularization involves the results for the elastic net optimization in corollary 1. In Panahi and Hassibi (2017) (section 3.3 of supplement), it is shown that for a small value of μ in the elastic net regularization $\lambda \|\cdot\|_1 + \frac{\mu}{2} \|\cdot\|_2^2$, the ℓ_2 term can be removed and the change of the solution is negligible, if the matrix \mathbf{X} satisfies a proper restricted isometry property (RIP). In Panahi and Hassibi (2017) (lemma 8 in supplement), the RIP is shown for i.i.d. sub-Gaussian features. We extend this result and show that a similar RIP condition holds for random features model. The condition on M_0 ensures that the optimal solution is sufficiently stable, which otherwise is not guaranteed with the lack of strong convexity. The full proof is presented in Appendix C.5.

5 Elastic Net Regularization

In this section, we apply our results to the case of elastic net regularization, for which asymptotic learning curves has not been previously proposed. We consider the regularization function

$$r(\boldsymbol{\theta}) = \lambda \|\boldsymbol{\theta}\|_1 + \frac{\alpha}{2} \|\boldsymbol{\theta}\|_2^2, \quad (34)$$

where λ and α are two regularization parameters. We note that in the case of $\lambda = 0$ we obtain ridge regularization and in the case of $\alpha = 0$ we obtain ℓ_1 regularization (LASSO). Due to the continuity of the asymptotic expressions, the analysis of elastic net may be directly used for the study of ridge or LASSO regression simply by setting either $\lambda = 0$ or $\alpha = 0$. Our interest in studying elastic net stems from

the sparsity-promoting effect of the ℓ_1 regularizer on the solution vector. When viewing the RF model as a shallow neural network, the effect of a sparse solution is to disable a number of nodes in the hidden layer. As a result, elastic net finds a subnetwork of the original NN with a minimal degradation in performance, in effect a form of network compression. For similar attempts, see for example (Tang et al., 2022; Oyedotun et al., 2021; Yu et al., 2014).

The asymptotic equivalent solution to the elastic net regularized problem is given by

$$(\hat{\boldsymbol{\theta}}_3)_i = \begin{cases} \frac{2\tilde{c}_1\theta_i^*}{2\tilde{c}_1+\alpha} + \frac{\tilde{c}_2\sqrt{\gamma}}{(2\tilde{c}_1+\alpha)}\phi_i - \frac{\lambda}{2\tilde{c}_1+\alpha} & \phi_i < -\zeta_{1i} \\ \frac{2\tilde{c}_1\theta_i^*}{2\tilde{c}_1+\alpha} + \frac{\tilde{c}_2\sqrt{\gamma}}{(2\tilde{c}_1+\alpha)}\phi_i + \frac{\lambda}{2\tilde{c}_1+\alpha} & \phi_i > \zeta_{2i} \\ 0 & -\zeta_{1i} \leq \phi_i \leq \zeta_{2i} \end{cases}, \quad (35)$$

in which ζ_{1i} and ζ_{2i} are given by

$$\zeta_{1i} = \frac{(\lambda - 2\tilde{c}_1\theta_i^*)}{\sqrt{\gamma}\tilde{c}_2} \quad \zeta_{2i} = \frac{(\lambda + 2\tilde{c}_1\theta_i^*)}{\sqrt{\gamma}\tilde{c}_2} \quad (36)$$

and \tilde{c}_1, \tilde{c}_2 are the constants described in Theorem 2. The solution may also be expressed more succinctly by means of a soft thresholding operator. A full derivation of this solution may be found in the supplement section D. We note that in the limit of $\lambda \rightarrow 0$, we obtain $-\zeta_{1i} = \zeta_{2i}$ and the solution collapses into a single case, that being the result for ridge regression.

According to theorem 5 and 6, the characteristics of the solution vector $\hat{\boldsymbol{\theta}}_1$, reflected by a suitable function h , asymptotically becomes close to that of $\hat{\boldsymbol{\theta}}_3$ (the ℓ_1 case is under sparsity condition). Here, we consider the sparsity of the solution. For this reason, we take a separable function $h_\epsilon(\boldsymbol{\theta}) = \sum_i \bar{h}_\epsilon(\theta_i)$, where $\bar{h}_\epsilon(\theta)$ is a positive C^∞ bump function such that $\bar{h}_\epsilon(0) = 1$ and $\bar{h}_\epsilon(\theta) = 0$ for $|\theta| > \epsilon$. Our results apply to this function and we note that

$$n_0(\boldsymbol{\theta}) \leq h_\epsilon(\boldsymbol{\theta}) \leq n_\epsilon(\boldsymbol{\theta}), \quad (37)$$

where $n_\epsilon(\boldsymbol{\theta})$ is the number of the elements θ_i in $\boldsymbol{\theta}$ with $|\theta_i| \leq \epsilon$. In particular, n_0 is the number of zeros. We may show that by theorem 5 and the law of large numbers, the value of $\frac{1}{m} h_\epsilon(\boldsymbol{\theta})$ converges in probability to a constant s_ϵ calculated by analyzing $\hat{\boldsymbol{\theta}}_3$. We refer to $s := \lim_{\epsilon \rightarrow 0} s_\epsilon$ as the "effective sparsity" of $\hat{\boldsymbol{\theta}}_1$. Roughly speaking, s counts not only the zero entries of $\hat{\boldsymbol{\theta}}_1$, but also the vanishing entries as the problem size grows.

By direct calculation, we shown in the supplement section D that

$$s \rightarrow \frac{2\tilde{c}_1 + \alpha}{\sqrt{\gamma}\tilde{c}_2} \frac{1}{m} \sum_i \mathbb{E} \left[\left(\hat{\boldsymbol{\theta}}_3 \right)_i \phi_i \right], \quad (38)$$

where ϕ_i and $\left(\hat{\boldsymbol{\theta}}_3 \right)_i$ are defined in (35). We note that for pure ℓ_1 regularization this formula may still be used by setting $\alpha = 0$, although we can theoretically support it for

small values of sparsity. In this case, $s = 1 - M_0$ where M_0 is given in (33). Experimental results for effective sparsity maybe be found in the supplement E.1

6 EXPERIMENTS

6.1 Experimental setup

Using the expressions derived in the previous section, we examine the case of elastic net regularization experimentally. We choose the tanh activation for the non linearity of the feature map. We consider a deterministic vector θ^* that consists of half ones and half zeros, . We set the noise power $\sigma_\epsilon^2 = 0.1$ and let $\delta = 1$. We consider multiple cases, where for each case we solve the problem P_3 (equation (14)) using an iterative refining grid search algorithm. We compare the results to an experimental simulation in which $n + m = 1000$, with the relative ratio varied for different values of $\gamma = m/n$. Each empirical data point was averaged over 100 random realizations of the weights \mathbf{W} , and the data \mathbf{z} . More details maybe found in appendix E.

6.2 Elastic net model


We compare the experimental and theoretically derived values for training and generalization error of the elastic net model for two cases. Firstly we vary the ratio $\gamma = \frac{m}{n}$ for fixed values of the regularization parameters, and secondly we vary the regularization parameter λ for all other parameters being fixed.

The case of varying γ is shown in figure 1. Here, we fix $\lambda = 10^{-3}$ and choose several values of α including 0, the case of pure ℓ_1 regularization. Our expressions accurately predict the expected behavior of a network, the small deviation explained by the fact that n, m are finite. However, the discrepancy is only notable in a small range near the interpolation peak, suggesting the validity of our expressions in a wide range of networks of a non asymptotic size. We observe that small values of α result in a spike in the generalization error at the interpolation threshold, which in this model, is slightly more than $\gamma = 1$. We note that as the regularization parameter increases in strength, the interpolation peak diminishes. This is consistent with other results on the study of the double descent phenomenon (d’Ascoli et al., 2020).


In figure 2, we choose $\alpha = 10^{-3}$ and vary the value of the regularization parameter λ at constant γ . We note that that the generalization error suggests that at each ratio of $\gamma = \frac{m}{n}$ there is an optimal value of λ that minimizes the expected error.

7 CONCLUSION

We derived expressions to determine the exact asymptotic learning curves for square loss random feature models, subject to strongly convex regularization, or ℓ_1 regularization. These expressions consist of a 4-dimensional scalar optimization with two min-max pairs that is computable using standard techniques. We proved in two steps that these expressions coincide with the asymptotic learning curves: First, we demonstrate that the scalar optimization is asymptotically equivalent to a surrogate Gaussian model whose first two moments match that of the RF models. For this, we proposed a novel multi-stage application of the CGMT. Then, we extended the results of the universality of RF models to a broader family, including elastic net and ℓ_1 regularization, thereby demonstrating an asymptotic equivalence between the Gaussian model and the non linear RF



(a) Training Error



(b) Generalization Error

Figure 1: Theoretically predicted (solid line) and numerically determined (markers) values of the training error (a) and generalization error (b) for the random features model with $\ell_1 + \ell_2$ regularization as a function of $\gamma = \frac{m}{n}$, for varying values of regularization strengths of α at constant value of $\lambda = 10^{-3}$.

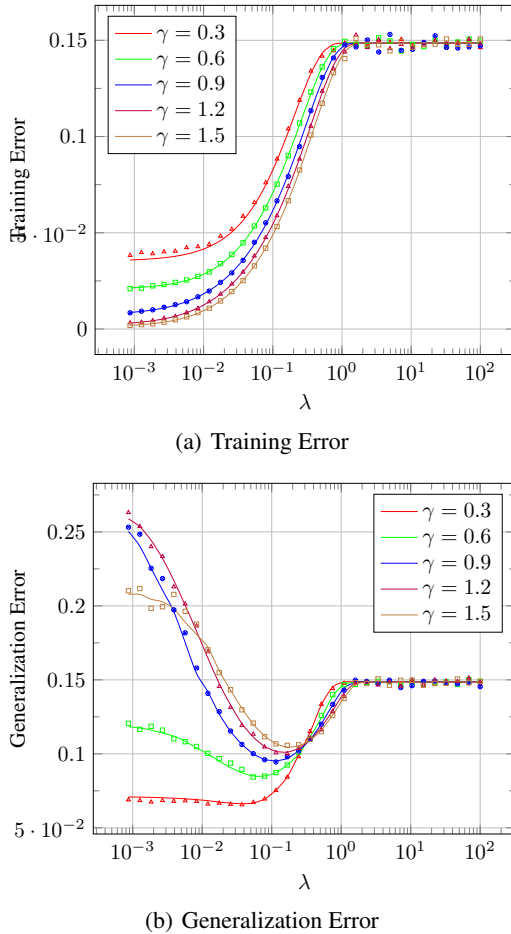


Figure 2: Theoretically predicted (solid line) and numerically determined (markers) values of the training error (a) and generalization error (b) for the random features model with $\ell_1 + \ell_2$ regularization as a function of the regularization parameter λ , for varying values of the ratio $\gamma = \frac{m}{n}$ constant value of $\alpha = 10^{-3}$.

model. Our results for universality hold not only for the cases of training and generalization error, but also for test functions h from a wide family.

There are several potential directions to extend our study. A particularly interesting direction is to use our methodology to obtain refined expressions for a more generic loss functions, extending the existing studies, e.g Loureiro et al. (2021).

References

Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning (still) requires rethinking generalization. *Communications of the ACM*, 64(3):107–115, 2021.

Mikhail Belkin, Daniel Hsu, Siyuan Ma, and Soumik Mandal. Reconciling modern machine-learning practice and

the classical bias–variance trade-off. *Proceedings of the National Academy of Sciences*, 116(32):15849–15854, 2019.

Mikhail Belkin, Daniel Hsu, and Ji Xu. Two models of double descent for weak features. *SIAM Journal on Mathematics of Data Science*, 2(4):1167–1180, 2020.

Trevor Hastie, Andrea Montanari, Saharon Rosset, and Ryan J Tibshirani. Surprises in high-dimensional ridgeless least squares interpolation. *arXiv preprint arXiv:1903.08560*, 2019.

Peter L. Bartlett, Philip M. Long, Gábor Lugosi, and Alexander Tsigler. Benign overfitting in linear regression. *Proceedings of the National Academy of Sciences*, 117(48):30063–30070, 2020a. ISSN 0027-8424. doi: 10.1073/pnas.1907378117. URL <https://www.pnas.org/content/117/48/30063>.

A. Tsigler and P. L. Bartlett. Benign overfitting in ridge regression. *arXiv e-prints*, art. arXiv:2009.14286, September 2020.

Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In *NIPS*, pages 1177–1184, 2007.

Song Mei and Andrea Montanari. The generalization error of random features regression: Precise asymptotics and the double descent curve. *Communications on Pure and Applied Mathematics*, n/a(n/a), 2019. doi: <https://doi.org/10.1002/cpa.22008>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/cpa.22008>.

Sebastian Goldt, Marc Mézard, Florent Krzakala, and Lenka Zdeborová. Modeling the influence of data structure on learning in neural networks: The hidden manifold model. *Physical Review X*, 10(4):041044, 2020a.

Stéphane d’Ascoli, Maria Refinetti, Giulio Biroli, and Florent Krzakala. Double trouble in double descent: Bias and variance (s) in the lazy regime. In *International Conference on Machine Learning*, pages 2280–2290. PMLR, 2020.

Oussama Dhifallah and Yue M. Lu. A precise performance analysis of learning with random features, 2020. URL <https://arxiv.org/abs/2008.11904>.

Christos Thrampoulidis, Samet Oymak, and Babak Hassibi. The Gaussian min-max theorem in the Presence of Convexity. *arXiv e-prints*, art. arXiv:1408.4837, August 2014.

Andrea Montanari, Feng Ruan, Youngtak Sohn, and Jun Yan. The generalization error of max-margin linear classifiers: High-dimensional asymptotics in the overparametrized regime. *arXiv preprint arXiv:1911.01544*, 2019.

Sebastian Goldt, Galen Reeves, Marc Mézard, Florent Krzakala, and Lenka Zdeborová. The gaussian equiv-

- alence of generative models for learning with two-layer neural networks. 2020b.
- Ashkan Panahi and Babak Hassibi. A universal analysis of large-scale regularized least squares solutions. In *NIPS*, pages 3384–3393, 2017.
- Samet Oymak and Joel A Tropp. Universality laws for randomized dimension reduction, with applications. *Information and Inference: A Journal of the IMA*, 7(3):337–446, 2018.
- Hong Hu and Yue M. Lu. Universality laws for high-dimensional learning with random features. *CoRR*, abs/2009.07669, 2020. URL <https://arxiv.org/abs/2009.07669>.
- Bruno Loureiro, Cédric Gerbelot, Hugo Cui, Sebastian Goldt, Florent Krzakala, Marc Mézard, and Lenka Zdeborová. Learning curves of generic features maps for realistic datasets with a teacher-student model, 2021. URL <https://arxiv.org/abs/2102.08127>.
- Xiangyu Chang, Yingcong Li, Samet Oymak, and Christos Thrampoulidis. Provable benefits of overparameterization in model compression: From double descent to pruning neural networks. *arXiv preprint arXiv:2012.08749*, 2020.
- Hui Zou and Trevor Hastie. Regularization and variable selection via the elastic net. *Journal of the royal statistical society: series B (statistical methodology)*, 67(2): 301–320, 2005.
- Marco Loog, Tom Viering, Alexander Mey, Jesse H Krijthe, and David MJ Tax. A brief prehistory of double descent. *Proceedings of the National Academy of Sciences*, 117(20):10625–10626, 2020.
- Mario Geiger, Arthur Jacot, Stefano Spigler, Franck Gabriel, Levent Sagun, Stéphane d’Ascoli, Giulio Biroli, Clément Hongler, and Matthieu Wyart. Scaling description of generalization with number of parameters in deep learning. *Journal of Statistical Mechanics: Theory and Experiment*, 2020(2):023401, 2020.
- Peter L. Bartlett, Philip M. Long, Gábor Lugosi, and Alexander Tsigler. Benign overfitting in linear regression. *arxiv:1906.11300*, 2020b.
- Vidya Muthukumar, Kailas Vodrahalli, Vignesh Subramanian, and Anant Sahai. Harmless interpolation of noisy data in regression, 2019. URL <https://arxiv.org/abs/1903.09139>.
- Dmitry Kobak, Jonathan Lomond, and Benoit Sanchez. The optimal ridge penalty for real-world high-dimensional data can be zero or negative due to the implicit ridge regularization. *J. Mach. Learn. Res.*, 21:169–1, 2020.
- Zeyu Deng, Abla Kammoun, and Christos Thrampoulidis. A model of double descent for high-dimensional binary linear classification. *arXiv preprint arXiv:1911.05822*, 2019.
- Hossein Taheri, Ramtin Pedarsani, and Christos Thrampoulidis. Fundamental limits of ridge-regularized empirical risk minimization in high dimensions. In *International Conference on Artificial Intelligence and Statistics*, pages 2773–2781. PMLR, 2021.
- Panagiotis Lolias. Regularization in high-dimensional regression and classification via random matrix theory. *arXiv preprint arXiv:2003.13723*, 2020.
- Francesca Mignacco, Florent Krzakala, Yue Lu, Pierfrancesco Urbani, and Lenka Zdeborova. The role of regularization in classification of high-dimensional noisy gaussian mixture. In *International Conference on Machine Learning*, pages 6874–6883. PMLR, 2020.
- Ganesh Ramachandra Kini and Christos Thrampoulidis. Analytic study of double descent in binary classification: The impact of loss. In *2020 IEEE International Symposium on Information Theory (ISIT)*, pages 2527–2532. IEEE, 2020.
- Tengyuan Liang and Pragma Sur. A precise high-dimensional asymptotic theory for boosting and minimum- ℓ_1 -norm interpolated classifiers. *arXiv preprint arXiv:2002.01586*, 2020.
- Hossein Taheri, Ramtin Pedarsani, and Christos Thrampoulidis. Sharp asymptotics and optimal performance for inference in binary models. In *International Conference on Artificial Intelligence and Statistics*, pages 3739–3749. PMLR, 2020.
- Fariborz Salehi, Ehsan Abbasi, and Babak Hassibi. The impact of regularization on high-dimensional logistic regression. *Advances in Neural Information Processing Systems*, 32, 2019.
- Yehoram Gordon. Some inequalities for gaussian processes and applications. *Israel Journal of Mathematics*, 50(4): 265–289, 1985.
- Yehoram Gordon. On milman’s inequality and random subspaces which escape through a mesh in \mathbb{R}^n . In *Geometric aspects of functional analysis*, pages 84–106. Springer, 1988.
- Christos Thrampoulidis, Ashkan Panahi, and Babak Hassibi. Asymptotically exact error analysis for the generalized equation-lasso. In *2015 IEEE International Symposium on Information Theory (ISIT)*, pages 2021–2025. IEEE, 2015.
- David Bosch, Ashkan Panahi, and Ayca Özcelikkale. Double descent in feature selection: Revisiting lasso and basis pursuit. In *International Conference on Machine Learning (ICML) 2021 Workshop on Overparameterization: Pitfalls & Opportunities*, 2021.
- Christos Thrampoulidis, Ehsan Abbasi, and Babak Hassibi. Precise error analysis of regularized m-estimators

in high-dimensions, 2016. URL <https://arxiv.org/abs/1601.06233>.

Sebastian Goldt, Bruno Loureiro, Galen Reeves, Florent Krzakala, Marc Mézard, and Lenka Zdeborová. The gaussian equivalence of generative models for learning with shallow neural networks. In *Mathematical and Scientific Machine Learning*, pages 426–471. PMLR, 2022.

Andrea Montanari and Basil Saeed. Universality of empirical risk minimization, 2022. URL <https://arxiv.org/abs/2202.08832>.

Jimmy Ba, Murat A. Erdogdu, Taiji Suzuki, Zhichao Wang, Denny Wu, and Greg Yang. High-dimensional asymptotics of feature learning: How one gradient step improves the representation, 2022. URL <https://arxiv.org/abs/2205.01445>.

Anda Tang, Pei Quan, Lingfeng Niu, and Yong Shi. A survey for sparse regularization based compression methods. *Annals of Data Science*, pages 1–28, 2022.

Oyebade K Oyedotun, Abd El Rahman Shabayek, Djamila Aouada, and Björn Ottersten. Deep network compression with teacher latent subspace learning and lasso. *Applied Intelligence*, 51(2):834–853, 2021.

Jun Yu, Yong Rui, and Dacheng Tao. Click prediction for web image reranking using multimodal sparse coding. *IEEE Transactions on Image Processing*, 23(5):2019–2032, 2014.

Omiros Papaspiliopoulos. High-dimensional probability: An introduction with applications in data science, 2020.

Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration inequalities: A nonasymptotic theory of independence*. Oxford university press, 2013.

Heinz H Bauschke, Patrick L Combettes, et al. *Convex analysis and monotone operator theory in Hilbert spaces*, volume 408. Springer, 2011.

Jean Honorio and Tommi Jaakkola. Tight bounds for the expected risk of linear classifiers and pac-bayes finite-sample guarantees. In *Artificial Intelligence and Statistics*, pages 384–392. PMLR, 2014.

Emmanuel J Candes, Justin K Romberg, and Terence Tao. Stable signal recovery from incomplete and inaccurate measurements. *Communications on Pure and Applied Mathematics: A Journal Issued by the Courant Institute of Mathematical Sciences*, 59(8):1207–1223, 2006.

A Gaussian Min Max Theorems

We make use of the Gaussian Min max theorem as well as the Convex Gaussian min max theorem in the proof of theorem 1. The Gaussian min max theorem was originally proven by Gordon (Gordon, 1985, 1988). The CGMT was developed by (Thrapoulidis et al., 2014), we state the theorem here for completeness.

The Gaussian Min Max theorem states the following:

Theorem 7. *Let $\mathbf{A} \in \mathbb{R}^{m \times n}$, $g \in \mathbb{R}$, $\mathbf{g} \in \mathbb{R}^m$ and $\mathbf{h} \in \mathbb{R}^n$ be independent of each other and have entries distributed i.i.d according to $\mathcal{N}(0, 1)$. Let $\mathcal{S}_1 \subset \mathbb{R}^n$ and $\mathcal{S}_2 \subset \mathbb{R}^m$ be nonempty compact sets. Let $f(\cdot, \cdot)$ be a continuous function on $\mathcal{S}_1 \times \mathcal{S}_2$. We define*

$$\mathbf{P}_1(\mathbf{A}) := \min_{\mathbf{x} \in \mathcal{S}_1} \max_{\mathbf{y} \in \mathcal{S}_2} \mathbf{y}^T \mathbf{A} \mathbf{x} + g \|\mathbf{x}\|_2 \|\mathbf{y}\|_2 + f(\mathbf{x}, \mathbf{y}), \quad (39)$$

$$\mathbf{P}_2(\mathbf{g}, \mathbf{h}) := \min_{\mathbf{x} \in \mathcal{S}_1} \max_{\mathbf{y} \in \mathcal{S}_2} \|\mathbf{x}\|_2 \mathbf{g}^T \mathbf{y} + \|\mathbf{y}\|_2 \mathbf{h}^T \mathbf{x} + f(\mathbf{x}, \mathbf{y}). \quad (40)$$

Then for any $c \in \mathbb{R}$:

$$\mathbb{P}(\mathbf{P}_1(\mathbf{A}, g) \leq c) \leq \mathbb{P}(\mathbf{P}_2(\mathbf{g}, \mathbf{h}) \leq c) \quad (41)$$

The Convex Gaussian Min Max theorem extends these results to the following:

Theorem 8. *Let $\mathbf{A} \in \mathbb{R}^{m \times n}$, $\mathbf{g} \in \mathbb{R}^m$ and $\mathbf{h} \in \mathbb{R}^n$ be independent of each other and have entries distributed i.i.d according to $\mathcal{N}(0, 1)$. Let $\mathcal{S}_1 \subset \mathbb{R}^n$ and $\mathcal{S}_2 \subset \mathbb{R}^m$ be nonempty compact sets. Let $f(\cdot, \cdot)$ be a continuous function on $\mathcal{S}_1 \times \mathcal{S}_2$. We define*

$$\mathbf{P}_1(\mathbf{A}) := \min_{\mathbf{x} \in \mathcal{S}_1} \max_{\mathbf{y} \in \mathcal{S}_2} \mathbf{y}^T \mathbf{A} \mathbf{x} + f(\mathbf{x}, \mathbf{y}), \quad (42)$$

$$\mathbf{P}_2(\mathbf{g}, \mathbf{h}) := \min_{\mathbf{x} \in \mathcal{S}_1} \max_{\mathbf{y} \in \mathcal{S}_2} \|\mathbf{x}\|_2 \mathbf{g}^T \mathbf{y} + \|\mathbf{y}\|_2 \mathbf{h}^T \mathbf{x} + f(\mathbf{x}, \mathbf{y}). \quad (43)$$

Then for any $c_1 \in \mathbb{R}$ we have that

$$\mathbb{P}(\mathbf{P}_1(\mathbf{A}) < c_1) \leq 2\mathbb{P}(\mathbf{P}_2(\mathbf{g}, \mathbf{h}) \leq c_1), \quad (44)$$

Under the further assumptions that \mathcal{S}_1 and \mathcal{S}_2 are convex sets and f is concave-convex on $\mathcal{S}_1 \times \mathcal{S}_2$ then for all $c_2 \in \mathbb{R}$

$$\mathbb{P}(\mathbf{P}_1(\mathbf{A}) > c_2) \leq 2\mathbb{P}(\mathbf{P}_2(\mathbf{g}, \mathbf{h}) \geq c_2). \quad (45)$$

We note that if in the limit of $n, m \rightarrow \infty$ the value of $\mathbf{P}_2(\mathbf{g}, \mathbf{h})$ concentrates on a value a then similarly $\mathbf{P}_1(\mathbf{A})$ converges to the same value.

B Proof of Theorem 2

To prove, theorem 2, we shall apply the CGMT (supplement theorem 8) to obtain an alternative problem formulation for (10). Subsequently, we will simplify the alternative problem, and then express it once again in the form that is suitable for a second CGMT application. Applying the CGMT for a second time, we obtain a second alternative problem. After simplifying this second alternative problem, we will demonstrate the results in Theorem 2. To begin with the first application of the CGMT, we fix \mathbf{W} and change the variable $\boldsymbol{\theta}$ in (10) to $\mathbf{e} = \boldsymbol{\theta} - \boldsymbol{\theta}^*$ to obtain

$$\tilde{P}_2(\tau_1, \tau_2) = \min_{\mathbf{e}} \frac{1}{2n} \|\boldsymbol{\epsilon} - \frac{1}{\sqrt{m}} \tilde{\mathbf{X}} \mathbf{e}\|_2^2 + \frac{1}{m} r(\mathbf{e} + \boldsymbol{\theta}^*) + \frac{\tau_1}{m} \mathbf{e}^T \tilde{\mathbf{R}} \mathbf{e} + \frac{\tau_2}{m} h(\mathbf{e} + \boldsymbol{\theta}^*), \quad (46)$$

note that the rows $\tilde{\mathbf{x}}_i$ of $\tilde{\mathbf{X}}$ are i.i.d, centered and Gaussian with the covariance matrix $\tilde{\mathbf{R}} = \frac{\rho_d^2}{d} \mathbf{W}^T \mathbf{W} + \rho_*^2 I$. Hence, we may write $\tilde{\mathbf{X}} = \mathbf{U} \tilde{\mathbf{R}}^{\frac{1}{2}}$ where \mathbf{U} has i.i.d. standard Gaussian entries. Next, using the Legendre transform of the square function, we may write (10) as

$$\tilde{P}_2(\tau_1, \tau_2) = \min_{\mathbf{e}} \max_{\boldsymbol{\lambda}} \frac{1}{n} \boldsymbol{\lambda}^T \boldsymbol{\epsilon} - \frac{1}{n\sqrt{m}} \boldsymbol{\lambda}^T \mathbf{U} \tilde{\mathbf{R}}^{1/2} \mathbf{e} - \frac{1}{2n} \|\boldsymbol{\lambda}\|_2^2 + \frac{1}{m} r(\mathbf{e} + \boldsymbol{\theta}^*) + \frac{1}{m} \tau_1 \mathbf{e}^T \tilde{\mathbf{R}} \mathbf{e} + \frac{1}{m} \tau_2 h(\mathbf{e} + \boldsymbol{\theta}^*) \quad (47)$$

In here $\tau_1 \in \mathbb{R}$ and $\tau_2 \in \mathbb{R}$ are constants and by the assumption, $r(\mathbf{e} + \boldsymbol{\theta}^*) + \tau_2 h(\mathbf{e} + \boldsymbol{\theta}^*)$ is $\frac{\mu}{2}$ -strongly convex. We require that τ_1 is chosen sufficiently small, to ensure that the entire optimization problem remains strongly convex in \mathbf{e} . In particular, we ensure that the term $B(\mathbf{e}) := r(\mathbf{e} + \boldsymbol{\theta}^*) + \tau_1 \mathbf{e}^T \tilde{\mathbf{R}} \mathbf{e} + \tau_2 h(\mathbf{e} + \boldsymbol{\theta}^*)$ is $\frac{\mu}{4}$ -strongly convex. First, we show that

$$|\tau_1| \leq \tau_1^* = \frac{\mu/8}{\rho_1^2(1 + 2\sqrt{\delta})^2 + \rho_*^2} \quad (48)$$

will satisfy this condition with high probability³. For this reason, we introduce the following lemma:

Lemma 1. Define $C_{\tilde{\mathbf{R}}} = \rho_1^2(1 + 2\sqrt{\delta})^2 + \rho_*^2$. For a random matrix \mathbf{W} with i.i.d. standard Gaussian entries and $\tilde{\mathbf{R}} = \frac{\rho_1^2}{d} \mathbf{W} \mathbf{W}^T + \rho_*^2 \mathbf{I}$, the following relation holds:

$$\mathbb{P} \left[\|\tilde{\mathbf{R}}\|_2 > C_{\tilde{\mathbf{R}}} \right] < 2e^{-cm} \quad (49)$$

for a universal constant $c > 0$, where $\|\cdot\|_2$ denotes the spectral norm.

Proof. We note that by the definition of $\tilde{\mathbf{R}}$, we have that

$$\|\tilde{\mathbf{R}}\|_2 = \left\| \frac{\rho_1^2}{d} \mathbf{W} \mathbf{W}^T + \rho_*^2 \mathbf{I} \right\|_2 \leq \frac{\rho_1^2}{d} \|\mathbf{W}\|_2^2 + \rho_*^2. \quad (50)$$

The elements of $\mathbf{W} \in \mathbb{R}^{m \times d}$ are i.i.d normally distributed. From a standard result in matrix theory (Papaspiliopoulos, 2020)[Corollary 7.3.3] we obtain

$$\mathbb{P} \left(\frac{1}{\sqrt{d}} \|\mathbf{W}\|_2 \geq 1 + \sqrt{m/d} + t \right) \leq 2e^{-cdt^2}. \quad (51)$$

Choosing $t = \sqrt{m/d}$ yields

$$\mathbb{P} \left(\frac{1}{\sqrt{d}} \|\mathbf{W}\|_2 \geq 1 + 2\sqrt{\delta} \right) \leq 2e^{-cm}, \quad (52)$$

where we recall that $\delta = \frac{m}{d}$. This provides the desired result. \square

According to lemma 1, the term $\mathbf{e}^T \tilde{\mathbf{R}} \mathbf{e}$ is $2C_{\tilde{\mathbf{R}}}$ -smooth, and hence for $\tau_1 \leq \frac{\mu}{2C_{\tilde{\mathbf{R}}}}$, the term B is $\frac{\mu}{2} - \frac{\mu}{4} = \frac{\mu}{4}$ -convex. This is the same as the condition in (48). Hence, in the rest of this proof we assume that B is strongly convex.

Next, we note that applying the CGMT requires that both $\boldsymbol{\lambda}$ and \mathbf{e} are in compact feasibility sets. Here we employ a similar strategy to (Thrampoulidis et al., 2016; Dhifallah and Lu, 2020; Loureiro et al., 2021) by showing that with high probability, the solutions of both the original problem and the alternative problem can be bound in fixed compact sets, hence restricting the optimizations to these sets will not affect the result. As a result, we may apply the CGMT.

Lemma 2. Consider the following two optimization problems that correspond to the primary optimization and to the first alternative optimization in CGMT.

$$\tilde{P}_{2,1} = \min_{\mathbf{e}} \max_{\boldsymbol{\lambda}} \frac{1}{n} \boldsymbol{\lambda}^T \boldsymbol{\epsilon} - \frac{1}{n\sqrt{m}} \boldsymbol{\lambda}^T \tilde{\mathbf{X}} \mathbf{e} - \frac{1}{2n} \|\boldsymbol{\lambda}\|_2^2 + \frac{1}{m} B(\mathbf{e}) \quad (53)$$

$$\tilde{P}_{2,2} = \min_{\mathbf{e}} \max_{\boldsymbol{\lambda}} \frac{1}{n} \boldsymbol{\lambda}^T \boldsymbol{\epsilon} - \frac{1}{n\sqrt{m}} \|\tilde{\mathbf{R}}^{1/2} \mathbf{e}\|_2 \boldsymbol{\lambda}^T \mathbf{g} - \frac{1}{n\sqrt{m}} \|\boldsymbol{\lambda}\|_2 \mathbf{h}^T \tilde{\mathbf{R}}^{1/2} \mathbf{e} - \frac{1}{2n} \|\boldsymbol{\lambda}\|_2^2 + \frac{1}{m} B(\mathbf{e}) \quad (54)$$

In these equations \mathbf{g}, \mathbf{h} are standard normal vectors of size n, m , respectively. Denote by $\tilde{\mathbf{e}}_{2,1}, \tilde{\mathbf{e}}_{2,2}$ the optimal solutions of $\tilde{P}_{2,1}$ and $\tilde{P}_{2,2}$, respectively. Furthermore, respectively denote by $\tilde{\boldsymbol{\lambda}}_1(\mathbf{e}), \tilde{\boldsymbol{\lambda}}_2(\mathbf{e})$ the solution of their inner optimization (over $\boldsymbol{\lambda}$) for a given vector \mathbf{e} . Let B be strongly convex with constant $\frac{\mu}{4}$ and $\max \{\|\nabla r(\boldsymbol{\theta}^*)\|, \|\nabla h(\boldsymbol{\theta}^*)\|\} = O(\sqrt{m})$. Then, there exist positive constants C_e, C_λ only depending on μ such that the following hold true:

³Throughout this paper, the term "high probability" means a probability converging to 1 as the problem size grows.

- The solutions $\tilde{\mathbf{e}}_{2,i}$ for $i = 1, 2$ satisfy

$$\lim_{m \rightarrow \infty} \mathbb{P} \left(\max \{ \|\tilde{\mathbf{e}}_{2,1}\|_2, \|\tilde{\mathbf{e}}_{2,2}\|_2 \} \leq C_e \sqrt{m} \right). \quad (55)$$

- It also holds that

$$\lim_{m \rightarrow \infty} \mathbb{P} \left(\sup_{\mathbf{e} \|\mathbf{e}\| \leq C_e \sqrt{m}} \max \{ \|\tilde{\boldsymbol{\lambda}}_1(\mathbf{e})\|, \|\tilde{\boldsymbol{\lambda}}_2(\mathbf{e})\| \} \leq C_\lambda \sqrt{m} \right) = 1 \quad (56)$$

Proof. We note that B is $\frac{\mu}{4}$ strongly convex. Solving for $\boldsymbol{\lambda}$ in both optimization, we may write the optimization over \mathbf{e} as

$$\min_{\mathbf{e}} F_i(\mathbf{e}) \quad i = 1, 2 \quad (57)$$

Where $F_i(\mathbf{e})$ is the optimal value over $\boldsymbol{\lambda}$. We note that setting $\boldsymbol{\lambda} = \mathbf{0}$ in both optimizations, we obtain that $F(\mathbf{e}) \geq \frac{1}{m} B(\mathbf{e})$. Then we see that

$$B(\mathbf{e}) \geq B(\mathbf{0}) + \mathbf{d}^T \mathbf{e} + \frac{\mu}{4} \|\mathbf{e}\|_2^2 \quad (58)$$

where $\mathbf{d} = \nabla B(\mathbf{0}) = \nabla r(\boldsymbol{\theta}^*) + \tau_2 \nabla h(\boldsymbol{\theta}^*)$ and by the assumption $\|\mathbf{d}\| = O(\sqrt{m})$.

For optimization P_1 , we note that

$$F(\mathbf{0}) = \frac{1}{m} B(\mathbf{0}) + \frac{1}{2n} \|\boldsymbol{\epsilon}\|_2^2. \quad (59)$$

This implies that for the optimal solution $\hat{\mathbf{e}}$ we have

$$\frac{1}{m} B(\mathbf{0}) + \frac{1}{2n} \|\boldsymbol{\epsilon}\|_2^2 = F(\mathbf{0}) \geq F(\tilde{\mathbf{e}}_1) \geq \frac{1}{m} B(\mathbf{0}) + \frac{1}{m} \mathbf{d}^T \tilde{\mathbf{e}}_1 + \frac{\mu}{4m} \|\tilde{\mathbf{e}}_1\|_2^2, \quad (60)$$

which yields

$$\frac{\mu}{4m} \left\| \tilde{\mathbf{e}}_1 + \frac{1}{\mu} \mathbf{d} \right\|_2^2 \leq \frac{1}{2n} \|\boldsymbol{\epsilon}\|_2^2 + \frac{1}{4\mu m} \|\mathbf{d}\|_2^2. \quad (61)$$

Then, we obtain

$$\|\tilde{\mathbf{e}}_1\|_2 \leq \left\| \frac{1}{\mu} \mathbf{d} \right\|_2 + \sqrt{\frac{2m}{n\mu} \|\boldsymbol{\epsilon}\|_2^2 + \frac{1}{\mu^2} \|\mathbf{d}\|_2^2}. \quad (62)$$

From the standard matrix theory (Papaspiliopoulos, 2020)[Theorem 2.8.1] we know that $\|\boldsymbol{\epsilon}\|_2^2 < cn$ for some c , with high probability. We observe that there must exist some constant C_{e_1} such that

$$\lim_{m \rightarrow \infty} \mathbb{P}(\|\tilde{\mathbf{e}}_1\|_2 \geq C_{e_1} \sqrt{m}) = 0. \quad (63)$$

Now we consider (54). Our strategy is similar to the previous case. We note that if we let $\beta = \|\boldsymbol{\lambda}\|_2$ we can solve the optimization over $\boldsymbol{\lambda}$ to obtain:

$$F(\mathbf{e}) = \max_{\beta \geq 0} \frac{\beta}{n} \left\| \boldsymbol{\epsilon} - \frac{1}{\sqrt{m}} \|\tilde{\mathbf{R}}^{1/2} \mathbf{e}\|_2 \mathbf{g} \right\| - \frac{\beta}{n\sqrt{m}} \mathbf{h}^T \tilde{\mathbf{R}}^{1/2} \mathbf{e} - \frac{\beta^2}{2n} + \frac{1}{m} B(\mathbf{e}). \quad (64)$$

The optimization is limited to $\beta \geq 0$. Hence, its optimal value will be increased when the constant is lifted, leading to a quadratic optimization and the following result

$$F(\mathbf{e}) \leq \frac{1}{m} B(\mathbf{e}) + \frac{1}{2n} \left(\left\| \boldsymbol{\epsilon} - \frac{1}{\sqrt{m}} \|\tilde{\mathbf{R}}^{1/2} \mathbf{e}\|_2 \mathbf{g} \right\| - \frac{\beta}{\sqrt{m}} \mathbf{h}^T \tilde{\mathbf{R}}^{1/2} \mathbf{e} \right)^2, \quad (65)$$

and in particular

$$F(\mathbf{0}) \leq \frac{1}{m} B(\mathbf{0}) + \frac{1}{2n} \|\boldsymbol{\epsilon}\|_2^2. \quad (66)$$

By applying the same inequality as in (60) we obtain that

$$\|\tilde{\mathbf{e}}_2\|_2 \leq \left\| \frac{1}{\mu} \mathbf{d} \right\| + \sqrt{\frac{2m}{n\mu} \|\boldsymbol{\epsilon}\|_2^2 + \frac{1}{\mu^2} \|\mathbf{d}\|_2^2}. \quad (67)$$

As such, we obtain

$$\lim_{m \rightarrow \infty} \mathbb{P}(\|\tilde{\mathbf{e}}_2\| \geq C_{e_2} \sqrt{m}) = 0. \quad (68)$$

Now, let $C_e = \max(C_{e_1}, C_{e_2})$, and use this to define the set $A_e = \{\mathbf{e} \in \mathbb{R}^m \mid \|\mathbf{e}\|_2 \leq C_e \sqrt{m}\}$.

Next, we note from the optimality condition of the inner optimization in Eq. (53) that

$$\tilde{\boldsymbol{\lambda}}_1(\mathbf{e}) = \boldsymbol{\epsilon} - \frac{1}{\sqrt{m}} \mathbf{U} \tilde{\mathbf{R}}^{1/2} \mathbf{e}. \quad (69)$$

As such, for all $\mathbf{e} \in A_e$ we have

$$\|\tilde{\boldsymbol{\lambda}}_1(\mathbf{e})\|_2 \leq \|\boldsymbol{\epsilon}\|_2 + \left\| \frac{1}{\sqrt{m}} \mathbf{U} \tilde{\mathbf{R}}^{1/2} \right\| \|\mathbf{e}\|_2 \leq \|\boldsymbol{\epsilon}\|_2 + \left\| \frac{1}{\sqrt{m}} \mathbf{U} \right\|_2 \|\tilde{\mathbf{R}}^{1/2}\|_2 \|\mathbf{e}\|_2 \quad (70)$$

We note from lemma 1 that $\|\mathbf{R}^{1/2}\|_2$ is bounded, and we make use of standard random matrix theory to conclude $\|\frac{1}{\sqrt{m}} \mathbf{U}\|_2 < C$ with high probability. Then, making use of the same arguments as before we can see that there must exist a constant C_{λ_1} such that for all $\mathbf{e} \in A_e$

$$\lim_{n \rightarrow \infty} \mathbb{P} \left(\sup_{\mathbf{e} \in A_e} \|\tilde{\boldsymbol{\lambda}}_1(\mathbf{e})\|_2 \geq C_{\lambda_1} \sqrt{n} \right) = 0 \quad (71)$$

Finally we note that the optimality condition over β of problem 54 gives that for all $\mathbf{e} \in A_e$

$$\begin{aligned} \hat{\beta} &= \|\tilde{\boldsymbol{\lambda}}_1(\mathbf{e})\|_2 = \left\| \boldsymbol{\epsilon} - \frac{1}{\sqrt{m}} \|\tilde{\mathbf{R}}^{1/2} \mathbf{e}\|_2 \mathbf{g} \right\|_2 - \frac{1}{\sqrt{m}} \tilde{\mathbf{R}}^{1/2} \mathbf{h} \\ &\leq \|\boldsymbol{\epsilon}\|_2 + \frac{1}{\sqrt{m}} \|\mathbf{g}\|_2 \|\tilde{\mathbf{R}}^{1/2}\|_2 \|\mathbf{e}\|_2 + \frac{1}{\sqrt{m}} \|\tilde{\mathbf{R}}^{1/2}\|_2 \|\mathbf{h}\|_2 \end{aligned} \quad (72)$$

We note that with high probability $\|\boldsymbol{\epsilon}\|_2 < C\sqrt{n}$, $\|\mathbf{g}\|_2 < C\sqrt{n}$ and $\|\mathbf{h}\|_2 < C\sqrt{m}$. From this we can see that there exists a constant C_{λ_2} such that

$$\lim_{n \rightarrow \infty} \mathbb{P} \left(\sup_{\mathbf{e} \in A_e} \|\tilde{\boldsymbol{\lambda}}_2(\mathbf{e})\|_2 \geq C_{\lambda_2} \sqrt{n} \right) = 0 \quad (73)$$

Taking $C_\lambda = \max(C_{\lambda_1}, C_{\lambda_2})$ completes the proof. \square

We use the definition of the sets $A_e = \{\mathbf{e} \in \mathbb{R}^m \mid \|\mathbf{e}\|_2 \leq C_e \sqrt{m}\}$ and $A_\lambda = \{\boldsymbol{\lambda} \in \mathbb{R}^n \mid \|\boldsymbol{\lambda}\|_2 \leq C_\lambda \sqrt{n}\}$ in the rest of this study. By the lemma above, we can with high probability, restrict ourselves to the following problem

$$\tilde{P}'_{2,1}(\tau_1, \tau_2) = \min_{\mathbf{e} \in A_e} \max_{\boldsymbol{\lambda} \in A_\lambda} \frac{1}{n} \boldsymbol{\lambda}^T \boldsymbol{\epsilon} - \frac{1}{n\sqrt{m}} \boldsymbol{\lambda}^T \mathbf{U} \tilde{\mathbf{R}}^{1/2} \mathbf{e} - \frac{1}{2n} \|\boldsymbol{\lambda}\|_2^2 + \frac{1}{m} B(\mathbf{e}), \quad (74)$$

and be certain that the solution vector and the optimal value to the problem \tilde{P}_2 will be equal to those of the problem $\tilde{P}'_{2,1}$. We now make use of the CGMT, (Thm. 8). From which we obtain the following optimization problem

$$\tilde{P}'_{2,2} = \min_{\mathbf{e} \in A_e} \max_{\boldsymbol{\lambda} \in A_\lambda} \frac{1}{n} \boldsymbol{\lambda}^T \boldsymbol{\epsilon} - \frac{1}{n\sqrt{m}} \|\tilde{\mathbf{R}}^{1/2} \mathbf{e}\|_2 \mathbf{g}^T \boldsymbol{\lambda} - \frac{1}{n\sqrt{m}} \|\boldsymbol{\lambda}\|_2 \mathbf{h}^T \tilde{\mathbf{R}}^{1/2} \mathbf{e} - \frac{1}{2n} \|\boldsymbol{\lambda}\|_2^2 + \frac{1}{m} B(\mathbf{e}) \quad (75)$$

In which $\mathbf{g} \sim \mathcal{N}(0, I_n)$ and $\mathbf{h} \sim \mathcal{N}(0, I_m)$. Note that by lemma (2), $\tilde{P}'_{2,1}$ is also identical to $\tilde{P}'_{2,2}$. We now let $\beta = \frac{1}{\sqrt{n}} \|\boldsymbol{\lambda}\|_2$. We further note that $0 \leq \beta \leq \beta_{max}$ in which β_{max} can be arbitrarily larger than C_λ . We can solve the optimization over $\boldsymbol{\lambda}$ to obtain

$$A_2 = \min_{\mathbf{e} \in A_e} \max_{0 \leq \beta \leq \beta_{max}} \beta \left\| \frac{1}{\sqrt{n}} \boldsymbol{\epsilon} - \frac{1}{\sqrt{nm}} \|\tilde{\mathbf{R}}^{1/2} \mathbf{e}\|_2 \mathbf{g} \right\|_2 - \frac{\beta}{\sqrt{nm}} \mathbf{h}^T \tilde{\mathbf{R}}^{1/2} \mathbf{e} - \frac{\beta^2}{2} + \frac{1}{m} B(\mathbf{e}) \quad (76)$$

We now note that the first term of this problem concentrates. We prove this in the following lemma

Lemma 3. Consider the term

$$F(\mathbf{e}, \beta) = \beta \left\| \frac{1}{\sqrt{n}} \boldsymbol{\epsilon} - \frac{1}{\sqrt{nm}} \|\tilde{\mathbf{R}}^{1/2} \mathbf{e}\|_2 \mathbf{g} \right\|_2 - \frac{\beta}{\sqrt{nm}} \mathbf{h}^T \tilde{\mathbf{R}}^{1/2} \mathbf{e} - \frac{\beta^2}{2} + \frac{1}{m} B(\mathbf{e}) \quad (77)$$

and let \bar{F} be given by

$$\bar{F}(\mathbf{e}, \beta) = \beta \sqrt{\sigma_\epsilon^2 - \frac{1}{m} \|\tilde{\mathbf{R}}^{1/2} \mathbf{e}\|_2^2} - \frac{\beta}{\sqrt{nm}} \mathbf{h}^T \tilde{\mathbf{R}}^{1/2} \mathbf{e} - \frac{\beta^2}{2} + \frac{1}{m} B(\mathbf{e}) \quad (78)$$

Then there exist positive constants C, c such that for any $\epsilon > 0$.

$$\mathbb{P} \left(\sup_{\mathbf{e} \in A_\epsilon, 0 \leq \beta \leq \beta_{max}} |F(\mathbf{e}, \beta) - \bar{F}(\mathbf{e}, \beta)| \geq \epsilon \right) \leq C e^{-c n \epsilon} \quad (79)$$

Proof. We see that F can be expressed as

$$F = \beta \sqrt{\frac{1}{n} \|\boldsymbol{\epsilon}\|_2^2 + \frac{1}{nm} \|\tilde{\mathbf{R}}^{1/2} \mathbf{e}\|_2^2 \|\mathbf{g}\|_2^2} - \frac{1}{n\sqrt{m}} \|\tilde{\mathbf{R}}^{1/2} \mathbf{e}\|_2 \boldsymbol{\epsilon}^T \mathbf{g} - \frac{\beta}{\sqrt{nm}} \mathbf{h}^T \tilde{\mathbf{R}}^{1/2} \mathbf{e} - \frac{\beta^2}{2} + \frac{1}{m} B(\mathbf{e}) \quad (80)$$

Or equivalently

$$F = \beta \sqrt{\left(\frac{1}{n} \|\boldsymbol{\epsilon}\|_2^2 - \sigma_\epsilon^2 \right) + \sigma_\epsilon^2 + \frac{1}{m} \|\tilde{\mathbf{R}}^{1/2} \mathbf{e}\|_2^2 \left(\frac{1}{n} \|\mathbf{g}\|_2^2 - 1 \right) + \frac{1}{m} \|\tilde{\mathbf{R}}^{1/2} \mathbf{e}\|_2 - \frac{2}{\sqrt{m}} \|\tilde{\mathbf{R}}^{1/2} \mathbf{e}\|_2 \frac{\boldsymbol{\epsilon}^T \mathbf{g}}{n} - \frac{\beta}{\sqrt{nm}} \mathbf{h}^T \tilde{\mathbf{R}}^{1/2} \mathbf{e} - \frac{\beta^2}{2} + \frac{1}{m} B(\mathbf{e})} \leq \bar{F} + \beta \sqrt{\delta} \leq \bar{F} + \beta_{max} \sqrt{\delta} \quad (81)$$

in which

$$\begin{aligned} \delta &= \left(\frac{1}{n} \|\boldsymbol{\epsilon}\|_2^2 - \sigma_\epsilon^2 \right) + \frac{1}{m} \|\tilde{\mathbf{R}}^{1/2} \mathbf{e}\|_2^2 \left(\frac{1}{n} \|\mathbf{g}\|_2^2 - 1 \right) - \frac{2}{\sqrt{m}} \|\tilde{\mathbf{R}}^{1/2} \mathbf{e}\|_2 \frac{\boldsymbol{\epsilon}^T \mathbf{g}}{n} \\ &\leq \left(\frac{1}{n} \|\boldsymbol{\epsilon}\|_2^2 - \sigma_\epsilon^2 \right) + C_\epsilon^2 C_{\tilde{\mathbf{R}}} \left(\frac{1}{n} \|\mathbf{g}\|_2^2 - 1 \right) + 2\sqrt{C_{\tilde{\mathbf{R}}} C_\epsilon} \left| \frac{\boldsymbol{\epsilon}^T \mathbf{g}}{n} \right| \stackrel{def}{=} \bar{\delta} \end{aligned} \quad (82)$$

We know that $C_{\tilde{\mathbf{R}}}$ and C_ϵ are universal constants. It is also clear that the probability that $\mathbb{P}(|\bar{\delta}| \geq \epsilon) \leq C e^{-c n \epsilon}$ for some constants $C, c > 0$. From this we can see that

$$\mathbb{P} \left(\sup_{\mathbf{e} \in A_\epsilon, 0 \leq \beta \leq \beta_{max}} |F(\mathbf{e}, \beta) - \bar{F}(\mathbf{e}, \beta)| \geq \epsilon \right) \leq \mathbb{P} \left(\sup_{\mathbf{e} \in A_\epsilon, 0 \leq \beta \leq \beta_{max}} |\delta \beta| \geq \epsilon \right) \leq \mathbb{P} (|\beta_{max} \bar{\delta}| \geq \epsilon) \leq C e^{-c n \epsilon} \quad (83)$$

For some constants $C, c > 0$.

□

Because of this we can with high probability, examine instead the problem

$$\bar{P}_2 = \min_{\mathbf{e} \in A_\epsilon} \max_{0 \leq \beta \leq \beta_{max}} \beta \sqrt{\sigma_\epsilon^2 + \frac{1}{m} \|\tilde{\mathbf{R}}^{1/2} \mathbf{e}\|_2^2} - \frac{\beta}{\sqrt{nm}} \mathbf{h}^T \tilde{\mathbf{R}}^{1/2} \mathbf{e} - \frac{\beta^2}{2} + \frac{1}{m} B(\mathbf{e}) \quad (84)$$

We now note that the optimization problem is convex in \mathbf{e} and concave in β , and both optimizations are over convex sets, as such we can interchange the order of min and max

$$\bar{P}_2 = \max_{0 \leq \beta \leq \beta_{max}} \min_{\mathbf{e} \in A_e} \beta \sqrt{\sigma_\epsilon^2 + \frac{1}{m} \|\tilde{\mathbf{R}}^{1/2} \mathbf{e}\|_2^2} - \frac{\beta}{\sqrt{nm}} \mathbf{h}^T \tilde{\mathbf{R}}^{1/2} \mathbf{e} - \frac{\beta^2}{2} + \frac{1}{m} B(\mathbf{e}) \quad (85)$$

We now note that for any scalar value a , we can express $\sqrt{a} = \min_{q>0} \frac{q}{2} + \frac{a}{2q}$. Making use of this "square root trick" we can obtain the problem

$$\bar{P}_2 = \max_{0 \leq \beta \leq \beta_{max}} \min_{\mathbf{e} \in A_e} \min_{q_{min} \leq q \leq q_{max}} \frac{\beta q}{2} + \frac{\beta}{2q} \sigma_\epsilon^2 + \frac{\beta}{2qm} \|\tilde{\mathbf{R}}^{1/2} \mathbf{e}\|_2^2 - \frac{\beta}{\sqrt{nm}} \mathbf{h}^T \tilde{\mathbf{R}}^{1/2} \mathbf{e} - \frac{\beta^2}{2} + \frac{1}{m} B(\mathbf{e}) \quad (86)$$

We note also that q can be bounded between $q_{min} = \sigma_\epsilon$, which is obtained when $\mathbf{e} = \mathbf{0}$ and $q_{max} = \sqrt{\sigma_\epsilon^2 + C_{\tilde{\mathbf{R}}} C_e^2}$. We can also swap the order of the two mins obtaining

$$\bar{P}_2 = \max_{0 \leq \beta \leq \beta_{max}} \min_{q_{min} \leq q \leq q_{max}} \min_{\mathbf{e} \in A_e} \frac{\beta q}{2} + \frac{\beta}{2q} \sigma_\epsilon^2 + \frac{\beta}{2qm} \|\tilde{\mathbf{R}}^{1/2} \mathbf{e}\|_2^2 - \frac{\beta}{\sqrt{nm}} \mathbf{h}^T \tilde{\mathbf{R}}^{1/2} \mathbf{e} - \frac{\beta^2}{2} + \frac{1}{m} B(\mathbf{e}) \quad (87)$$

At this point we will consider only the inner optimization problem over \mathbf{e} and consider β and q to be fixed. We shall return to the outer optimization later, and instead only consider

$$D_2 = D_2(\beta, q) = \min_{\mathbf{e} \in A_e} \frac{\beta q}{2} + \frac{\beta}{2q} \sigma_\epsilon^2 + \frac{\beta}{2qm} \|\tilde{\mathbf{R}}^{1/2} \mathbf{e}\|_2^2 - \frac{\beta}{\sqrt{nm}} \mathbf{h}^T \tilde{\mathbf{R}}^{1/2} \mathbf{e} - \frac{\beta^2}{2} + \frac{1}{m} B(\mathbf{e}) \quad (88)$$

We now make use of the definition of $\tilde{\mathbf{R}}$. We note specifically that

$$\tilde{\mathbf{R}}^{1/2} \mathbf{h} = \tilde{\mathbf{h}} \sim \mathcal{N}(0, \tilde{\mathbf{R}} = \frac{\rho_1^2}{d} \mathbf{W} \mathbf{W}^T + \rho_*^2 \mathbf{I}) \quad (89)$$

Which by the additivity of Gaussians can be expressed as

$$\tilde{\mathbf{h}} = \frac{\rho_1}{\sqrt{d}} \mathbf{W} \phi_1 + \rho_* \phi_2 \quad (90)$$

In which $\phi_1 \sim \mathcal{N}(0, I_d)$ and $\phi_2 \sim \mathcal{N}(0, I_m)$, we also pull the relevant factor of $\mathbf{e}^T \tilde{\mathbf{R}} \mathbf{e}$ out of $B(\mathbf{e})$. We make a new definition $\tilde{B}(\mathbf{e}) = r(\mathbf{e} + \boldsymbol{\theta}^*) + \tau_2 h(\mathbf{e} + \boldsymbol{\theta}^*)$, we remind that $\tilde{B}(\mathbf{e})$ is by assumption $\frac{\mu}{2}$ strongly convex. Making the relevant substitutions we obtain

$$D_2 = \min_{\mathbf{e} \in A_e} \frac{\beta q}{2} + \frac{\beta}{2q} \sigma_\epsilon^2 + \frac{\beta \rho_1^2}{2qmd} \|\mathbf{W}^T \mathbf{e}\|_2^2 + \frac{\beta \rho_*^2}{2qm} \|\mathbf{e}\|_2^2 - \frac{\beta \rho_1}{\sqrt{nmd}} \mathbf{e}^T \mathbf{W} \phi_1 - \frac{\beta \rho_*}{\sqrt{nm}} \phi_2^T \mathbf{e} - \frac{\beta^2}{2} + \frac{\tau_1 \rho_1^2}{md} \|\mathbf{W}^T \mathbf{e}\|_2^2 + \frac{\tau_1 \rho_*^2}{m} \|\mathbf{e}\|_2^2 + \frac{1}{m} \tilde{B}(\mathbf{e}) \quad (91)$$

We complete the square over the terms that contain $\mathbf{W}^T \mathbf{e}$, obtaining:

$$D_2 = \min_{\mathbf{e} \in A_e} \frac{\rho_1^2 (\beta + 2q\tau_1)}{2qmd} \left\| \mathbf{W}^T \mathbf{e} - \frac{\beta \sqrt{md}}{\rho_1 (\beta + 2qm\tau_1) \sqrt{n}} \phi_1 \right\|_2^2 - \frac{\beta^2 q}{2n(\beta + 2q\alpha_1)} \|\phi_1\|_2^2 + \frac{\beta q}{2} + \frac{\beta}{2q} \sigma_\epsilon^2 + \frac{\rho_*^2 (\beta + 2q\tau_1)}{2qm} \|\mathbf{e}\|_2^2 - \frac{\beta \rho_*}{\sqrt{nm}} \phi_2^T \mathbf{e} - \frac{\beta^2}{2} + \frac{1}{m} \tilde{B}(\mathbf{e}) \quad (92)$$

We now introduce another maximization over \mathbf{p} as the convex conjugate of the ℓ_2^2 norm. We obtain

$$D_2 = \min_{\mathbf{e} \in A_e} \max_{\mathbf{p}} \frac{\rho_1^2(\beta + 2q\tau_1)}{qmd} \mathbf{p}^T \mathbf{W}^T \mathbf{e} - \frac{\beta\rho_1}{\sqrt{nm}d} \mathbf{p}^T \phi_1 - \frac{\rho_1^2(\beta + 2q\tau_1)}{2qmd} \|\mathbf{p}\|_2^2 - \frac{\beta^2 q}{2n(\beta + 2q\tau_1)} \|\phi_1\|_2^2 + \frac{\beta q}{2} + \frac{\beta}{2q} \sigma_\epsilon^2 + \frac{\rho_*^2(\beta + 2q\tau_1)}{2qm} \|\mathbf{e}\|_2^2 - \frac{\beta\rho_*}{\sqrt{nm}} \phi_2^T \mathbf{e} - \frac{\beta^2}{2} + \frac{1}{m} \tilde{B}(\mathbf{e}) \quad (93)$$

Our goal now is to apply the CGMT again to this problem. We note that the problem is convex in \mathbf{e} and concave in \mathbf{p} . However we need to show that \mathbf{e} and \mathbf{p} can be bound to compact and convex sets, and that the optimal points of both optimizations fall within these sets. We prove this in the following lemma

Lemma 4. *Consider the following two optimization problems that correspond to the first alternative and second alternative optimization by the CGMT*

$$D_2 = \min_{\mathbf{e} \in A_e} \max_{\mathbf{p}} \frac{\rho_1^2(\beta + 2q\tau_1)}{qmd} \mathbf{p}^T \mathbf{W}^T \mathbf{e} - \frac{\beta\rho_1}{\sqrt{nm}d} \mathbf{p}^T \phi_1 - \frac{\rho_1^2(\beta + 2q\tau_1)}{2qmd} \|\mathbf{p}\|_2^2 - \frac{\beta^2 q}{2n(\beta + 2q\tau_1)} \|\phi_1\|_2^2 + \frac{\beta q}{2} + \frac{\beta}{2q} \sigma_\epsilon^2 + \frac{\rho_*^2(\beta + 2q\tau_1)}{2qm} \|\mathbf{e}\|_2^2 - \frac{\beta\rho_*}{\sqrt{nm}} \phi_2^T \mathbf{e} - \frac{\beta^2}{2} + \frac{1}{m} \tilde{B}(\mathbf{e}) \quad (94)$$

$$D_3 = \min_{\mathbf{e}} \max_{\mathbf{p}} \frac{\rho_1^2(\beta + 2q\tau_1)}{qmd} \|\mathbf{p}\|_2 \mathbf{e}^T \phi_3 + \frac{\rho_1^2(\beta + 2q\tau_1)}{qmd} \|\mathbf{e}\|_2 \mathbf{p}^T \phi_4 - \frac{\beta\rho_1}{\sqrt{nm}d} \mathbf{p}^T \phi_1 - \frac{\rho_1^2(\beta + 2q\tau_1)}{2qmd} \|\mathbf{p}\|_2^2 - \frac{\beta^2 q}{2n(\beta + 2q\tau_1)} \|\phi_1\|_2^2 + \frac{\beta q}{2} + \frac{\beta}{2q} \sigma_\epsilon^2 + \frac{\rho_*^2(\beta + 2q\tau_1)}{2qm} \|\mathbf{e}\|_2^2 - \frac{\beta\rho_*}{\sqrt{nm}} \phi_2^T \mathbf{e} - \frac{\beta^2}{2} + \frac{1}{m} \tilde{B}(\mathbf{e}) \quad (95)$$

where ϕ_3 and ϕ_4 are standard normals of dimension m, d respectively. Denote $\hat{\mathbf{e}}_2, \hat{\mathbf{e}}_3$ as optimal points of D_2 and D_3 respectively and $\hat{\mathbf{p}}_2(\mathbf{e}), \hat{\mathbf{p}}_3(\mathbf{e})$ as their inner optimization solution for a fixed \mathbf{e} . Let \tilde{B} be $\frac{\mu}{2}$ strongly convex and $\max\{\|\nabla r(\theta^*)\|, \|\nabla h(\theta^*)\|\} = O(\sqrt{m})$. Then there exist positive constants C_e, C_p only depending on μ such that

$$\lim_{m \rightarrow \infty} \mathbb{P}(\|\hat{\mathbf{e}}_i\|_2 \leq C_e \sqrt{m}) = 1 \quad i = 2, 3 \quad (96)$$

and

$$\lim_{m \rightarrow \infty} \mathbb{P}\left(\sup_{\mathbf{e} \|\mathbf{e}\| \leq C_e} \|\hat{\mathbf{p}}_i(\mathbf{e})\|_2 \leq C_p \sqrt{md}\right) = 1 \quad i = 2, 3 \quad (97)$$

Proof. We know that C_{e_2} exists from the fact that in D_2 \mathbf{e} is already in a bounded set. For both optimizations, we solve the optimization over \mathbf{p} , and write this optimization over \mathbf{e} as

$$\min_{\mathbf{e}} F_i(\mathbf{e}) \quad i = 2, 3, \quad (98)$$

where $F_i(\mathbf{e})$ is the optimal value over \mathbf{p} . We note that setting $\mathbf{p} = \mathbf{0}$ in both optimizations we obtain that

$$F(\mathbf{e}) \geq \frac{1}{m} T(\mathbf{e}) := -\frac{\beta^2 q}{2n(\beta + 2q\tau_1)} \|\phi_1\|_2^2 + \frac{\beta q}{2} + \frac{\beta}{2q} \sigma_\epsilon^2 + \frac{\rho_*^2(\beta + 2q\tau_1)}{2qm} \|\mathbf{e}\|_2^2 - \frac{\beta\rho_*}{\sqrt{nm}} \phi_2^T \mathbf{e} - \frac{\beta^2}{2} + \frac{1}{m} \tilde{B}(\mathbf{e}) \quad (99)$$

On the other hand, by taking the second derivative, we observe that $T(\mathbf{e})$ is $\nu = \frac{\rho_* \beta}{2q} + \frac{\mu}{2}$ strongly convex with respect to \mathbf{e} . As such, we find that

$$T(\mathbf{e}) \geq T(\mathbf{0}) + \mathbf{d}^T \mathbf{e} + \frac{\nu}{2} \|\mathbf{e}\|_2^2, \quad (100)$$

where $\mathbf{d} = \nabla T(\mathbf{0})$. We note that by the assumption, $\mathbf{d} = O(\sqrt{m})$. For the optimization D_3 , we let $\xi = \|\mathbf{p}\|_2$ and solve the optimization over \mathbf{p} to obtain that

$$F_3(\mathbf{e}) = \max_{\xi > 0} \frac{\rho_1^2(\beta + 2q\tau_1)\xi}{qmd} \mathbf{e}^T \phi_3 + \xi \left\| \frac{\rho_1^2(\beta + 2q\tau_1)}{qmd} \|\mathbf{e}\|_2 \phi_4 - \frac{\beta\rho_1}{\sqrt{nm}d} \phi_1 \right\|_2 - \frac{\rho_1^2(\beta + 2q\tau_1)\xi^2}{2qmd} + \frac{1}{m} T(\mathbf{e}) \quad (101)$$

We note that dropping the constraint over ξ will not decrease the optimal value, as such

$$F_3(\mathbf{e}) \leq \max_{\xi} \frac{\rho_1^2(\beta + 2q\tau_1)\xi}{qmd} \mathbf{e}^T \phi_3 + \xi \left\| \frac{\rho_1^2(\beta + 2q\tau_1)}{qmd} \|\mathbf{e}\|_2 \phi_4 - \frac{\beta\rho_1}{\sqrt{nm}d} \phi_1 \right\|_2 - \frac{\rho_1^2(\beta + 2q\tau_1)\xi^2}{2qmd} + \frac{1}{m} T(\mathbf{e}) \quad (102)$$

From which we see that

$$\begin{aligned} F_3(\mathbf{0}) &\leq \max_{\xi} \frac{\xi\beta\rho_1}{\sqrt{nm}d} \|\phi_1\|_2 - \frac{\rho_1^2(\beta + 2q\tau_1)\xi^2}{2qmd} + \frac{1}{m}T(\mathbf{0}) \\ &= \frac{\beta\sqrt{md}q\|\phi_1\|_2}{2\sqrt{n}\rho_1(\beta + 2q\tau_1)^2} + \frac{1}{m}T(\mathbf{0}) \end{aligned} \quad (103)$$

From this we obtain that

$$\frac{\beta\sqrt{md}q\|\phi_1\|_2}{2\sqrt{n}\rho_1(\beta + 2q\tau_1)^2} + \frac{1}{m}T(\mathbf{0}) \geq F(\mathbf{0}) \geq F_3(\hat{\mathbf{e}}) \geq \frac{1}{m}T(\mathbf{0}) + \frac{1}{m}\mathbf{d}^T\mathbf{e} + \frac{\nu}{2m}\|\mathbf{e}\|_2^2, \quad (104)$$

and hence

$$\frac{\nu}{2m} \left\| \mathbf{e} + \frac{1}{\nu}\mathbf{d} \right\|_2^2 \leq \frac{1}{\nu m} \|\mathbf{d}\|_2^2 + \frac{\beta\sqrt{md}q\|\phi_1\|_2}{2\sqrt{n}\rho_1(\beta + 2q\tau_1)^2} \quad (105)$$

or

$$\|\mathbf{e}\|_2 \leq \frac{1}{\nu} \|\mathbf{d}\|_2 + \sqrt{\frac{2}{\nu^2} \|\mathbf{d}\|_2^2 + \frac{m\beta\sqrt{md}q\|\phi_1\|_2}{\nu\sqrt{n}\rho_1(\beta + 2q\tau_1)^2}} \quad (106)$$

Noting that with high probability $\|\phi\|_2 \leq C\sqrt{d}$ and recalling that n, m, d all grow at constant ratios, we can see that there must exist a constant C_{e_3} such that

$$\mathbb{P}(\|\hat{\mathbf{e}}_3\| > C_{e_3}\sqrt{m}) \rightarrow 0 \quad (107)$$

We then let $C_e = \max(C_{e_2}, C_{e_3})$ and define the set $\tilde{A}_e = \{\mathbf{e} \in \mathbb{R}^m \mid \|\mathbf{e}\|_2 \leq C_e\sqrt{m}\}$. Then from the optimality condition over \mathbf{p} for eq (94) we know that

$$\hat{\mathbf{p}}_2(\mathbf{e}) = \mathbf{W}^T\mathbf{e} - \frac{\beta\sqrt{md}}{\rho_1(\beta + 2qm\tau_1)\sqrt{n}}\phi_1 \quad (108)$$

and as such for all $\mathbf{e} \in A_e$ we must have that

$$\|\hat{\mathbf{p}}_2(\mathbf{e})\|_2 \leq \|\mathbf{W}\|_2\|\mathbf{e}\|_2 + \frac{\beta\sqrt{md}}{\rho_1(\beta + 2qm\tau_1)\sqrt{n}}\|\phi_1\|_2 \quad (109)$$

We know as a standard result that $\|\mathbf{W}\|_2 \leq C\sqrt{d}$ and that $\|\phi_1\|_2 \leq C\sqrt{d}$ with high probability. As such the constant C_{p_2} must exist.

Finally examining the optimality condition over ξ of problem (95) we find that for all $\mathbf{e} \in A_e$ we have that

$$\begin{aligned} \hat{\xi} = \|\hat{\mathbf{p}}_3(\mathbf{e})\|_2 &= \mathbf{e}^T\phi_3 + \left\| \|\mathbf{e}\|_2\phi_4 - \frac{q\beta\sqrt{md}}{\rho_1(\beta + 2q\tau_1)\sqrt{n}}\|\phi_1\|_2 \right\|_2 \\ &\leq \|\mathbf{e}\|_2\|\phi_3\|_2 + \|\mathbf{e}\|_2\|\phi_4\|_2 + \frac{q\beta\sqrt{md}}{\rho_1(\beta + 2q\tau_1)\sqrt{n}}\|\phi_1\|_2 \end{aligned} \quad (110)$$

We note that with high probability $\|\phi_1\|_2 < \sqrt{d}C$, $\|\phi_4\|_2 < \sqrt{d}C$ and $\|\phi_3\|_2 < \sqrt{m}C$. Recalling that m, d grow at constant ratio we see that the constant C_{p_3} exists. \square

We can therefore define the constants $C_e := \max(C_{e_i})$ from $i = 1, 2, 3$ and $C_p = \max(C_{p_2}, C_{p_3})$, and by doing so define the sets $A_e = \{\mathbf{e} \in \mathbb{R}^m \mid \|\mathbf{e}\|_2 < C_e\sqrt{m}\}$ and $A_p = \{\mathbf{p} \in \mathbb{R}^d \mid \|\mathbf{p}\|_2 < C_p\sqrt{md}\}$. From this we can see that with high probability the optimal value of the optimization \hat{P}_2 will be equal to that of

$$\begin{aligned}
 D_2 = & \min_{\mathbf{e} \in A_{\mathbf{e}}} \max_{\mathbf{p} \in A_{\mathbf{p}}} \frac{\rho_1^2(\beta + 2q\tau_1)}{qmd} \mathbf{p}^T \mathbf{W}^T \mathbf{e} - \frac{\beta\rho_1}{\sqrt{nm}d} \mathbf{p}^T \phi_1 - \frac{\rho_1^2(\beta + 2q\tau_1)}{2qmd} \|\mathbf{p}\|_2^2 \\
 & - \frac{\beta^2q}{2n(\beta + 2q\tau_1)} \|\phi_1\|_2^2 + \frac{\beta q}{2} + \frac{\beta}{2q} \sigma_\epsilon^2 + \frac{\rho_*^2(\beta + 2q\tau_1)}{2qm} \|\mathbf{e}\|_2^2 - \frac{\beta\rho_*}{\sqrt{nm}} \phi_2^T \mathbf{e} - \frac{\beta^2}{2} + \frac{1}{m} \tilde{B}(\mathbf{e})
 \end{aligned} \quad (111)$$

We now apply the CGMT to the problem D_2 for fixed values of β, q , we obtain the following problem

$$\begin{aligned}
 D_3 = & \min_{\mathbf{e} \in A_{\mathbf{e}}} \max_{\mathbf{p} \in A_{\mathbf{p}}} \frac{\rho_1^2(\beta + 2q\tau_1)}{qmd} \|\mathbf{p}\|_2 \mathbf{e}^T \phi_3 + \frac{\rho_1^2(\beta + 2q\tau_1)}{qmd} \|\mathbf{e}\|_2 \mathbf{p}^T \phi_4 - \frac{\beta\rho_1}{\sqrt{nm}d} \mathbf{p}^T \phi_1 - \frac{\rho_1^2(\beta + 2q\tau_1)}{2qmd} \|\mathbf{p}\|_2^2 \\
 & - \frac{\beta^2q}{2n(\beta + 2q\tau_1)} \|\phi_1\|_2^2 + \frac{\beta q}{2} + \frac{\beta}{2q} \sigma_\epsilon^2 + \frac{\rho_*^2(\beta + 2q\tau_1)}{2qm} \|\mathbf{e}\|_2^2 - \frac{\beta\rho_*}{\sqrt{nm}} \phi_2^T \mathbf{e} - \frac{\beta^2}{2} + \frac{1}{m} \tilde{B}(\mathbf{e})
 \end{aligned} \quad (112)$$

Let $\xi = \frac{\rho_1}{\sqrt{md}} \|\mathbf{p}\|_2$ and solve the optimization over \mathbf{p} . We note that $\xi \geq 0$ and that $\xi \leq \xi_{max} = \frac{\rho_1}{\sqrt{dm}} \sup_{\mathbf{p} \in A_{\mathbf{p}}} \|\mathbf{p}\|_2$. From this we obtain the problem,

$$\begin{aligned}
 D_3 = & \min_{\mathbf{e} \in A_{\mathbf{e}}} \max_{0 \leq \xi \leq \xi_{max}} \frac{\rho_1(\beta + 2q\tau_1)}{q\sqrt{md}} \mathbf{e}^T \phi_3 + \xi \left\| \frac{\rho_1(\beta + 2q\tau_1)}{q\sqrt{md}} \|\mathbf{e}\|_2 \phi_4 - \frac{\beta}{\sqrt{n}} \phi_1 \right\|_2 - \frac{(\beta + 2q\tau_1)\xi^2}{2q} \\
 & - \frac{\beta^2q}{2n(\beta + 2q\tau_1)} \|\phi_1\|_2^2 + \frac{\beta q}{2} + \frac{\beta}{2q} \sigma_\epsilon^2 + \frac{\rho_*^2(\beta + 2q\tau_1)}{2qm} \|\mathbf{e}\|_2^2 - \frac{\beta\rho_*}{\sqrt{nm}} \phi_2^T \mathbf{e} - \frac{\beta^2}{2} + \frac{1}{m} \tilde{B}(\mathbf{e})
 \end{aligned} \quad (113)$$

We now show that this term concentrates in the following lemma

Lemma 5. Let $F(\mathbf{e}, \xi)$ be given by

$$\begin{aligned}
 F(\mathbf{e}, \xi) = & \frac{\rho_1(\beta + 2q\tau_1)}{q\sqrt{md}} \mathbf{e}^T \phi_3 + \xi \left\| \frac{\rho_1(\beta + 2q\tau_1)}{q\sqrt{md}} \|\mathbf{e}\|_2 \phi_4 - \frac{\beta}{\sqrt{n}} \phi_1 \right\|_2 - \frac{(\beta + 2q\tau_1)\xi^2}{2q} \\
 & - \frac{\beta^2q}{2n(\beta + 2q\tau_1)} \|\phi_1\|_2^2 + \frac{\beta q}{2} + \frac{\beta}{2q} \sigma_\epsilon^2 + \frac{\rho_*^2(\beta + 2q\tau_1)}{2qm} \|\mathbf{e}\|_2^2 - \frac{\beta\rho_*}{\sqrt{nm}} \phi_2^T \mathbf{e} - \frac{\beta^2}{2} + \frac{1}{m} \tilde{B}(\mathbf{e})
 \end{aligned} \quad (114)$$

and let $\bar{F}(\mathbf{e}, \xi)$

$$\begin{aligned}
 \bar{F}(\mathbf{e}, \xi) = & \frac{\rho_1(\beta + 2q\tau_1)}{q\sqrt{md}} \mathbf{e}^T \phi_3 + \xi \sqrt{\frac{\rho_1^2(\beta + 2q\tau_1)^2}{q^2m} \|\mathbf{e}\|_2^2 + \frac{\beta^2d}{n}} - \frac{(\beta + 2q\tau_1)\xi^2}{2q} \\
 & - \frac{\beta^2qd}{2n(\beta + 2q\tau_1)} + \frac{\beta q}{2} + \frac{\beta}{2q} \sigma_\epsilon^2 + \frac{\rho_*^2(\beta + 2q\tau_1)}{2qm} \|\mathbf{e}\|_2^2 - \frac{\beta\rho_*}{\sqrt{nm}} \phi_2^T \mathbf{e} - \frac{\beta^2}{2} + \frac{1}{m} \tilde{B}(\mathbf{e})
 \end{aligned} \quad (115)$$

Then

$$\mathbb{P} \left(\sup_{\mathbf{e} \in A_{\mathbf{e}}, 0 \leq \xi \leq \xi_{max}} |F(\mathbf{e}, \xi) - \bar{F}(\mathbf{e}, \xi)| > \epsilon \right) \xrightarrow[m, d \rightarrow \infty]{P} 0 \quad (116)$$

Proof. The lemma is proven in the same manner as lemma 3. \square

By this lemma we can with high probability consider the following problem instead:

$$\begin{aligned}
 \bar{D}_3 = & \min_{\mathbf{e} \in A_{\mathbf{e}}} \max_{0 \leq \xi \leq \xi_{max}} \frac{\rho_1(\beta + 2q\tau_1)}{q\sqrt{md}} \mathbf{e}^T \phi_3 + \xi \sqrt{\frac{\rho_1^2(\beta + 2q\tau_1)^2}{q^2m} \|\mathbf{e}\|_2^2 + \frac{\beta^2d}{n}} - \frac{(\beta + 2q\tau_1)\xi^2}{2q} \\
 & - \frac{\beta^2qd}{2n(\beta + 2q\tau_1)} + \frac{\beta q}{2} + \frac{\beta}{2q} \sigma_\epsilon^2 + \frac{\rho_*^2(\beta + 2q\tau_1)}{2qm} \|\mathbf{e}\|_2^2 - \frac{\beta\rho_*}{\sqrt{nm}} \phi_2^T \mathbf{e} - \frac{\beta^2}{2} + \frac{1}{m} \tilde{B}(\mathbf{e})
 \end{aligned} \quad (117)$$

We now interchange the order of the min and max. As the problem is clearly convex in \mathbf{e} and concave in ξ and the problem is over convex sets this interchange is admissible.

$$\begin{aligned} \bar{D}_3 = & \max_{0 \leq \xi \leq \xi_{max}} \min_{\mathbf{e} \in A_{\mathbf{e}}} \frac{\rho_1(\beta + 2q\tau_1)}{q\sqrt{md}} \mathbf{e}^T \phi_3 + \xi \sqrt{\frac{\rho_1^2(\beta + 2q\tau_1)^2}{q^2 m} \|\mathbf{e}\|_2^2 + \frac{\beta^2 d}{n} - \frac{(\beta + 2q\tau_1)\xi^2}{2q}} \\ & - \frac{\beta^2 qd}{2n(\beta + 2q\tau_1)} + \frac{\beta q}{2} + \frac{\beta}{2q} \sigma_{\epsilon}^2 + \frac{\rho_*^2(\beta + 2q\tau_1)}{2qm} \|\mathbf{e}\|_2^2 - \frac{\beta \rho_*}{\sqrt{nm}} \phi_2^T \mathbf{e} - \frac{\beta^2}{2} + \frac{1}{m} \tilde{B}(\mathbf{e}) \end{aligned} \quad (118)$$

We now make use of the square root trick one more time, introducing new parameter t , we note that t can be bounded by $t_{min} = \frac{\beta\sqrt{d}}{\sqrt{n}}$ and $t_{max} = \sqrt{\frac{\beta^2 d}{n} + \frac{\rho_1^2(\beta + 2q\tau_1)^2}{q^2} C_{\mathbf{e}}^2}$.

$$\begin{aligned} \bar{D}_3 = & \max_{0 \leq \xi \leq \xi_{max}} \min_{t_{min} \leq t \leq t_{max}} \min_{\mathbf{e} \in A_{\mathbf{e}}} \frac{\rho_1(\beta + 2q\tau_1)}{q\sqrt{md}} \mathbf{e}^T \phi_3 + \frac{\xi \rho_1^2(\beta + 2q\tau_1)^2}{2tq^2 m} \|\mathbf{e}\|_2^2 + \frac{\beta^2 \xi d}{2tn} - \frac{(\beta + 2q\tau_1)\xi^2}{2q} \\ & - \frac{\beta^2 qd}{2n(\beta + 2q\tau_1)} + \frac{\xi t}{2} + \frac{\beta q}{2} + \frac{\beta}{2q} \sigma_{\epsilon}^2 + \frac{\rho_*^2(\beta + 2q\tau_1)}{2qm} \|\mathbf{e}\|_2^2 - \frac{\beta \rho_*}{\sqrt{nm}} \phi_2^T \mathbf{e} - \frac{\beta^2}{2} + \frac{1}{m} \tilde{B}(\mathbf{e}) \end{aligned} \quad (119)$$

Where we have changed the order of the two min operations. We can now define the constants,

$$c_1 = \frac{\xi \rho_1^2(\beta + 2q\tau_1)^2}{2tq^2} + \frac{\rho_*^2(\beta + 2q\tau_1)}{2q} \quad c_2 = \sqrt{\frac{\rho_1^2(\beta + 2q\tau_1)^2 \eta}{q^2} + \rho_*^2 \beta^2} \quad (120)$$

and we note that by the additivity of Gaussians we have that

$$\frac{c_2}{\sqrt{nm}} \phi = \frac{\rho_1(\beta + 2q\tau_1)}{q\sqrt{md}} \phi_3 - \frac{\beta \rho_*}{\sqrt{nm}} \phi_2 \quad (121)$$

We obtain

$$\begin{aligned} \bar{D}_3 = & \max_{0 \leq \xi \leq \xi_{max}} \min_{t_{min} \leq t \leq t_{max}} \min_{\mathbf{e} \in A_{\mathbf{e}}} \frac{c_1}{m} \|\mathbf{e}\|_2^2 + \frac{c_2}{\sqrt{nm}} \phi^T \mathbf{e} + \frac{\beta^2 \xi d}{2tn} - \frac{(\beta + 2q\tau_1)\xi^2}{2q} \\ & - \frac{\beta^2 qd}{2n(\beta + 2q\tau_1)} + \frac{\xi t}{2} + \frac{\beta q}{2} + \frac{\beta}{2q} \sigma_{\epsilon}^2 - \frac{\beta^2}{2} + \frac{1}{m} \tilde{B}(\mathbf{e}) \end{aligned} \quad (122)$$

Completing the square over \mathbf{e} we find

$$\begin{aligned} \bar{D}_3 = & \max_{0 \leq \xi \leq \xi_{max}} \min_{t_{min} \leq t \leq t_{max}} \min_{\mathbf{e} \in A_{\mathbf{e}}} \frac{c_1}{m} \left\| \mathbf{e} + \frac{c_2 \sqrt{m}}{2c_1 \sqrt{n}} \phi \right\|_2^2 - \frac{c_2^2}{4c_1 n} \|\phi\|_2^2 + \frac{\beta^2 \xi d}{2tn} - \frac{(\beta + 2q\tau_1)\xi^2}{2q} \\ & - \frac{\beta^2 qd}{2n(\beta + 2q\tau_1)} + \frac{\xi t}{2} + \frac{\beta q}{2} + \frac{\beta}{2q} \sigma_{\epsilon}^2 - \frac{\beta^2}{2} + \frac{1}{m} \tilde{B}(\mathbf{e}) \end{aligned} \quad (123)$$

Finally noting that in the asymptotic limit $\|\phi\|_2^2$ concentrates to m with high probability, and then recognizing the Moreau envelope over \mathbf{e} (see definition 1 below) we obtain the problem

$$\begin{aligned} \bar{D}_3 = & \max_{0 \leq \xi \leq \xi_{max}} \min_{t_{min} \leq t \leq t_{max}} \frac{1}{m} \mathcal{M}_{\frac{1}{2c_1} \tilde{B}} \left(-\frac{c_2^2 \sqrt{m}}{2c_1 \sqrt{n}} \phi \right) - \frac{c_2 m}{4c_1 n} + \frac{\beta^2 \xi d}{2tn} - \frac{(\beta + 2q\tau_1)\xi^2}{2q} \\ & - \frac{\beta^2 qd}{2n(\beta + 2q\tau_1)} + \frac{\xi t}{2} + \frac{\beta q}{2} + \frac{\beta}{2q} \sigma_{\epsilon}^2 - \frac{\beta^2}{2} \end{aligned} \quad (124)$$

We can recall that $\tilde{B}(\mathbf{e}) = r(\mathbf{e} + \boldsymbol{\theta}^*) + \tau_2 h(\mathbf{e} + \boldsymbol{\theta}^*)$, and letting $\boldsymbol{\theta} = \mathbf{e} + \boldsymbol{\theta}^*$, we obtain

$$\begin{aligned} \bar{D}_3 = & \max_{0 \leq \xi \leq \xi_{max}} \min_{t_{min} \leq t \leq t_{max}} \frac{1}{m} \mathcal{M}_{\frac{1}{2c_1}(r + \tau_2 h)} \left(\boldsymbol{\theta}^* - \frac{c_2^2 \sqrt{m}}{2c_1 \sqrt{n}} \phi \right) - \frac{c_2 m}{4c_1 n} + \frac{\beta^2 \xi d}{2tn} - \frac{(\beta + 2q\tau_1)\xi^2}{2q} \\ & - \frac{\beta^2 qd}{2n(\beta + 2q\tau_1)} + \frac{\xi t}{2} + \frac{\beta q}{2} + \frac{\beta}{2q} \sigma_{\epsilon}^2 - \frac{\beta^2}{2} \end{aligned} \quad (125)$$

Finally we show in Lemma 8 that the Moreau envelope will concentrate in the asymptotic limit on its expected value. As such we finally obtain:

$$\begin{aligned} \bar{D}_3 = \max_{0 \leq \xi \leq \xi_{max}} \min_{t_{min} \leq t \leq t_{max}} \frac{1}{m} \mathbb{E} \mathcal{M}_{\frac{1}{2c_1}(r+\tau_2 h)} \left(\boldsymbol{\theta}^* - \frac{c_2 \sqrt{m}}{2c_1 \sqrt{n}} \boldsymbol{\phi} \right) - \frac{c_2 m}{4c_1 n} + \frac{\beta^2 \xi d}{2tn} - \frac{(\beta + 2q\tau_1) \xi^2}{2q} \\ - \frac{\beta^2 q d}{2n(\beta + 2q\tau_1)} + \frac{\xi t}{2} + \frac{\beta q}{2} + \frac{\beta}{2q} \sigma_\epsilon^2 - \frac{\beta^2}{2} \end{aligned} \quad (126)$$

We know by the properties of the CGMT that for any fixed choice of β, q that $D_3(\beta, q)$ converges pointwise to $D_2(\beta, q)$. However to determine the properties that we are interested in we require uniform convergence. For this, we simply show that $D_2(\beta, q), \bar{D}_3(\beta, q)$ are Lipschitz continuous for $\beta \in [0, \beta_{max}]$ and $q \in [q_{min}, q_{max}]$.

Lemma 6. *The problem D_2 as given in (88) and problem \bar{D}_3 as given in equations 117, (119) and (124) are C -Lipschitz on the compact set $K = [0, \beta_{max}] \times [q_{min}, q_{max}]$ for some constant $C < \infty$, with high probability.*

Proof. We first consider problem D_2 given in equation (88).

$$D_2 = \min_{A_e} \frac{\beta q}{2} + \frac{\beta \sigma_\epsilon^2}{2q} + \frac{\beta}{2qm} \|\tilde{\mathbf{R}}^{1/2} \mathbf{e}\|_2^2 - \frac{\beta}{\sqrt{nm}} \mathbf{h}^T \tilde{\mathbf{R}} \mathbf{e} - \frac{\beta^2}{2} + \frac{\tau_1}{m} \mathbf{e} \tilde{\mathbf{R}} \mathbf{e} + \frac{1}{m} \tilde{B}(\mathbf{e}) \quad (127)$$

We note that the objective $D(\beta, q, \mathbf{e})$ is strongly convex, the solution is hence unique, and D_2 is continuously differentiable on the compact set K . We simply bound its gradient, which is given by

$$\frac{\partial D_2}{\partial \beta} = \frac{\partial D}{\partial \beta} \Big|_{\mathbf{e}=\hat{\mathbf{e}}} = \frac{q}{2} + \frac{\sigma_\epsilon^2}{2q} + \frac{1}{2qm} \|\tilde{\mathbf{R}}^{1/2} \hat{\mathbf{e}}\|_2^2 - \frac{1}{\sqrt{nm}} \mathbf{h}^T \tilde{\mathbf{R}} \hat{\mathbf{e}} - \beta \quad (128)$$

$$\frac{\partial D_2}{\partial q} = \frac{\partial D}{\partial q} \Big|_{\mathbf{e}=\hat{\mathbf{e}}} = \frac{\beta}{2} - \frac{\beta \sigma_\epsilon^2}{2q^2} - \frac{\beta}{2q^2 m} \|\tilde{\mathbf{R}}^{1/2} \hat{\mathbf{e}}\|_2^2 \quad (129)$$

where $\hat{\mathbf{e}}$ is the optimal solution. Noting that $\hat{\mathbf{e}} \in A_e$ and β, q are bounded, we obtain the result for D_2 .

For problem \bar{D}_3 we make use of the same strategy by calculating the gradient. Defining $\hat{\mathbf{e}}, \hat{\xi}$ as the optimal solution of (117), we observe that

$$\hat{\xi} = \sqrt{\frac{\rho_1^2}{4m} \|\hat{\mathbf{e}}\|_2^2 + \frac{\beta^2 q^2 d}{4(\beta + 2q\tau_1)^2 n}} \quad (130)$$

Further, we define

$$\hat{t} = \sqrt{\frac{\rho_1^2 (\beta + 2q\tau_1)^2}{q^2 m} \|\hat{\mathbf{e}}\|_2^2 + \frac{\beta^2 d}{n}} \quad (131)$$

Finally we examine the partial derivatives of problem D_3 with respect to β and q ,

$$\begin{aligned} \frac{\partial \bar{D}_3}{\partial \beta} = \frac{\rho_1}{q\sqrt{md}} \mathbf{e}^T \boldsymbol{\phi}_3 + \frac{\xi \rho_1^2 (\beta + 2q\tau_1)}{t q^2 m} \|\mathbf{e}\|_2^2 + \frac{\beta \xi d}{tn} - \frac{\xi^2}{2q} - \frac{\beta q d}{n(\beta + 2q\tau_1)} + \frac{\beta^2 q d}{2n(\beta + 2q\tau_1)^2} \\ + \frac{q}{2} + \frac{\sigma_\epsilon^2}{2q} + \frac{\rho_*^2}{2qm} \|\mathbf{e}\|_2^2 - \frac{\beta \rho_*}{\sqrt{nm}} \boldsymbol{\phi}_2^T \mathbf{e} - \beta \end{aligned} \quad (132)$$

$$\begin{aligned} \frac{\partial \bar{D}_3}{\partial q} = -\frac{\rho_1 (\beta + 2q\tau_1)}{q^2 \sqrt{md}} \mathbf{e}^T \boldsymbol{\phi}_3 - \frac{\xi \rho_1^2 (\beta + 2q\tau_1)^2}{t q^3 m} \|\mathbf{e}\|_2^2 + \frac{\xi \rho_1^2 \tau_1 (\beta + 2q\tau_1)}{2t q^2 m} \|\mathbf{e}\|_2^2 + \frac{(\beta + 2q\tau_1) \xi^2}{2q^2} - \frac{2\tau_1 \xi^2}{2q} \\ - \frac{\beta^2 d}{2n(\beta + 2q\tau_1)} + \frac{\beta^2 q d \tau_1}{n(\beta + 2q\tau_1)^2} + \frac{\beta}{2} - \frac{\beta}{2q^2} \sigma_\epsilon^2 - \frac{\rho_*^2 (\beta + 2q\tau_1)}{2q^2 m} \|\mathbf{e}\|_2^2 + \frac{\rho_*^2 \tau_1}{qm} \|\mathbf{e}\|_2^2 \end{aligned} \quad (133)$$

Noting the boundedness of the involved terms, we conclude the result. \square

We have established that both D_2 and \bar{D}_3 are Lipschitz, we now create a rectangular ϵ net \mathcal{N} on the set $[0, \beta_{max}] \times [q_{min}, q_{max}]$ consisting of $k = \frac{\beta_{max}(q_{max}-q_{min})}{\epsilon^2}$ points. We can then see that

$$|D_2(\beta, q) - D_3(\beta, q)| \leq |D_2(\beta, q) - D_2(\beta_k, q_k)| + |D_2(\beta_k, q_k) - D_3(\beta_k, q_k)| + |D_3(\beta_k, q_k) - D_3(\beta, q)| \leq C\epsilon\sqrt{2} + |D_2(\beta_k, q_k) - D_3(\beta_k, q_k)| + C\epsilon\sqrt{2}, \quad (134)$$

β_k, q_k is the closes element of the ϵ -net to β, q . The second inequality is due to the fact that both D_2 and D_3 are C -Lipschitz with respect to both β and q and the distance of between β, q and β_k, q_k cannot be more than $\epsilon\sqrt{2}$. From this we can see that

$$\sup_{0 \leq \beta \leq \beta_{max}, q_{min} \leq q \leq q_{max}} |D_2(\beta, q) - D_3(\beta, q)| \leq 2C\epsilon\sqrt{2} + \sup_{\beta, q \in \mathcal{N}} |D_2(\beta, q) - D_3(\beta, q)| \quad (135)$$

As a result,

$$\mathbb{P} \left(\sup_{0 \leq \beta \leq \beta_{max}, q_{min} \leq q \leq q_{max}} |D_2(\beta, q) - D_3(\beta, q)| \geq 4C\epsilon\sqrt{2} \right) \leq \mathbb{P} \left(\sup_{\beta, q \in \mathcal{N}} |D_2(\beta, q) - D_3(\beta, q)| \geq 2C\epsilon\sqrt{2} \right) \quad (136)$$

For a fixed and k , the right hand side goes to zero by the union bound and the second CGMT. Therefore the convergence is uniform in the sense that

$$\mathbb{P} \left(\sup_{0 \leq \beta \leq \beta_{max}, q_{min} \leq q \leq q_{max}} |D_2(\beta, q) - D_3(\beta, q)| \geq \delta \right) \rightarrow 0 \quad (137)$$

for any $\delta > 0$. Finally we can obtain the following optimization problem:

$$\tilde{P}_3 = \max_{0 \leq \beta \leq \beta_{max}} \min_{q_{min} \leq q \leq q_{max}} \max_{0 \leq \xi \leq \xi_{max}} \min_{t_{min} \leq t \leq t_{max}} \mathbb{E} \frac{1}{m} \mathcal{M}_{\frac{1}{2c_1}}(r + \tau_2 h) \left(\theta^* - \frac{c_2^2 \sqrt{m}}{2c_1 \sqrt{n}} \phi \right) - \frac{c_2 m}{4c_1 n} + \frac{\beta^2 \xi d}{2tn} - \frac{(\beta + 2q\tau_1)\xi^2}{2q} - \frac{\beta^2 q d}{2n(\beta + 2q\tau_1)} + \frac{\xi t}{2} + \frac{\beta q}{2} + \frac{\beta}{2q} \sigma_\epsilon^2 - \frac{\beta^2}{2} \quad (138)$$

We have now demonstrated that \tilde{P}_3 converges in probability to \tilde{P}_2 , which subsequently converges to \tilde{P}_1 . This establishes the first part of Theorem 2, about the optimal values. We show the asymptotic equivalence of the generalization error and test functions by following lemma

Lemma 7. Let $\hat{\theta}_2(\tau_1, \tau_2)$ be the solution of P_2 (10) and let $\hat{\theta}_3(\tau_1, \tau_2)$ be the solution of \tilde{P}_3 as given in (138), then

$$\mathcal{E}_{gen}(\hat{\theta}_2(0, 0)) \xrightarrow[n \rightarrow \infty]{P} \tilde{\mathcal{E}}_{gen} \quad (139)$$

$$\frac{1}{m} h(\hat{\theta}_2(0, 0)) \xrightarrow[n \rightarrow \infty]{P} \frac{1}{m} h(\hat{\theta}_3(0, 0)) \quad (140)$$

Proof. We note that for any optimization $P(\tau) = \min_{\mathbf{e}} F(\mathbf{e}) + \tau G(\mathbf{e})$ with optimal solution \mathbf{e}_τ it holds that

$$P(\tau) \leq F(\mathbf{e}_0) + \tau G(\mathbf{e}_0) \quad (141)$$

Applying this observation to our problem with $\tau_1 = \tau$ and $\tau_2 = 0$, we obtain

$$P_2(\tau, 0) \leq P_2(0, 0) + \tau \frac{(\hat{\theta} - \theta^*)^T \tilde{\mathbf{R}}(\hat{\theta} - \theta^*)}{m} \quad (142)$$

From which we obtain that

$$\begin{aligned} \frac{P_2(\tau, 0) - P_2(0, 0)}{\tau} &\leq \frac{(\hat{\theta} - \theta^*)^T \tilde{\mathbf{R}}(\hat{\theta} - \theta^*)}{m} \quad \tau > 0 \\ \frac{(\hat{\theta} - \theta^*)^T \tilde{\mathbf{R}}(\hat{\theta} - \theta^*)}{m} &\leq \frac{P_2(0, 0) - P_2(\tau, 0)}{\tau} \quad \tau < 0 \end{aligned} \quad (143)$$

Take an arbitrary $\delta > 0$. For sufficiently small values of τ and from the convergence of the optimal value we have that

$$\mathbb{P} \left(\frac{(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*)^T \tilde{\mathbf{R}}(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*)}{m} < \frac{\tilde{P}_3(\tau, 0) - \tilde{P}_3(0, 0)}{\tau} + \frac{\delta}{2} \right) \rightarrow 0, \quad \tau > 0 \quad (144)$$

$$\mathbb{P} \left(\frac{(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*)^T \tilde{\mathbf{R}}(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*)}{m} > \frac{\tilde{P}_3(0, 0) - \tilde{P}_3(\tau, 0)}{\tau} - \frac{\delta}{2} \right) \quad \tau < 0 \rightarrow 0 \quad (145)$$

Where this relationship follows from the fact that $\tilde{P}_2(\tau_1, \tau_2)$ converges to $\tilde{P}_3(\tau_1, \tau_2)$ for all $\tau_1 \in [-\tau_1^*, \tau_1^*]$ and $\tau_2 \in [-\tau_2^*, \tau_2^*]$. We also know that for sufficiently small values of $|\tau|$ we have that

$$\left| \frac{\tilde{P}_3(\tau, 0) - \tilde{P}_3(0, 0)}{\tau} - \frac{\partial \tilde{P}_3(\tau_1, 0)}{\partial \tau_1} \Bigg|_{\tau_1=0} \right| \leq \frac{\delta}{2} \quad (146)$$

The uniqueness of the solutions $\hat{t}, \hat{\xi}, \hat{q}, \hat{\beta}$ guarantees that the derivatives exist. We then obtain that

$$\mathbb{P} \left(\left| \frac{(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*)^T \tilde{\mathbf{R}}(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*)}{m} - \frac{\tilde{P}_3(\tau_1, 0)}{\partial \tau_1} \Bigg|_{\tau_1=0} \right| > \delta \right) \rightarrow 0 \quad (147)$$

from which we finally obtain that

$$\frac{(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*)^T \tilde{\mathbf{R}}(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*)}{m} \xrightarrow[n \rightarrow \infty]{P} \frac{\tilde{P}_3(\tau_1, 0)}{\partial \tau_1} \Bigg|_{\tau_1=0} \quad (148)$$

This provides the first result, but we can also compute that

$$\begin{aligned} \frac{\tilde{P}_3(\tau_1, 0)}{\partial \tau_1} \Bigg|_{\tau_1=0} &= \frac{1}{m} \mathbb{E} \left[\left\| \boldsymbol{\theta}^* - \frac{c_2^2 \sqrt{m}}{2c_1 \sqrt{n}} \boldsymbol{\phi} - \text{prox}_{\frac{1}{2c_1}} \left(\boldsymbol{\theta}^* - \frac{c_2^2 \sqrt{m}}{2c_1 \sqrt{n}} \boldsymbol{\phi} \right) \right\|_2^2 \frac{\partial c_1}{\partial \tau_1} \right. \\ &+ \left. \left(\boldsymbol{\theta}^* - \frac{c_2^2 \sqrt{m}}{2c_1 \sqrt{n}} \boldsymbol{\phi} - \text{prox}_{\frac{1}{2c_1}} \left(\boldsymbol{\theta}^* - \frac{c_2^2 \sqrt{m}}{2c_1 \sqrt{n}} \boldsymbol{\phi} \right) \right)^T \left(\frac{c_2^2 \sqrt{m}}{c_1 \sqrt{n}} \frac{\partial c_1}{\partial \tau_1} - \frac{c_2 \sqrt{m}}{\sqrt{n}} \frac{\partial c_2}{\partial \tau_1} \right) \boldsymbol{\phi} \Bigg|_{\tau_1=0} \right] \\ &\quad - \hat{\xi}^2 - \frac{\hat{q}^2 d}{n} \end{aligned} \quad (149)$$

where c_1 and c_2 are evaluated at $\hat{\beta}, \hat{q}, \hat{\xi}, \hat{t}$ and $\tau_1 = 0, \tau_2 = 0$. In this computation we have made use of the following rules for the derivatives of Moreau envelopes

$$\nabla_{\mathbf{x}} \mathcal{M}_{\tau f}(\mathbf{x}) = \frac{1}{\tau} (\mathbf{x} - \text{prox}_{\tau f}(\mathbf{x})) \quad (150)$$

$$\frac{\partial}{\partial \tau} \mathcal{M}_{\tau f}(\mathbf{x}) = -\frac{1}{2\tau^2} \|\mathbf{x} - \text{prox}_{\tau f}(\mathbf{x})\|_2^2 \quad (151)$$

Using the same symmetric logic for the case of τ_2 we find that

$$\frac{h(\hat{\boldsymbol{\theta}}(0, 0)_2)}{m} \xrightarrow[n \rightarrow \infty]{P} \frac{\partial \tilde{P}_3(0, \tau_2)}{\partial \tau_2} \Bigg|_{\tau_2=0} \quad (152)$$

where we find that

$$\frac{\partial \tilde{P}_3(0, \tau_2)}{\partial \tau_2} \Bigg|_{\tau_2=0} = h(\hat{\boldsymbol{\theta}}_3(\hat{\beta}, \hat{q}, \hat{\xi}, \hat{t})) \quad (153)$$

From this we see that

$$\mathbb{E} \frac{1}{m} h(\hat{\boldsymbol{\theta}}_2(0, 0)) \xrightarrow[n \rightarrow \infty]{P} \mathbb{E} \frac{1}{m} h(\hat{\boldsymbol{\theta}}_3(0, 0)) \quad (154)$$

Finally to demonstrate the generalization error we note that

$$\mathcal{E}_{gen}(\hat{\theta}_2) = \mathbb{E} \left(y_{new} - \frac{1}{\sqrt{m}} \tilde{\varphi}(\mathbf{z}_{new})^T \hat{\theta}_2 \right)^2 = \mathbb{E} \left(\epsilon_{new} - \frac{1}{\sqrt{m}} \tilde{\varphi}(\mathbf{z}_{new})^T (\hat{\theta}_2 - \theta^*) \right)^2 \quad (155)$$

in which we have made use of the definition of $y_{new} = \frac{1}{\sqrt{m}} \tilde{\varphi}(\mathbf{z}_{new})^T \theta^* + \epsilon_{new}$. We recall that $\mathbb{E}[\tilde{\varphi}(\mathbf{z}_{new}) \tilde{\varphi}(\mathbf{z}_{new})^T] = \tilde{\mathbf{R}}$. As such we obtain that

$$\mathcal{E}_{gen}(\hat{\theta}_2) = \sigma_\epsilon^2 + \frac{(\hat{\theta} - \theta^*)^T \tilde{\mathbf{R}} (\hat{\theta} - \theta^*)}{m} \quad (156)$$

By the calculation above we see that

$$\mathcal{E}_{gen}(\hat{\theta}_2) \rightarrow \sigma_\epsilon^2 + \left. \frac{\partial \tilde{P}_3(\tau_1, 0)}{\partial \tau_1} \right|_{\tau_1=0} = \tilde{\mathcal{E}}_{gen} \quad (157)$$

□

B.1 Non Deterministic True Vector

In the previous analysis we have assumed that the true vector θ^* has been deterministic. In the case of θ^* being random, we can freeze its value by conditioning on θ^* . The proof holds for a random θ^* with high probability, according to the assumptions. This shows that the results hold for a suitable random θ^* .

B.2 Moreau Envelopes

We remind the reader of the definition of the Moreau Envelope and the proximal operator.

Definition 1. Let $f : \mathcal{X} \rightarrow (-\infty, \infty]$ be a proper, lower semi-continuous function on a Hilbert space \mathcal{X} . Then the Moreau envelope with step size τ of the function is given by

$$\mathcal{M}_{\tau f}(\mathbf{y}) = \min_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}) + \frac{1}{2\tau} \|\mathbf{x} - \mathbf{y}\| \quad (158)$$

The proximal operator of the function f with step size τ is given by

$$\text{prox}_{\tau f}(\mathbf{y}) = \arg \min_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}) + \frac{1}{2\tau} \|\mathbf{x} - \mathbf{y}\| \quad (159)$$

Here we give a lemma concerning the concentration of Moreau envelopes.

Lemma 8 (Gaussian Concentration of Moreau Envelopes, extension of (Loureiro et al. (2021), lemma 5)). *Consider a proper convex function $f : \mathbb{R}^n \rightarrow \mathbb{R}$. Furthermore, let $\mathbf{g} \in \mathbb{R}^n$ be a standard Gaussian random vector and $\mathbf{a} \in \mathbb{R}^n$ a constant vector with finitely bounded norm. Then for any parameter $\tau > 0$ and for any $\epsilon > 0$, there exists a constant c such that*

$$\mathbb{P} \left(\left| \frac{1}{n} \mathcal{M}_{\tau f}(\mathbf{a} + \mathbf{g}) - \mathbb{E} \left[\frac{1}{n} \mathcal{M}_{\tau f}(\mathbf{a} + \mathbf{g}) \right] \right| \geq \epsilon \right) \leq \frac{c}{n\tau^2\epsilon^2} \quad (160)$$

The original lemma as given by (Loureiro et al., 2021) does not have the constant vector \mathbf{a} and instead only considers a Moreau envelope over a Gaussian. We give a proof here for this case but note that the original proof may be applied by instead considering the shifted function $\tilde{f}(\cdot) = f(\cdot - \mathbf{a})$. We give the proof here for completeness.

Proof. First, we show that the Moreau envelope of a convex proper function f is integrable with respect to the Gaussian measure. By making use of the convexity of the optimization problem that defines the Moreau envelope, and because f is proper, there exists a $\mathbf{z}_0 \in \mathbb{R}^n$ and finite constant κ such that

$$\begin{aligned} \frac{1}{n} \mathcal{M}_{\tau f}(\mathbf{g} + \mathbf{a}) &\leq \frac{1}{n} f(\mathbf{z}_0) + \frac{1}{2n\tau} \|\mathbf{z}_0 - \mathbf{g} - \mathbf{a}\|^2 \\ &\leq \kappa + \frac{1}{2n\tau} \|\mathbf{z}_0 - \mathbf{g} - \mathbf{a}\|^2 \end{aligned} \quad (161)$$

The second line is integrable with respect to a Gaussian measure. By means of the Gaussian Poincare inequality (see for example, (Boucheron et al., 2013)).

$$\text{Var} \left[\frac{1}{n} \mathcal{M}_{\tau f}(\mathbf{a} + \mathbf{g}) \right] \leq \frac{c}{n^2} \mathbb{E}_{\mathbf{g}} [\|\nabla_{\mathbf{g}} \mathcal{M}_{\tau f}(\mathbf{a} + \mathbf{g})\|_2^2] = \frac{c}{n^2} \mathbb{E}_{\mathbf{g}} \left\| \frac{1}{\tau} (\mathbf{g} + \mathbf{a} - \text{prox}_{\tau f}(\mathbf{g} + \mathbf{a})) \right\|_2^2 \quad (162)$$

From (Bauschke et al., 2011)[Proposition 12.28 and Proposition 4.4], the function $f(\mathbf{g} + \mathbf{a}) = \mathbf{z} - \text{prox}_{\tau f}(\mathbf{g} + \mathbf{a})$ is firmly non-expansive and

$$\|\mathbf{g} + \mathbf{a} - \text{prox}_{\tau f}(\mathbf{g} + \mathbf{a})\|_2^2 \leq \langle \mathbf{g} + \mathbf{a} | \mathbf{g} - \text{prox}_{\tau f}(\mathbf{g} + \mathbf{a}) \rangle \quad (163)$$

which implies that

$$\|\mathbf{g} + \mathbf{a} - \text{prox}_{\tau f}(\mathbf{g} + \mathbf{a})\|_2^2 \leq \|\mathbf{g} + \mathbf{a}\|_2^2 \quad (164)$$

by means of the Cauchy Swarchz inequality. \square

This implies that

$$\text{var} \left[\frac{1}{n} \mathcal{M}_{\tau f}(\mathbf{a} + \mathbf{g}) \right] \leq \frac{c}{n^2 \tau^2} \mathbb{E} \|\|\mathbf{g} + \mathbf{a}\|_2\| = \frac{c(n + \|\mathbf{a}\|_2^2)}{n^2 \tau^2} \leq \frac{C}{n \tau^2} \quad (165)$$

in which we have used the fact that the norm of \mathbf{a} is bounded. By making use of Chebyshev's inequality we obtain that

$$\mathbb{P} \left(\left| \frac{1}{n} \mathcal{M}_{\tau f}(\mathbf{a} + \mathbf{g}) - \mathbb{E} \left[\frac{1}{n} \mathcal{M}_{\tau f}(\mathbf{a} + \mathbf{g}) \right] \right| \geq \epsilon \right) \leq \frac{c}{n \tau^2 \epsilon^2} \quad (166)$$

C Analysis of Universality

We recall the definition of the perturbed optimization problem as a function of the feature map

$$P(\tau_1, \tau_2) = \min_{\mathbf{e}} \frac{1}{2n} \left\| \epsilon - \frac{1}{\sqrt{m}} \mathbf{X} \mathbf{e} \right\|_2^2 + \frac{1}{m} r(\mathbf{e} + \boldsymbol{\theta}^*) + \frac{1}{m} \tau_1 \mathbf{e} \mathbf{R} \mathbf{e} + \frac{1}{m} \tau_2 h(\mathbf{e} + \boldsymbol{\theta}^*) \quad (167)$$

and

$$\tilde{P}(\tau_1, \tau_2) = \min_{\mathbf{e}} \frac{1}{2n} \left\| \epsilon - \frac{1}{\sqrt{m}} \tilde{\mathbf{X}} \mathbf{e} \right\|_2^2 + \frac{1}{m} r(\mathbf{e} + \boldsymbol{\theta}^*) + \frac{1}{m} \tau_1 \mathbf{e} \mathbf{R} \mathbf{e} + \frac{1}{m} \tau_2 h(\mathbf{e} + \boldsymbol{\theta}^*), \quad (168)$$

where $\mathbf{X}, \tilde{\mathbf{X}}$ are respectively generated by the following two alternative feature maps

$$\boldsymbol{\varphi}(\mathbf{z}) = \sigma \left(\frac{1}{\sqrt{d}} \mathbf{W} \mathbf{z} \right) \quad (169)$$

$$\tilde{\boldsymbol{\varphi}}(\mathbf{z}) = \frac{\rho_1}{\sqrt{d}} \mathbf{W} \mathbf{z} + \rho_* \mathbf{g}, \quad (170)$$

which lead to the following two covariance matrices

$$\mathbf{R} = \mathbb{E}_{\mathbf{z}} [\boldsymbol{\varphi}(\mathbf{z}) \boldsymbol{\varphi}^T(\mathbf{z})] = \mathbb{E}_{\mathbf{z}} \left[\sigma \left(\frac{1}{\sqrt{d}} \mathbf{W} \mathbf{z} \right) \sigma^T \left(\frac{1}{\sqrt{d}} \mathbf{W} \mathbf{z} \right) \right] \quad (171)$$

$$\tilde{\mathbf{R}} = \mathbb{E}_{\mathbf{z}} [\tilde{\boldsymbol{\varphi}}(\mathbf{z}) \tilde{\boldsymbol{\varphi}}^T(\mathbf{z})] = \frac{\rho_1^2}{d} \mathbf{W} \mathbf{W}^T + \rho_*^2 \mathbf{I} \quad (172)$$

Now recall the function $B(\mathbf{e}) = r(\mathbf{e} + \boldsymbol{\theta}^*) + \tau_1 \mathbf{e} \mathbf{R} \mathbf{e} + \tau_2 h(\mathbf{e} + \boldsymbol{\theta}^*)$. We recall that r is assumed to be μ -strongly convex. The values $\tau_1 \in [-\tau_1^*, \tau_1^*]$ and $\tau_2 \in [-\tau_2^*, \tau_2^*]$, with the bounds τ_1^* and τ_2^* chosen to be sufficiently small such that B remains $\frac{\mu}{4}$ -strongly convex

We can now state a theorem concerning Universality that is an extension of Theorem 1 in Hu and Lu (2020)

Theorem 9 (Extension of Hu and Lu (2020)). *Assume that assumptions A3-A6 hold. Fix $\tau_1 \in [-\tau_1^*, \tau_1^*]$ and $\tau_2 \in [-\tau_2^*, \tau_2^*]$. Finally assume that the regularization function $r(\boldsymbol{\theta})$ is strongly convex, thrice differentiable with bounded third derivative.*

Then for every $\epsilon \in (0, 1)$ and every finite constant c , we have that

$$\mathbb{P}(|P(\tau_1, \tau_2) - c| \geq 2\epsilon) \leq \mathbb{P}(|\tilde{P}(\tau_1, \tau_2) - c| \geq \epsilon) + \frac{\text{polylog}m}{\epsilon\sqrt{m}} \quad (173)$$

and

$$\mathbb{P}(|\tilde{P}(\tau_1, \tau_2) - c| \geq 2\epsilon) \leq \mathbb{P}(|P(\tau_1, \tau_2) - c| \geq \epsilon) + \frac{\text{polylog}m}{\epsilon\sqrt{m}} \quad (174)$$

for $m \geq \frac{1}{\epsilon^2}$, in which $\text{polylog}m$ is a function that grows no faster than a polynomial of $\log m$. Consequently,

$$P(\tau_1, \tau_2) \xrightarrow[n, m, d \rightarrow \infty]{P} c \quad \text{iff} \quad \tilde{P}(\tau_1, \tau_2) \xrightarrow[n, m, d \rightarrow \infty]{P} c \quad (175)$$

This theorem is different than the one presented by (Hu and Lu, 2020) in two ways. Firstly we have restricted ourselves to the square loss function which simplifies this analysis, we discuss this difference in remark 2. Secondly, the term associated with τ_2 is different. We consider generic test functions $h(\boldsymbol{\theta})$ satisfying assumptions A2, (Hu and Lu, 2020) only consider one particular case of $h(\boldsymbol{\theta}) = \frac{\rho_1\sqrt{m}}{\sqrt{d}}\boldsymbol{\xi}^T\mathbf{W}\boldsymbol{\theta}$ in which $\boldsymbol{\xi}$ is their teacher vector. The changes required to their proof to apply to generic test functions are minimal, and we give an outline in proof sketch below.

Remark 2. We note that the conditions considered by Hu and Lu (2020) are slightly different than the case considered here. However the proof is sufficiently generic that it applies to the case considered here. Specifically, (Hu and Lu, 2020) consider a generic strongly convex and thrice differentiable loss function $l(\frac{1}{\sqrt{m}}\boldsymbol{\varphi}(\mathbf{z}_i)\boldsymbol{\theta}, y_i)$ for a particular data element i . For the labels y_i , (Hu and Lu, 2020) consider a function $\psi_{\text{teach}}(\mathbf{z}_i^T\boldsymbol{\xi})$ in which $\boldsymbol{\xi}$ is a teacher vector and ψ_{teach} is a differentiable function (except at a finite number of points) and is bounded by

$$\forall x \in \mathbb{R} \quad \psi_{\text{teach}}(x) \leq C(1 + |x|^K) \quad (176)$$

for some constants $C > 0$ and positive integer K . They then prove their results for the joint distribution $(\frac{1}{\sqrt{m}}\boldsymbol{\varphi}(\mathbf{z})^T\boldsymbol{\theta}; \mathbf{z}^T\boldsymbol{\xi})$, which is jointly Gaussian through the variable \mathbf{z} . In the case considered in this paper, we consider $y_i = \frac{1}{\sqrt{m}}\boldsymbol{\varphi}(\mathbf{z}_i)\boldsymbol{\theta}^* + \epsilon_i$ for some known vector $\boldsymbol{\theta}^*$ and noise ϵ_i , and specifically choose the square loss. This allows for the definition of the error vector $\mathbf{e} = \boldsymbol{\theta} - \boldsymbol{\theta}^*$, and allows us to instead consider the distribution $(\frac{1}{\sqrt{m}}\boldsymbol{\varphi}(\mathbf{z})\mathbf{e}; \epsilon_i)$ which simplifies the analysis in this case.

C.1 Proof sketch

Here we discuss how to extend the results of Hu and Lu (2020) to the case of generic test function $h(\boldsymbol{\theta})$, instead of their particular choice of $\frac{\rho_1\sqrt{m}}{\sqrt{d}}\boldsymbol{\xi}^T\mathbf{W}\boldsymbol{\theta}$. The structure and details of the entire proof remain almost unchanged, except for the following set of minor changes, where the equation numbers refers to Hu and Lu (2020):

- In equation 172 step (a) and in the proof of (Hu and Lu, 2020) lemma 19, the property that $H_{\setminus k} \succeq \frac{\mu}{2}\mathbf{I}$, where

$$H_{\setminus k} = \frac{1}{m} \sum_{i=0}^{k-1} l''\left(\frac{1}{\sqrt{m}}\tilde{\boldsymbol{\varphi}}(\mathbf{z}_i), \bar{\mathbf{e}}\right)\tilde{\boldsymbol{\varphi}}(\mathbf{z}_i)\tilde{\boldsymbol{\varphi}}^T(\mathbf{z}_i) + \frac{1}{m} \sum_{i=k+1}^n l''\left(\frac{1}{\sqrt{m}}\boldsymbol{\varphi}(\mathbf{z}_i)\bar{\mathbf{e}}\right)\boldsymbol{\varphi}(\mathbf{z}_i)\boldsymbol{\varphi}^T(\mathbf{z}_i) + \text{diag}\{r''(\bar{\mathbf{e}} + \boldsymbol{\theta}^*)\} + \nabla^2(\tau_1\bar{\mathbf{e}}^T\mathbf{R}\bar{\mathbf{e}} + \tau_2h(\bar{\mathbf{e}} + \boldsymbol{\theta}^*)), \quad (177)$$

where $\bar{\mathbf{e}}$ is the optimal solution to the problem given in Hu and Lu (2020) (equation 32), l is the loss function and l'' its second derivative, in our case the square loss. For the case of for our choice of $\tau_2 \in [-\tau_2^*, \tau_2^*]$ and assumptions A2, and recalling that r is μ strongly convex, this property holds.

- Similarly they require that $R_{\setminus k}$ given in equation 187, defined as

$$R_{\setminus k}(\boldsymbol{\theta}) = \sum_{i \neq k} l\left(\frac{1}{\sqrt{m}}\boldsymbol{\varphi}(\mathbf{z}_i)^T\boldsymbol{\theta}\right) + \sum_{j=1}^m r(\boldsymbol{\theta}) + \tau_1\boldsymbol{\theta}^T\mathbf{R}\boldsymbol{\theta} + \tau_2h(\boldsymbol{\theta}) \quad (178)$$

to be $\frac{\mu}{2}$ -strongly convex. Which obviously holds with our restrictions on h .

- In equation 210 they require that that $G(\mathbf{e}) = r(\mathbf{e} + \boldsymbol{\theta}^*) + \tau_1 \mathbf{e}^T \mathbf{R} \mathbf{e} + \tau_2 h(\mathbf{e} + \boldsymbol{\theta}^*)$ is $\frac{\mu}{2}$ strongly convex and that

$$\|\nabla G(\mathbf{0})\| \leq C\sqrt{m} \quad (179)$$

Which is clearly satisfied by assumption A2 and assumption A5 (see errata for updated assumption A5).

- Finally in equation 252 they require that $c, c', C > 0$

$$\mathbb{P}\left(\max_i |(\nabla h(\boldsymbol{\theta}^*))_i| > c \log m\right) \leq C e^{-c'(\log m)^2} \quad (180)$$

This boundedness is satisfied by assumption A5.

As these are the only changes necessary to prove (Hu and Lu, 2020) results for more generic test functions we do not reproduce the proof here in full.

C.2 Universality of Generalization Error and Test functions h

In this section we demonstrate that the universality of generalization error holds for strongly convex and thrice differentiable regularization functions, making use of the perturbation that we defined above in problem 167. We prove the following result based on results from Hu and Lu (2020). For this theorem we require the following definition

Definition 2. In Theorem 2, we showed that $\tilde{P}(\tau_1, \tau_2) \xrightarrow{P} \tilde{P}_3(\tau_1, \tau_2)$. Let the partial derivatives of $\tilde{P}(\tau_1, \tau_2)$ at $\tau_1 = \tau_2 = 0$ be denoted by $\frac{\partial}{\partial \tau_1} \tilde{P}_3(0, 0) = \hat{\kappa}$ and $\frac{\partial}{\partial \tau_2} \tilde{P}_3(0, 0) = \hat{\pi}$.

Note that derivatives may be readily computed as done in lemma 7. We now state the following result

Theorem 10 (Universality of Generalization Error). *Assume the same assumptions hold as in theorem 9 and let $\hat{\kappa}$ and $\hat{\pi}$ be given in definition 2. Take the Generalization error for a given feature map as*

$$\mathcal{E}_{gen}(\boldsymbol{\theta}, \boldsymbol{\varphi}) = \mathbb{E}\left(y_{new} - \frac{1}{\sqrt{m}} \boldsymbol{\varphi}(\mathbf{z}_{new})^T \boldsymbol{\theta}\right)^2, \quad (181)$$

where $\mathbf{z}_{new} \sim \mathcal{N}(0, I_d)$ and $y_{new} = \frac{1}{\sqrt{m}} \boldsymbol{\varphi}(\mathbf{z}_{new})^T \boldsymbol{\theta}^* + \epsilon_{new}$, where ϵ_{new} is noise. Then

$$\mathcal{E}_{gen}(\hat{\boldsymbol{\theta}}_1, \boldsymbol{\varphi}) \rightarrow \mathcal{E}_{gen}^* \quad \text{and} \quad \mathcal{E}_{gen}(\hat{\boldsymbol{\theta}}_2, \tilde{\boldsymbol{\varphi}}) \rightarrow \mathcal{E}_{gen}^* \quad (182)$$

in which

$$\mathcal{E}_{gen}^* = \sigma_\epsilon^2 + \hat{\kappa} \quad (183)$$

Proof. We let $\tau_2 = 0$ and let $\mathbf{z}_{new} \sim \mathcal{N}(0, I_n)$ be a new Gaussian vector that is independent of all other training samples, and let $y_{new} = \frac{1}{\sqrt{m}} \boldsymbol{\theta}^{*T} \boldsymbol{\varphi}(\mathbf{z}) + \epsilon_{new}$. We can then express the generalization errors as

$$\mathcal{E}_{gen}(\hat{\boldsymbol{\theta}}_1, \boldsymbol{\varphi}) = \mathbb{E}_{\epsilon_{new}, \mathbf{z}_{new}} \left[\epsilon_{new} - \frac{1}{\sqrt{m}} \boldsymbol{\varphi}(\mathbf{z}_{new})^T \hat{\boldsymbol{\theta}}_1 \right]^2 = \sigma_\epsilon^2 + \frac{1}{m} \hat{\mathbf{e}}_1^T \mathbf{R} \hat{\mathbf{e}}_1 \quad (184)$$

$$\mathcal{E}_{gen}(\hat{\boldsymbol{\theta}}_2, \tilde{\boldsymbol{\varphi}}) = \mathbb{E}_{\epsilon_{new}, \mathbf{z}_{new}} \left[\epsilon_{new} - \frac{1}{\sqrt{m}} \tilde{\boldsymbol{\varphi}}(\mathbf{z}_{new})^T \hat{\boldsymbol{\theta}}_2 \right]^2 = \sigma_\epsilon^2 + \frac{1}{m} \hat{\mathbf{e}}_2^T \tilde{\mathbf{R}} \hat{\mathbf{e}}_2 \quad (185)$$

Let $\kappa_2 = \frac{1}{m} \hat{\mathbf{e}}_2^T \tilde{\mathbf{R}} \hat{\mathbf{e}}_2$, from which we see that $\mathcal{E}_{gen}(\tilde{\boldsymbol{\varphi}}) = \sigma_\epsilon^2 + \kappa_2$. We start by noting that by lemma 7 that $\mathcal{E}_{gen}(\hat{\boldsymbol{\theta}}_2, \boldsymbol{\varphi}_2) = \sigma_\epsilon^2 + \kappa_2 \rightarrow \sigma_\epsilon^2 + \hat{\kappa}$, which proves the second claim. Now, we consider the value of $\kappa_1 = \frac{1}{m} \hat{\mathbf{e}}_1^T \mathbf{R} \hat{\mathbf{e}}_1$. By the definition of the optimization problem we have

$$P(\tau_1, \tau_2 = 0) \leq P(0, 0) + \tau_1 \hat{\mathbf{e}}_2^T \mathbf{R} \hat{\mathbf{e}}_2 \quad (186)$$

For any τ_1 . From this it follows that for any $\tau > 0$ we have

$$\frac{P(\tau, 0) - P(0, 0)}{\tau} \leq \kappa_1 \leq \frac{P(-\tau, 0) - P(0, 0)}{-\tau} \quad (187)$$

We choose an $\epsilon > 0$. By definition 2 the limit function $\tilde{P}_3(\tau_1, \tau_2)$ is differentiable at the origin and we know from theorem 9 and 2 that $P(\tau_1, \tau_2) \xrightarrow{P} \tilde{P}_3(\tau_1, \tau_2)$. On the other hand there exists some $\delta > 0$ such that

$$\left| \frac{\tilde{P}_3(\delta, 0) - \tilde{P}_3(0, 0)}{\delta} - \hat{\kappa} \right| \leq \frac{\epsilon}{3}. \quad (188)$$

Substituting this into the first inequality of equation 187 above and letting $\tau = \delta$ we obtain

$$\begin{aligned} \mathbb{P}(\kappa_1 - \hat{\kappa} < -\epsilon) &\leq \mathbb{P}\left(\frac{P(\delta, 0) - P(0, 0)}{\delta} - \hat{\kappa} < -\epsilon\right) \\ &\leq \mathbb{P}(|P(\delta, 0) - \tilde{P}_3(\delta, 0)| \geq \delta\epsilon/3) + \mathbb{P}(|P(0, 0) - \tilde{P}_3(0, 0)| \geq \delta\epsilon/3) \end{aligned} \quad (189)$$

Now by assumption we have that $P(\delta, 0) \xrightarrow{P} \tilde{P}_3(\delta, 0)$ and $P(0, 0) \xrightarrow{P} \tilde{P}_3(0, 0)$. It then follows from Eq 189 that $\lim_{n \rightarrow \infty} \mathbb{P}(\kappa_1 - \hat{\kappa} < -\epsilon) = 0$. The exact same reasoning may be applied to second inequality 187 to obtain that $\lim_{n \rightarrow \infty} \mathbb{P}(\kappa_1 - \hat{\kappa} > \epsilon) = 0$ as such $\kappa_1 \xrightarrow{P} \hat{\kappa}$. □

We now prove the universality of the test functions $h(\boldsymbol{\theta})$.

Theorem 11 (Universality of Test Functions). *Assume that the same assumptions hold as in theorem 9 and let $\hat{\pi}$ be given in definition 2. Then*

$$\frac{1}{m}h(\hat{\boldsymbol{\theta}}_2) \rightarrow \hat{\pi} \quad \text{and} \quad \frac{1}{m}h(\hat{\boldsymbol{\theta}}_1) \rightarrow \hat{\pi} \quad (190)$$

Proof. Our proof takes a similar form to the proof of theorem 10. We let $\tau_1 = 0$. Then we note that by the definition of the optimization problems

$$P(\tau_1 = 0, \tau_2) \leq P(0, 0) + \tau_2 h(\hat{\boldsymbol{\theta}}_1) \quad (191)$$

for any τ_2 . It follows that for any $\tau > 0$ we have that

$$\frac{P(0, \tau) - P(0, 0)}{\tau} \leq h(\hat{\boldsymbol{\theta}}_1) \leq \frac{P(0, 0) - P(0, -\tau)}{-\tau} \quad (192)$$

We choose $\epsilon > 0$. By definition 2 the limit function $\tilde{P}_3(\tau_1, \tau_2)$ is differentiable at the origin. Therefore there exists some δ such that

$$\left| \frac{\tilde{P}_3(0, \delta) - \tilde{P}_3(0, 0)}{\delta} - \hat{\pi} \right| \leq \frac{\epsilon}{3} \quad (193)$$

we substitute this into the first inequality of equation (192) above and let $\tau = \delta$. We obtain

$$\begin{aligned} \mathbb{P}\left(\frac{1}{m}h(\hat{\boldsymbol{\theta}}_1) - \hat{\pi} < -\epsilon\right) &\leq \mathbb{P}\left(\frac{P(0, \delta) - P(0, 0)}{\delta} - \hat{\pi} < -\epsilon\right) \\ &\leq \mathbb{P}(|P(0, \delta) - \tilde{P}_3(0, \delta)| > \delta\epsilon/3) + \mathbb{P}(|P(0, 0) - \tilde{P}_3(0, 0)| > \delta\epsilon/3) \end{aligned} \quad (194)$$

Because of the universality laws given in theorem (9) and Theorem 2, we know that $P(\tau_1, \tau_2) \xrightarrow{P} \tilde{P}_3(\tau_1, \tau_2)$. It then follows that $\lim_{n \rightarrow \infty} \mathbb{P}\left(\frac{1}{m}h(\hat{\boldsymbol{\theta}}_1) - \hat{\pi} < -\epsilon\right) = 0$. The exact same reasoning may be applied to the second inequality of (192) to obtain that $\lim_{n \rightarrow \infty} \mathbb{P}\left(\frac{1}{m}h(\hat{\boldsymbol{\theta}}_1) - \hat{\pi} > \epsilon\right) = 0$. As such we conclude that $\frac{1}{m}h(\hat{\boldsymbol{\theta}}_1) \rightarrow \hat{\pi}$. Similar argument proves the result for $\hat{\boldsymbol{\theta}}_2$. □

C.3 Proof of Theorem 4

We now consider our theorem 4. We first show that the thrice differentiability condition of $r(\boldsymbol{\theta})$ can be lifted in the case that there exist a sequence of function that are differentiable and converge to r .

Lemma 9. *Let $r^{(k)}(\boldsymbol{\theta})$ be a sequence of functions that are each thrice differentiable and strongly convex. Assume further that $r^{(k)}(\boldsymbol{\theta})$ converge uniformly to the regularization function $r(\boldsymbol{\theta})$ in the limit of $k \rightarrow \infty$. Then, the results of theorem 9 hold for this regularization function $r(\boldsymbol{\theta})$.*

Proof. We define $P^{(k)}(\tau_1, \tau_2), \tilde{P}^{(k)}(\tau_1, \tau_2)$ to be the optimal cost for the regularization function $r^{(k)}$ respectively with feature map $\boldsymbol{\varphi}, \tilde{\boldsymbol{\varphi}}$. We choose k to be sufficiently large such that $|r^{(k)}(\boldsymbol{\theta}) - r(\boldsymbol{\theta})| < m\epsilon$ for every $\boldsymbol{\theta}$. This implies that $|P^{(k)}(\tau_1, \tau_2) - P(\tau_1, \tau_2)| < \epsilon$ and $|\tilde{P}^{(k)}(\tau_1, \tau_2) - \tilde{P}(\tau_1, \tau_2)| < \epsilon$. Furthermore, by theorem 9, we have

$$\mathbb{P}(|P^{(k)}(\tau_1, \tau_2) - c| > 4\epsilon) \leq \mathbb{P}(|\tilde{P}(\tau_1, \tau_2) - c| > 2\epsilon) + \frac{\text{polylog}(m)}{\sqrt{m}} \quad (195)$$

and hence

$$\mathbb{P}(|P(\tau_1, \tau_2) - c| > 5\epsilon) \leq \mathbb{P}(|\tilde{P}(\tau_1, \tau_2) - c| > \epsilon) + \frac{\text{polylog}(m)}{\sqrt{m}} \quad (196)$$

The other case is similarly proven. \square

We note that this proof hold analogously for $r(\boldsymbol{\theta}) + \tau_2 h(\boldsymbol{\theta})$ for a test function $h(\boldsymbol{\theta})$ that satisfies the conditions of assumption A2. This completes the proof of theorem 4.

C.4 Proof of Corollary 1

We now consider elastic net regularization

$$r(\boldsymbol{\theta}) = \lambda \|\boldsymbol{\theta}\|_1 + \frac{\epsilon}{2} \|\boldsymbol{\theta}\|_2^2 \quad (197)$$

The following lemma demonstrates that can construct a sequence of regularization function that uniformly converges to Eq. 197. This result together with Theorem 4 shows that for all $\epsilon > 0$ universality is established.

Lemma 10. *There exists a sequence of function $r^{(k)}(\boldsymbol{\theta})$ that are separable, strongly convex and thrice differentiable that converge uniformly to the elastic net regularization function given in Eq. 197.*

Proof. Define $h^{(k)}(x)$ as

$$h^{(k)}(x) = x \operatorname{erf}\left(\frac{\sqrt{k}x}{\sqrt{2}}\right) + \frac{\sqrt{\frac{2}{\pi}} e^{-\frac{kx^2}{2}}}{\sqrt{k}} \quad (198)$$

in which erf is the error function. It is simple to verify that $h^{(k)}(x)$ is thrice differentiable and has bounded third derivative. The maximum difference between $|x|$ and $h^{(k)}(x)$ is at $x = 0$ and is $\sqrt{\frac{2}{k\pi}}$. A such in the limit of $k \rightarrow \infty$, $h^{(k)}(x)$ converges uniformly to the absolute value function. We choose

$$r^{(k)}(\boldsymbol{\theta}) = \frac{\epsilon}{2} \|\boldsymbol{\theta}\|_2^2 + \sum_i^m h^{(k)}(\theta_i) \quad (199)$$

This regularization function satisfies the conditions of the lemma. \square

C.5 Proof of Theorem 6

First consider the universality with elastic net which is proven in Section C.4. To demonstrate universality with respect to the ℓ_1 norm, we take the case of elastic net in (197) with a sufficiently small ϵ drop the quadratic part of $r(\boldsymbol{\theta})$. We note that for any $\epsilon > 0$, Theorem 4 holds. Our goal will be to show that for very small values of ϵ removing ϵ does not substantially change the value of the training and testing error. We first make the following definitions

Definition 3. Consider an $m \times n$ matrix \mathbf{A} .

1. For any $k \in \mathbb{N}$ such that $k < n$, the RIP constant $\delta_k(\mathbf{A})$ is the smallest number δ , such that for any index subset $I \subset \{1, 2, \dots, n\}$ with $|I| \leq k$

$$1 - \delta \leq \sigma_{\min}^2(\mathbf{A}_I) \leq \sigma_{\max}^2(\mathbf{A}_I) \leq 1 + \delta \quad (200)$$

in which σ_{\min} and σ_{\max} are the minimum and maximum singular values.

2. Let $\theta_k(\mathbf{A})$ for any $k < n/2$ be the smallest number θ such that for any disjoint subsets $I, I' \subset \{1, 2, \dots, n\}$ with $|I|, |I'| \leq k$ it holds that

$$\sigma_{\max}(\mathbf{A}_{I'}^T \mathbf{A}_I) \leq \theta \quad (201)$$

In which $\sigma_{\max}(\mathbf{X})$ is the maximum singular value of a matrix \mathbf{X} . It is known that $\theta_k \leq \delta_{2k}$.

3. We define the admissible sparsity $M_{adm}(\mathbf{A})$ as

$$M_{adm}(\mathbf{A}) = \sup_k \frac{k[1 - \delta_k(\mathbf{A})]_+}{2n} \quad (202)$$

in which $[\cdot]_+$ represents the positive part.

This admissible sparsity M_{adm} is the constant ρ given in Theorem 6. Provided that M_0 , the effective sparsity given in (33), is strictly less than M_{adm} the theorem holds.

For our proof we require the following lemma. The original lemma is given in Panahi and Hassibi (2017) but we have extended it here to a slightly more general setup.

Lemma 11 (Extension of Panahi and Hassibi (2017) lemma 8). *Suppose that σ is 1-Lipschitz. For the feature matrix $\mathbf{X} = \left(\sigma \left(\frac{1}{\sqrt{d}} \mathbf{w}_j^T \mathbf{z}_i \right) \right)_{ij}$ and $\mathbf{A} = \frac{1}{\sqrt{n(\rho_*^2 + \rho_1^2)}} \mathbf{X}$ there exist constants $\alpha, \beta > 0$ and $1 > \epsilon > 0$ such that*

$$\lim_{n \rightarrow \infty} \mathbb{P}(\delta_{\alpha n}(\mathbf{A}) + \theta_{\alpha n}(\mathbf{A}) > 1 - \epsilon) = 0 \quad (203)$$

$$\lim_{n \rightarrow \infty} \mathbb{P}(\sigma_{\max}(\mathbf{A}) > \beta) = 0 \quad (204)$$

Proof. Consider an arbitrary subset $I \subset [m]$ with $|I| = k$. Let $\mathbf{u} \in \mathbb{S}^{k-1}$, where \mathbb{S}^{k-1} is the surface of the unit sphere in \mathbb{R}^k . We note given \mathbf{W} that $\mathbf{y} = \mathbf{A}_I \mathbf{u}$ is an i.i.d, centered vector and defining

$$f(\mathbf{x}) = \frac{1}{\sqrt{(\rho_*^2 + \rho_1^2)}} \sum_{i=1}^k \sigma \left(\frac{\mathbf{w}_j^T \mathbf{x}}{\sqrt{d}} \right) u_j, \quad (205)$$

we have $y_i = \frac{1}{\sqrt{n}} f(\mathbf{x}_i)$. We observe that

$$\nabla f(\mathbf{x}) = \frac{1}{\sqrt{d(\rho_*^2 + \rho_1^2)}} \mathbf{W} \mathbf{v} \quad (206)$$

where $\mathbf{v} = \left(\sigma' \left(\frac{\mathbf{w}_j^T \mathbf{x}}{\sqrt{d}} \right) u_j \right)_j$. Note that by Lipschitz continuity $\|\mathbf{v}\|_2 \leq 1$ and hence

$$\|\nabla f\| \leq \frac{1}{\sqrt{d(\rho_*^2 + \rho_1^2)}} \|\mathbf{W}\|_2 \quad (207)$$

Hence, by the standard random matrix results, f is $\mu' = \mu \sqrt{1 + \frac{m}{d}}$ -Lipschitz, where $\mu := \frac{1}{\sqrt{(\rho_*^2 + \rho_1^2)}}$, with high probability. Hence, y_i is $\frac{\mu'}{\sqrt{n}}$ -sub-Gaussian. The rest of the argument is conditioned on the event that $\|\mathbf{W}\|_2$ is bounded and hence leads to μ' -Sub-Gaussian variables. Define

$$\sigma^2 := n \text{Var}(y_i) = \mu^2 \sum_{j,j'}^k u_j u_{j'} R_{jj'}, \quad (208)$$

where $\mathbf{R} = (R_{jj'})$ is the exact covariance matrix of the features. It has been shown in the previous works e.g. Hu and Lu (2020), that exists a constant c such that

$$\mathbb{P}(|\sigma - 1| > \epsilon) \leq \frac{1}{c} e^{-nc\epsilon^2} \quad (209)$$

Now, since y_i are μ -sub-Gaussian, it is standard to show that (Honorio and Jaakkola (2014))

$$\forall \lambda : |\lambda| < \frac{n}{4\mu^2}, \quad \mathbb{E} \left[e^{\lambda(y_i^2 - \frac{1}{n}\sigma^2)} \mid \mathbf{W} \right] \leq e^{\frac{16\mu'^4\lambda^2}{n^2}} \quad (210)$$

and hence, conditioned on \mathbf{W} (which satisfies μ' -sub-Gaussianity) we have

$$\mathbb{P}(\|\mathbf{y}\|_2^2 \geq \sigma^2 + \epsilon) = \mathbb{P} \left(\sum_i^n (y_i^2 - \frac{1}{n}\sigma^2) \geq \epsilon \right) \leq \min_{0 < \lambda < \frac{n}{4\mu^2}} \left(\mathbb{E} \left[e^{\lambda \sum_{i=1}^n (y_i^2 - \frac{1}{n}\sigma^2)} \right] e^{-\lambda\epsilon} \right) \leq \min_{0 < \lambda < \frac{n}{4\mu^2}} e^{\frac{16\mu'^4\lambda^2}{n} - \lambda\epsilon} \quad (211)$$

Hence, for sufficiently small ϵ , we may choose $\lambda = \frac{n\epsilon}{32\mu'^4}$ and obtain

$$\mathbb{P}(\|\mathbf{y}\|_2^2 \geq \sigma^2 + \epsilon) \leq e^{-c n \epsilon^2} \quad (212)$$

where c is a suitable constant that may grow in each appearance. We conclude that for a random \mathbf{W} we have

$$\mathbb{P}(\|\mathbf{y}\|_2^2 \geq 1 + 2\epsilon) \leq \mathbb{P}(\|\mathbf{y}\|_2^2 \geq \sigma^2 + \epsilon) + \mathbb{P}(\sigma^2 \geq 1 + \epsilon) \leq \frac{1}{c} e^{-c n \epsilon^2} \quad (213)$$

We may repeat the above Chernoff bound on the event $\|\mathbf{y}\|_2^2 \leq \sigma^2 - \epsilon$, to conclude that

$$\mathbb{P}(|\|\mathbf{y}\|_2^2 - 1| \geq 2\epsilon) \leq \frac{1}{c} e^{-c n \epsilon^2} \quad (214)$$

The rest of the proof is similar to Panahi and Hassibi (2017). We note that for every $\Delta > 0$ there exists a set $G_k \subset \mathbb{S}^{k-1}$ of maximally $(\frac{3}{\Delta})^k$ points such that for any $\mathbf{u} \in \mathbb{S}^{k-1}$ there exists a point $\mathbf{u}_1 \in G_k$ such that $\|\mathbf{u} - \mathbf{u}_1\|_2 \leq \Delta$. We denote $B = \max_{\mathbf{u} \in G_k} \|\mathbf{A}_I \mathbf{u}\|_2$ and $A = \sigma_{max}(\mathbf{A}_I) = \max_{x \in \mathbb{S}^{k-1}} \|\mathbf{A} \mathbf{u}\|_2$ with its maximum being at \mathbf{u}_0 . From this we see that

$$A = \|\mathbf{A}_I \mathbf{u}_0\|_2 \leq \|\mathbf{A}_I \mathbf{u}_1\|_2 + \|\mathbf{A}_I(\mathbf{u}_1 - \mathbf{u}_0)\|_2 \leq B + \Delta A \quad (215)$$

in which \mathbf{u}_1 is the point in G_n closest to \mathbf{u}_0 . If $\Delta < 1$ we obtain

$$\sigma_{max}(\mathbf{A}_I) \leq \frac{\max_{x \in G_n} \|\mathbf{A}_I \mathbf{u}\|_2}{1 - \Delta} \quad (216)$$

This argument may be repeated for the minimum singular value to obtain

$$\sigma_{min}(\mathbf{A}_I) \geq \min_{\mathbf{u} \in G_n} \|\mathbf{A}_I \mathbf{u}\|_2 - \sigma_{max}(\mathbf{A}_I) \Delta \quad (217)$$

From equation 216 we see that $\epsilon_0 = \delta - \Delta - \Delta\delta > 0$, we have

$$\mathbb{P}(\sigma_{max}(\mathbf{A}_I) > 1 + \delta) \leq \mathbb{P} \left(\min_{\mathbf{u} \in G_n} \|\mathbf{A} \mathbf{u}\|_2 > (1 - \Delta)(1 + \delta) \right) \leq \frac{1}{c} e^{-c n \epsilon_0^2} \left(\frac{3}{\Delta} \right)^k \quad (218)$$

and

$$\mathbb{P}(\sigma_{min}(\mathbf{A}_I) < 1 - \delta) \leq \mathbb{P}(\min_{\mathbf{u} \in G_n} \|\mathbf{A} \mathbf{u}\|_2 < 1 - \delta + (1 + \delta)\Delta) + \mathbb{P}(\sigma_{max}(\mathbf{A}_I) > 1 + \delta) \leq \frac{2}{c} e^{-c n \epsilon_0^2} \left(\frac{3}{\Delta} \right)^k \quad (219)$$

Choose $k = n$, $\Delta = \frac{1}{2}$, $\epsilon_0^2 > \frac{\log 6}{c}$ and $\delta = 1 + 2\epsilon_0$. We observe that $\mathbb{P}(\sigma_{max}(\mathbf{A}_I) > 1 + \delta) \rightarrow 0$, which proves the second part. For the first part, note that by the union bound

$$\mathbb{P}(\delta_k(\mathbf{A}) > \delta) \leq \frac{3}{c} e^{-c n \epsilon_0^2} \left(\frac{3}{\Delta} \right)^k \binom{n}{k} \quad (220)$$

Take for example $\Delta = \frac{1}{5}$, $\delta = \frac{1}{3}$, hence $\epsilon_0 = \frac{1}{5}$. Furthermore, for $k = 2\alpha n$, we have $\binom{n}{k} \sim e^{nH(2\alpha)}$ where $H(p) = -p \log p - (1-p) \log(1-p)$ is the entropy function. Choosing α small enough such that $H(2\alpha) + 2\alpha \log 15 < c\epsilon_0^2$ will lead to $\mathbb{P}(\delta_k(\mathbf{A}) > \delta) \rightarrow 0$. We conclude the first result by noting that $\delta_{\alpha n} + \theta_{\alpha n} \leq 2\delta_{2\alpha n}$.

□

We note that the feature matrix \mathbf{X} satisfies this lemma for most practical choices of the activation function. \tanh and the error function are both odd activation functions that satisfy the assumptions A6 and produce a suitable matrix \mathbf{X} .

By the above lemma, we conclude Theorem 6 as the rest of the proof in Panahi and Hassibi (2017) will hold true. For the sake of completeness we repeat these proofs in full.

Theorem 12. *Let assumptions A2-A6 hold and let $r(\boldsymbol{\theta}) = \lambda \|\boldsymbol{\theta}\|_1$. Denote*

$$P_\lambda = \min_{\mathbf{e}} \frac{1}{2n} \|\boldsymbol{\epsilon} + \mathbf{X}\mathbf{e}\|_2^2 + \frac{\lambda}{m} \|\mathbf{e} + \boldsymbol{\theta}^*\|_1 \quad (221)$$

Furthermore assume that there exist constants α, β, ϵ , such that

$$\lim_{n \rightarrow \infty} \mathbb{P}(\delta_{\alpha n}(\mathbf{X}) + \theta_{\alpha n}(\mathbf{X}) > 1 - \epsilon) = 0 \quad (222)$$

$$\lim_{n \rightarrow \infty} \mathbb{P}(\sigma_{max}(\mathbf{X}) > \beta) = 0 \quad (223)$$

Then

$$P_\lambda \xrightarrow[n, m, d \rightarrow \infty]{P} \tilde{P}_{3, \lambda}(\beta, q, \xi, r) \quad (224)$$

In which $\tilde{P}_{3, \lambda}$ is \tilde{P}_3 as given in Eq. (138) for the case that $r(\boldsymbol{\theta}) = \lambda \|\boldsymbol{\theta}\|_1$

Proof. Let $\hat{\mathbf{e}}^{(0)}$ be the minimal point of the optimization

$$P_{\lambda, \mu} = \min_{\mathbf{e}} \frac{1}{2n} \|\boldsymbol{\epsilon} + \mathbf{X}\mathbf{e}\|_2^2 + \frac{\lambda}{m} \|\mathbf{e} + \boldsymbol{\theta}^*\|_1 + \frac{\mu}{2m} \|\mathbf{e}\|_2^2 \quad (225)$$

We know that from Lemma 2 and 4 that there exists a number $C_{\mathbf{e}}$ such that for every $\mu < 1$, $\|\hat{\mathbf{e}}^{(0)}\|_2^2 \leq C_{\mathbf{e}}^2 m$ with high probability. We define

$$p(\mathbf{e}) = \frac{1}{2} \|\boldsymbol{\epsilon} - \mathbf{X}\mathbf{e}\|_2^2 + \frac{\lambda}{\gamma} \|\mathbf{e} + \boldsymbol{\theta}^*\|_1 \quad (226)$$

From the KKT conditions we know that

$$-\mu \hat{\mathbf{e}}^{(0)} \in \partial p(\hat{\boldsymbol{\theta}}^{(0)}) \quad (227)$$

where ∂ represents the subdifferential. We define $\zeta^{(0)} = -\mu \hat{\mathbf{e}}^{(0)}$. We let $k = \alpha n$ and select k entries of $\hat{\mathbf{e}}^{(0)}$ with the largest absolute values and collect their indices in I_0 . We set $\mathbf{a}_0 = 0 \in \mathbb{R}^k$ and let $t = 0$. We now perform the subsequent iterative algorithm.

1. Define $\mathbf{A}_t = \mathbf{X}_{I_t}$ and let $\mathbf{h}_t = \boldsymbol{\epsilon} + \mathbf{X}_{I_t^c} \hat{\mathbf{e}}_{I_t^c}^{(t)}$ and solve

$$\min_{\mathbf{w}} \frac{1}{2} \|\mathbf{h}_t + \mathbf{A}_t \mathbf{w}\|_2^2 + \lambda \|\boldsymbol{\theta}_{I_t}^* + \mathbf{w}\|_1 - \mathbf{a}_t^T \mathbf{w} \quad (228)$$

define its cost function and optimal point by $p_t(\mathbf{w})$ and \mathbf{w}_t respectively

2. Find k elements in I_t^c with largest absolute values in $\mathbf{X}_{I_t^c}^T \mathbf{X}_{I_t}(\mathbf{w}_t - \mathbf{e}_{I_t}^{(t)})$. We denote the indices by I_{t+1} . We set $\mathbf{a}_{t+1} = \zeta_{I_{t+1}}^{(t)}$
3. We construct $\mathbf{e}^{(t+1)}$ and $\zeta^{(t+1)}$ such that $\mathbf{e}_{I_t}^{(t+1)} = \mathbf{w}_t$, $\mathbf{e}_{I_t^c}^{(t+1)} = \mathbf{e}_{I_t^c}^{(t)}$, $\zeta_{I_t}^{(t+1)} = \mathbf{a}_t$ and $\zeta_{I_t^c}^{(t+1)} = \zeta_{I_t^c}^{(t)} + \mathbf{X}_{I_t^c}^T \mathbf{X}_{I_t}(\mathbf{w}_t - \mathbf{e}_{I_t}^{(t)})$
4. we let $t \leftarrow t + 1$ and return to step 1.

In Lemma 12 below we show that this iterative process results in a point $\mathbf{e}^{(\infty)}$ with subgradient $\zeta^{(\infty)} \in \partial p(\mathbf{e}^{(\infty)})$, such that

$$\frac{1}{\sqrt{m}} \|\mathbf{e}^{(\infty)} - \mathbf{e}^{(0)}\|_2 \leq \frac{\mu C_{\mathbf{e}}}{1 - \delta_k - \theta_k} \stackrel{def}{=} \mu C_2 \quad (229)$$

$$\|\zeta^{(\infty)}\|_\infty \leq \mu C_{\mathbf{e}} \left(\sqrt{\frac{m}{k}} + \frac{\theta_k}{1 - \delta_k - \theta_k} \right) \stackrel{def}{=} \mu C_1 \quad (230)$$

We note that $\mathbf{e}^{(\infty)}$ is the optimal point of the optimization

$$\rho_{\mu,\lambda} = \min_{\mathbf{e}} \frac{1}{2} \|\boldsymbol{\epsilon} - \mathbf{X}\mathbf{e}\|_2^2 + \lambda \|\boldsymbol{\theta}^* + \mathbf{e}\|_1 + \mathbf{e}^T \zeta^{(\infty)} \quad (231)$$

We shall let the subscripts λ, μ denote that a particular value of $\zeta^{(\infty)}$ or $\mathbf{e}^{(\infty)}$ are computed for particular values of λ, μ . We now note that $\mathbf{e}^T \zeta^{(\infty)} \leq \|\mathbf{e}\|_1 \|\zeta^{(\infty)}\|_\infty \leq \mu C_1 \|\mathbf{e}\|_1$, and as such

$$\rho_{\mu,\lambda} \leq P_{\lambda+C_1\mu} \quad (232)$$

Or equivalently we can express this as

$$P_\lambda \geq \rho_{\mu,\lambda-C_1\mu} \quad (233)$$

We also note that

$$\begin{aligned} m\rho_{\mu,\lambda} &= \frac{1}{2} \|\boldsymbol{\epsilon} + \mathbf{X}\mathbf{e}^{(\infty)}\|_2^2 + \lambda \|\mathbf{e}^{(\infty)} - \boldsymbol{\theta}^*\|_1 + \mathbf{e}^{(\infty)T} \zeta^{(\infty)} \\ &\geq mP_{\lambda,\mu} + \mathbf{f}^T \mathbf{X}(\mathbf{e}^{(\infty)} - \mathbf{e}^{(0)}) - \lambda \|\mathbf{e}^{(\infty)} - \mathbf{e}^{(0)}\|_1 + \mathbf{e}^{(\infty)T} \zeta^{(\infty)} \\ &\geq mP_{\lambda,\mu} - (\|\mathbf{f}^T \mathbf{X}\|_2 + \lambda\sqrt{m}) \|\mathbf{e}^{(\infty)} - \mathbf{e}^{(0)}\|_2 - \|\mathbf{e}^{(\infty)}\|_2 \|\zeta^{(\infty)}\|_2 \\ &\geq mP_{\lambda,\mu} - (\sigma_{max}(\mathbf{X}) \|\mathbf{f}\|_2 + \lambda\sqrt{m}) \mu C_2 \sqrt{m} - \sqrt{m} \mu C_1 \|\mathbf{e}^{(\infty)}\|_2 \\ &\geq mP_{\lambda,\mu} - (\sigma_{max}(\mathbf{X}) \kappa + \lambda) \mu C_2 m - m \mu C_1 (C_e + C_2 \mu) \end{aligned} \quad (234)$$

in which $\mathbf{f} = \boldsymbol{\epsilon} - \mathbf{X}\mathbf{e}^{(0)}$, κ is a proper bound that is independent of all other parameters, such that $\|\boldsymbol{\epsilon}\|_2 \leq \sqrt{m} \kappa$ with high probability. This holds by the law of large numbers, and we note that $\|\mathbf{f}\|_2 \leq \|\boldsymbol{\epsilon}\|_2$. From this we find that

$$\rho_{\mu,\lambda} \geq P_{\lambda,\mu} - (\sigma_{max}(\mathbf{X})r + \lambda) \mu C_2 - \mu C_1 (C_e + C_2 \mu) \quad (235)$$

Noting that by Theorem 4 that $P_{\lambda,\mu} \rightarrow \tilde{P}_{3\lambda,\mu}(\beta, q, \xi, r)$. We note that by the continuity of $\tilde{P}_{3\lambda,\mu}$ at $\mu = 0$, and for any $\epsilon > 0$, we can select a value of μ small enough such that

$$\mathbb{P}(|P_\lambda - \tilde{P}_{3\lambda,\mu=0}| > \epsilon) \xrightarrow[n,m,d \rightarrow \infty]{P} 0 \quad (236)$$

We also note that for any sufficiently small value of δ we see that

$$\frac{P_\lambda - P_{\lambda-\delta}}{\delta} \leq \frac{\|\hat{\boldsymbol{\theta}}\|_1}{m} \leq \frac{P_{\lambda+\delta} - P_\lambda}{\delta} \quad (237)$$

From this we see that

$$\frac{\|\hat{\boldsymbol{\theta}}\|_1}{m} \xrightarrow[n,m,d \rightarrow \infty]{P} \frac{\partial P_\lambda}{\partial \lambda} \quad (238)$$

□

Lemma 12. *The iterative process defined in Theorem 12 produces a point $\mathbf{e}^{(\infty)}$ with subgradient $\zeta^{(\infty)}$ that are bounded as*

$$\frac{1}{\sqrt{m}} \|\mathbf{e}^{(\infty)} - \mathbf{e}^{(0)}\|_2 \leq \frac{\mu C_e}{1 - \delta_k - \theta_k} \stackrel{def}{=} \mu C_2 \quad (239)$$

$$\|\zeta^{(\infty)}\|_\infty \leq \mu C_e \left(\sqrt{\frac{m}{k}} + \frac{\theta_k}{1 - \delta_k - \theta_k} \right) \stackrel{def}{=} \mu C_1 \quad (240)$$

Proof. Firstly we show that $\zeta_t \in \partial p(\mathbf{e}^{(t)})$. We prove this by means of induction. We note that by definition $\zeta_0 \in \partial p(\mathbf{e}^{(0)})$. For the iteration step we assume that $\zeta_t \in \partial p(\mathbf{e}^{(t)})$. We note that by the KKT conditions of the problem 228 we see that

$$(\zeta_{t+1})_{I_t} = \mathbf{a}_t \in \mathbf{X}_{I_t}^T (\boldsymbol{\epsilon} + \mathbf{X}\mathbf{e}^{(t+1)}) + \partial \|\boldsymbol{\theta}_{I_t}^* + \mathbf{e}^{(t+1)_{I_t}}\|_1 \quad (241)$$

Furthermore we have that

$$(\zeta_t)_{I_t^c} \in \mathbf{X}_{I_t^c}^T (\boldsymbol{\epsilon} + \mathbf{X}\mathbf{e}^{(t)}) + \partial \|\boldsymbol{\theta}_{I_t^c}^* + \mathbf{e}_{I_t^c}^{(t)}\|_1 \quad (242)$$

From which we can see that

$$(\zeta_{t+1})_{I_t^c} \in -\mathbf{X}_{I_t^c}^T(\boldsymbol{\epsilon} + \mathbf{X}\mathbf{e}^{(t+1)}) + \partial\|\boldsymbol{\theta}_{I_t^c}^* + \mathbf{e}_{I_t^c}^{(t+1)}\|_1 \quad (243)$$

This shows that $\zeta_{t+1} \in \partial p(\mathbf{e}^{(t+1)})$. This completes the induction.

Next we will show by induction that

$$\frac{1}{\sqrt{m}}\|\mathbf{e}^{(t+1)} - \mathbf{e}^{(t)}\|_2 \leq \frac{\mu C_{\mathbf{e}}}{1 - \delta_k} \left(\frac{\theta_k}{1 - \delta_k}\right)^t \quad (244)$$

$$\zeta_{I_t}^{(t+1)} = \zeta_{I_t}^{(t-1)} \quad (245)$$

$$\|\zeta_{(I_t \cup I_{t+1})^c}^{(t+1)} - \zeta_{(I_t \cup I_{t+1})^c}^{(t)}\|_{\infty} \leq \mu C_{\mathbf{e}} \left(\frac{\theta_k}{1 - \delta_k}\right)^{t+1} \sqrt{\frac{m}{k}} \quad (246)$$

To prove this we first note that Eq 245 holds, as by definition $\zeta_{I_t}^{(t-1)} = \zeta_{I_t}^{(t+1)} = \mathbf{a}_t$. We then note that

$$\frac{1}{\sqrt{m}}\|\zeta_{I_0^c}^{(0)}\|_{\infty} = \min |\zeta_{I_0^c}^{(0)}| \leq \mu C_{\mathbf{e}} \sqrt{\frac{m}{k}} \quad (247)$$

We further note that $\zeta_{I_0}^{(0)} \in \partial p(\mathbf{w} = \mathbf{e}_{I_0}^{(0)})$. Therefore by lemma 15 we see that

$$\frac{1}{\sqrt{m}}\|\mathbf{w}_0 - \mathbf{v}_{I_0}^{(0)}\|_2 \leq \frac{\|\zeta_{I_0}^{(0)}\|_2}{\sigma_{max}^2(\mathbf{X}_{I_0})} \leq \frac{\mu C_{\mathbf{e}}}{1 - \delta_k} \quad (248)$$

We now note that at $t = 0$, by construction

$$\begin{cases} \zeta_{I_0}^{(0)} = \mathbf{0} \\ \zeta_{I_0}^{(1)} = \zeta_{I_0}^{(0)} + \mathbf{X}_{I_0^c}^T \mathbf{A}_{I_0}(\mathbf{w}_0 - \mathbf{e}_{I_0}^{(0)}) \end{cases} \quad (249)$$

and that $\mathbf{p}_1 = \zeta_{I_1}^{(0)}$. From this we see that

$$\frac{1}{\sqrt{m}}\|\mathbf{X}_{I_0^c}^T \mathbf{X}_{I_0}(\mathbf{w}_0 - \mathbf{e}_{I_0}^{(0)})\|_2 \leq \frac{\theta_k}{\sqrt{m}}\|\mathbf{w}_0 - \mathbf{e}_{I_0}^{(0)}\|_2 \leq \frac{\theta_k \mu C_{\mathbf{e}}}{1 - \delta_k} \quad (250)$$

From which we obtain

$$\|\mathbf{X}_{(I_0 \cup I_1)^c}^T \mathbf{X} \mathbf{X}_{I_0}(\mathbf{w}_0 - \mathbf{e}_{I_0}^{(0)})\|_{\infty} \leq \min |\mathbf{X}_{I_1}^T \mathbf{X}_{I_0}(\mathbf{w}_0 - \mathbf{e}_{I_0}^{(0)})| \leq \frac{\theta_k \mu C_{\mathbf{e}}}{1 - \delta_k} \sqrt{\frac{m}{k}} \quad (251)$$

Finally noting that for $t = 0$ we have that $\|\mathbf{e}^{(1)} - \mathbf{e}^{(0)}\|_2 = \|\mathbf{w}_0 - \mathbf{e}_{I_0}^{(0)}\|_2$. From this we see that the base case of the induction is satisfied.

We now assume that equations 244 - 246 hold for all $t' \leq t$ we now prove that they will hold for $t + 1$. We consider the optimization 228 at step t , we also showed above that $\zeta^{(t)} \in \partial p(\mathbf{e}^{(t)})$. From this we see that

$$\zeta_{I_t}^{(t)} - \mathbf{a}_t \in \partial p_t(\mathbf{e}_{I_t}^{(t)}) \quad (252)$$

From this we see that

$$\mathbf{X}_{I_t}^T \mathbf{X}_{I_{t-1}}(\mathbf{w}_{t-1} - \mathbf{e}_{I_{t-1}}^{(t-1)}) \in \partial p_t(\mathbf{e}_{I_t}^{(t)}) \quad (253)$$

By Lemma 15 below we see that

$$\begin{aligned} \frac{1}{\sqrt{m}}\|\mathbf{w}_t - \mathbf{e}_{I_t}^{(t)}\|_2 &\leq \frac{1}{(1 - \delta_k)\sqrt{m}}\|\mathbf{X}_{I_t}^T \mathbf{X}_{I_{t-1}}(\mathbf{w}_{t-1} - \mathbf{e}_{I_{t-1}}^{(t-1)})\|_2 \\ &\leq \frac{\theta_k}{(1 - \delta_k)\sqrt{m}}\|\mathbf{w}_{t-1} - \mathbf{e}_{I_{t-1}}^{(t-1)}\|_2 \\ &= \frac{\theta_k}{(1 - \delta_k)\sqrt{m}}\|\mathbf{e}^{(t)} - \mathbf{e}^{(t-1)}\|_2 \\ &\leq \frac{\theta_k}{1 - \delta_k} \frac{\mu C_{\mathbf{e}}}{1 - \delta_k} \left(\frac{\theta_k}{1 - \delta_k}\right)^{t-1} = \frac{\mu C_{\mathbf{e}}}{1 - \delta_k} \left(\frac{\theta_k}{1 - \delta_k}\right)^t \end{aligned} \quad (254)$$

This proves Eq 245. We also see that

$$\frac{1}{\sqrt{m}} \|\mathbf{X}_{I_{t+1}}^T \mathbf{X}_{I_t} (\mathbf{w}_t - \mathbf{e}_{I_t}^{(t)})\|_2 \leq \theta_k \|\mathbf{w}_t - \mathbf{e}_{I_t}^{(t)}\|_2 \leq \frac{\theta_k \mu C_e}{1 - \delta_k} \left(\frac{\theta_k}{1 - \delta_k} \right)^t \quad (255)$$

Therefore

$$\begin{aligned} \|\zeta_{(I_t \cup I_{t+1})^c}^{(t+1)} - \zeta_{(I_t \cup I_{t+1})^c}^{(t)}\|_\infty &= \|\mathbf{X}_{(I_t \cup I_{t+1})^c}^T \mathbf{X}_{I_t} (\mathbf{w}_t - \mathbf{e}_{I_t}^{(t)})\|_\infty \\ &\leq \min |\mathbf{X}_{I_{t+1}}^T \mathbf{X}_{I_t} (\mathbf{w}_t - \mathbf{e}_{I_t}^{(t)})| \leq \sqrt{\frac{1}{k}} \|\mathbf{X}_{I_{t+1}}^T \mathbf{X}_{I_t} (\mathbf{w}_t - \mathbf{e}_{I_t}^{(t)})\|_2 \\ &\leq \sqrt{\frac{m}{k}} \mu C_e \left(\frac{\theta_k}{1 - \delta_k} \right)^{t+1} \end{aligned} \quad (256)$$

This proves Eq 246.

We now see in eq 244 that if $\theta_k + \delta_k < 1$, then the sequence of $\mathbf{e}^{(t)}$ is absolutely convergent. Furthermore, from 245 and 246, in addition to the relation

$$\frac{1}{\sqrt{m}} \|\zeta_{I_{t+1}}^{(t+1)} - \zeta_{I_{t+1}}^t\|_2 = \|\mathbf{X}_{I_{t+1}}^T \mathbf{X}_{I_t} (\mathbf{w}_t - \mathbf{e}_{I_t}^{(t)})\|_2 \leq \mu C_e \left(\frac{\theta_k}{1 - \delta_k} \right)^{t+1} \quad (257)$$

From this we obtain

$$\begin{aligned} \frac{1}{\sqrt{m}} \|\zeta^{(t+1)} - \zeta^{(t)}\|_2 &= \sqrt{\|\zeta_{I_t}^{(t+1)} - \zeta_{I_t}^{(t)}\|_2 + \|\zeta_{I_{t+1}}^{(t+1)} - \zeta_{I_{t+1}}^{(t)}\|_2 + \|\zeta_{(I_t \cup I_{t+1})^c}^{(t+1)} - \zeta_{(I_t \cup I_{t+1})^c}^{(t)}\|_2} \\ &= \sqrt{\|\zeta_{I_t}^{(t-1)} - \zeta_{I_t}^{(t)}\|_2 + \|\zeta_{I_{t+1}}^{(t+1)} - \zeta_{I_{t+1}}^{(t)}\|_2 + \|\zeta_{(I_t \cup I_{t+1})^c}^{(t+1)} - \zeta_{(I_t \cup I_{t+1})^c}^{(t)}\|_2} \\ &\leq \sqrt{m} \sqrt{\mu^2 C_e^2 \left(\frac{\theta_k}{1 - \delta_k} \right)^{2t} + \mu^2 C_e^2 \left(\frac{\theta_k}{1 - \delta_k} \right)^{2t+2} + \mu^2 C_e^2 \left(\frac{m}{k} - 1 \right) \left(\frac{\theta_k}{1 - \delta_k} \right)^{2t+2}} \end{aligned} \quad (258)$$

As such we see that the sequence $\zeta^{(t)}$ is absolutely convergent. We denote the limits of $\zeta^{(t)}$ and $\mathbf{v}^{(t)}$ as $\zeta^{(\infty)}$ and $\mathbf{e}^{(\infty)}$ respectively.

We have that

$$\begin{aligned} \frac{1}{\sqrt{m}} \|\mathbf{e}^{(0)} - \mathbf{e}^{(\infty)}\|_2 &\leq \sum_{t=0}^{\infty} \|\mathbf{e}^{(t+1)} - \mathbf{e}^{(t)}\|_2 \leq \sum_{t=0}^{\infty} \frac{\mu C_e}{1 - \delta_k} \left(\frac{\theta_k}{1 - \delta_k} \right)^t \\ &\quad - \frac{\mu C_e}{1 - \delta_k - \theta_k} \end{aligned} \quad (259)$$

Finally we show that $\|\zeta^{(\infty)}\|_\infty$ is bounded as well. We consider an index i and denote by $t_1 < t_2 < \dots$ as the iterations of t for which $i \in I_t$. In the case that $i \notin I_0$ by equation 245 we see that

$$\zeta_i^{(\infty)} - \zeta_i^{(0)} = \sum_{t=0}^{\infty} \zeta_i^{(t+1)} - \zeta_i^{(t)} = \sum_{t|i \in (I_t \cup I_{t+1})^c} \zeta_i^{(t+1)} - \zeta_i^{(t)} \quad (260)$$

As such we obtain

$$|\zeta_i^{(\infty)}| \leq \zeta_i^{(0)} + \sum_{t|i \in (I_t \cup I_{t+1})^c} |\zeta_i^{(t+1)} - \zeta_i^{(t)}| \leq \mu C_e \sqrt{\frac{m}{k}} + \mu C_e \sum_t \left(\frac{\theta_k}{1 - \delta_k} \right)^t \leq \mu C_e \left(\sqrt{\frac{m}{k}} + \frac{\theta_k}{1 - \delta_k - \theta_k} \right) \quad (261)$$

For any $i \in I_0$ we have

$$\zeta_i^{(\infty)} - \zeta_i^{(1)} = \sum_{t \geq 1 | i \in (I_t \cup I_{t+1})^c} \zeta_i^{(t+1)} - \zeta_i^{(t)} \quad (262)$$

By recalling that $\zeta_{I_0}^{(1)} = 0$ we obtain

$$|\zeta_i^{(\infty)}| \leq \mu C_e \sum_t \left(\frac{\theta_k}{1 - \delta_k} \right)^t = \frac{\mu C_e \theta_k^2}{(1 - \delta_k - \theta_k)(1 - \delta_k)} \quad (263)$$

Combining the results in total we obtain

$$\|\zeta^{(\infty)}\|_\infty \leq \mu C_e \left(\sqrt{\frac{m}{k}} + \frac{\theta_k}{1 - \delta_k - \theta_k} \right) \quad (264)$$

Finally we note that because for each t , $\zeta^{(t)} \in \partial p(\mathbf{e}^{(t)})$ we see that $\zeta^{(\infty)} \in \partial p(\mathbf{e}^{(\infty)})$

□

C.6 Universality of Generalization Error and Test Functions for ℓ_1 regularization

We first demonstrate the universality of the Generalization error. We demonstrate that the 2-norm of the solution vector of the ℓ_1 regularized case is asymptotically equivalent to the case of the elastic net regularized case for small values of ℓ_2^2 regularization. We have already demonstrated that the generalization error for the elastic net case is universal, by showing that the ℓ_1 is asymptotically equivalent we prove universality for that case as well.

Lemma 13. Denote by $\hat{\boldsymbol{\theta}}^{\lambda, \epsilon}$ as the optimal point of

$$P_{\lambda, \epsilon} = \min_{\boldsymbol{\theta}} \frac{1}{2n} \|\mathbf{y} - \mathbf{X}\boldsymbol{\theta}\|_2^2 + \lambda \|\boldsymbol{\theta}\|_1 + \frac{\epsilon}{2} \|\boldsymbol{\theta}\|_2^2 \quad (265)$$

Under the conditions assumed in theorem 6, for each $\eta > 0$, there exists ϵ, ρ such that for $0 < \epsilon < \eta$ and $|\rho| < \eta$, such that

$$\mathbb{P} \left(\frac{\|\hat{\boldsymbol{\theta}}^{\lambda+\rho, \epsilon} - \hat{\boldsymbol{\theta}}^{\lambda, 0}\|_2^2}{n} > \eta \right) \rightarrow 0 \quad (266)$$

Proof. We first note that with a high degree of probability we have that

$$M_0 + \theta < \frac{l(1 - \delta_l(\mathbf{X}))}{2n} \quad (267)$$

in which M_0 is the effective sparsity given in equation 33, $\theta > 0$ is a fixed number and $l < n$ is natural number such that $\delta_l < 1$. From this we see that $(1 - \delta_l) > 2(M_0 + \theta)$ and $l/n > 2(M_0 + \theta)$. We let $0 < \alpha < \min(4M_0, 2\theta)$, and let $K = M_0 + \theta - \alpha/2$ and $k = \frac{l}{nK} - 1$. We note that $K > M_0$ and

$$k = \frac{l}{n(M_0 + \theta - \alpha/2)} - 1 > \frac{l}{n(M_0 + \theta)} - 1 > 1 \quad (268)$$

Furthermore,

$$K = M_0 + \theta - \frac{\alpha}{2} \leq \frac{l(1 - \delta_l(\mathbf{X}))}{2n} \leq \frac{l}{n} \left[\frac{1 - \alpha - \delta_l(\mathbf{X})}{2 - \alpha} + \frac{\alpha}{2} \right] - \frac{\alpha}{2} \leq \frac{l}{n} \left[\frac{1 - \alpha - \delta_l(\mathbf{X})}{2 - \alpha} \right] \quad (269)$$

from which we can see that

$$\alpha \leq \frac{k - 1 - (k + 1)\delta_l(\mathbf{X})}{k} \quad (270)$$

We define a function $M_{r, \psi}$ in which r is the regularization function and ψ are both functions given by:

$$M_{r, \psi}(\beta, q, \xi, r) = \mathbb{E} \left(\psi \left(\hat{\boldsymbol{\theta}}_r(\beta, q, \xi, r) \right) \right) \quad (271)$$

In which $\hat{\boldsymbol{\theta}}_r(\beta, q, \xi, t)$ is the optimal value of \tilde{P}_3 given in (138) with regularization function r . We now define

$$M^{\lambda, \epsilon} = M_{\lambda|\mathbf{x}| + \frac{\epsilon}{2}x^2, x^2} \quad N^{\lambda, \epsilon} = M_{\lambda|\mathbf{x}| + \frac{\epsilon}{2}, |\mathbf{x}|} \quad (272)$$

Now let $\delta > 0$. We can see that there exist value ρ, ϵ such that $0 < \epsilon < \delta$, $|\rho| < \delta$ such that $0 < N^{\lambda+\rho, \epsilon} - N^{\lambda, 0} < \delta$. Then let $\mu > 0$ be defined such that

$$2\mu < N^{\lambda+\rho, \epsilon} - N^{\lambda, 0} \quad (273)$$

We define $\mathbf{h} = \hat{\boldsymbol{\theta}}^{\lambda, 0} - \hat{\boldsymbol{\theta}}^{\lambda+\rho, \epsilon}$. We denote the objective function in Eq 265 as $P_{\lambda, \epsilon}(\boldsymbol{\theta})$. We have that

$$\begin{aligned}
 P_{\lambda+\rho,\epsilon}(\hat{\boldsymbol{\theta}}^{\lambda,0}) &= P_{\lambda,0}(\hat{\boldsymbol{\theta}}^{\lambda,0}) + \frac{1}{n} \left(\frac{\epsilon}{2} \|\hat{\boldsymbol{\theta}}^{\lambda,0}\|_2^2 - \rho \|\hat{\boldsymbol{\theta}}^{\lambda,0}\|_1 \right) \\
 &\leq P_{\lambda,0}(\hat{\boldsymbol{\theta}}^{\lambda+\rho,\epsilon}) + \frac{1}{n} \left(\frac{\epsilon}{2} \|\hat{\boldsymbol{\theta}}^{\lambda,0}\|_2^2 - \rho \|\hat{\boldsymbol{\theta}}^{\lambda,0}\|_1 \right) \\
 &= P_{\lambda+\rho,\epsilon}(\hat{\boldsymbol{\theta}}^{\lambda+\rho,\epsilon}) + \frac{1}{n} \left(\frac{\epsilon}{2} \|\hat{\boldsymbol{\theta}}^{\lambda,0}\|_2^2 - \rho \|\hat{\boldsymbol{\theta}}^{\lambda,0}\|_1 - \frac{\epsilon}{2} \|\hat{\boldsymbol{\theta}}^{\lambda+\rho,\epsilon}\|_2^2 - \rho \|\hat{\boldsymbol{\theta}}^{\lambda+\rho,\epsilon}\|_1 \right) \\
 &\leq P_{\lambda+\rho,\epsilon}(\hat{\boldsymbol{\theta}}^{\lambda+\rho,\epsilon}) + \frac{\epsilon}{2} \frac{\|\mathbf{h}\|_2^2}{n} + \epsilon \frac{\|\mathbf{h}\|_2}{\sqrt{n}} + \frac{\rho}{n} \left(\|\hat{\boldsymbol{\theta}}^{\lambda,0}\|_1 - \|\hat{\boldsymbol{\theta}}^{\lambda+\rho,\epsilon}\|_1 \right)
 \end{aligned} \tag{274}$$

From theorem 4 and part one of Theorem 6 we know that

$$\frac{\|\hat{\boldsymbol{\theta}}^{\lambda+\rho,\epsilon}\|_2^2}{n} \xrightarrow{P} M^{\lambda+\rho,\epsilon} \quad \frac{\|\hat{\boldsymbol{\theta}}^{\lambda+\rho,\epsilon}\|_1}{n} \xrightarrow{P} N^{\lambda+\rho,\epsilon} \quad \frac{\|\hat{\boldsymbol{\theta}}^{\lambda,0}\|_1}{n} \xrightarrow{P} M^{\lambda,0} \tag{275}$$

Choosing a value of $M > \sqrt{M^{\lambda+\rho,\epsilon}}$, we obtain

$$P_{\lambda+\rho,\epsilon}(\hat{\boldsymbol{\theta}}^{\lambda,0}) \leq P_{\lambda+\rho,\epsilon}(\hat{\boldsymbol{\theta}}^{\lambda+\rho,\epsilon}) + \frac{\epsilon}{2} \frac{\|\mathbf{h}\|_2^2}{n} + M \epsilon \frac{\|\mathbf{h}\|_2}{\sqrt{n}} + \rho \delta \tag{276}$$

We now define the following index sets

$$S = \{k \mid |\hat{\boldsymbol{\theta}}_k^{\lambda+\rho,\epsilon}| \geq \mu\} \quad L = \{k \mid 0 < |\hat{\boldsymbol{\theta}}_k^{\lambda+\rho,\epsilon}| < \mu\} \tag{277}$$

We also define

$$K_\mu^{\lambda,\epsilon} = M_{\lambda|x|+\epsilon x^2/2, \chi_{\mathbb{R} \setminus (-\mu, \mu)}} \tag{278}$$

In which χ_A is the indicator function on the set A . By theorem 4 we have that

$$\frac{|S|}{n} \xrightarrow{P} K_\mu^{\lambda+\rho,\epsilon} \tag{279}$$

we also see that

$$\lim_{(\mu,\rho,\epsilon) \rightarrow 0} K_\mu^{\lambda+\rho,\epsilon} = M_0 \tag{280}$$

Therefore, for small values of δ we know that $K_\mu^{\lambda+\rho,\epsilon} < K$ and as such with high probability

$$\frac{|S|}{n} < K \tag{281}$$

We also know from equation 275 that with high probability

$$\frac{\|\hat{\boldsymbol{\theta}}^{\lambda+\rho,\epsilon}\|_1}{n} - \frac{\|\hat{\boldsymbol{\theta}}^{\lambda,0}\|_1}{n} > 2\mu \tag{282}$$

This can equivalently be expressed as

$$\begin{aligned}
 &\frac{\|\hat{\boldsymbol{\theta}}_S^{\lambda+\rho,\epsilon}\|_1}{n} + \frac{\|\hat{\boldsymbol{\theta}}_L^{\lambda+\rho,\epsilon}\|_1}{n} > \frac{\|\hat{\boldsymbol{\theta}}_S^{\lambda+\rho,\epsilon} + \mathbf{h}_S\|_1}{n} + \frac{\|\hat{\boldsymbol{\theta}}_L^{\lambda+\rho,\epsilon} + \mathbf{h}_L\|_1}{n} + \frac{\|\mathbf{h}_{S \cup L}^c\|_1}{n} + 2\mu \\
 &\geq \frac{\|\hat{\boldsymbol{\theta}}_S^{\lambda+\rho,\epsilon}\|_1 - \|\mathbf{h}_S\|_1}{n} + \frac{\|\mathbf{h}_L\|_1 - \|\hat{\boldsymbol{\theta}}_L^{\lambda+\rho,\epsilon}\|_1}{n} + \frac{\|\hat{\boldsymbol{\theta}}_L^{\lambda+\rho,\epsilon} + \mathbf{h}_L\|_1}{n} + \frac{\|\mathbf{h}_{S \cup L}^c\|_1}{n} + 2\mu
 \end{aligned} \tag{283}$$

By definition $\|\hat{\boldsymbol{\theta}}_L^{\lambda+\rho,\epsilon}\|_1 \leq \mu$. As such with high probability we obtain

$$\|\mathbf{h}_S\|_1 \geq \|\mathbf{h}_{S^c}\|_1 \tag{284}$$

We now define $\mathbf{z} = \mathbf{y} - \mathbf{X}\hat{\boldsymbol{\theta}}^{\lambda+\rho,\epsilon}$. We wish to decompose the vector \mathbf{h}_{S^c} into block T_1, T_2, \dots . We let \mathbf{h}_{T_1} be the $k|S|$ elements of \mathbf{h}_{S^c} with largest absolute value, \mathbf{h}_{T_2} are the next $k|S|$ largest absolute values and so on. Let $U = S \cup T_1$. With that we have

$$nP_{\lambda+\rho,\epsilon}(\hat{\boldsymbol{\theta}}^{\lambda,0}) = \frac{1}{2} \|\mathbf{z} - \mathbf{X}\mathbf{h}\|_2^2 + \frac{(\lambda + \rho)}{\gamma} \|\hat{\boldsymbol{\theta}}^{\lambda+\rho,\epsilon} + \mathbf{h}\|_1 + \frac{\epsilon}{2\gamma} \|\hat{\boldsymbol{\theta}}^{\lambda+\rho,\epsilon} + \mathbf{h}\|_2^2 \tag{285}$$

We note that $\hat{\boldsymbol{\theta}}^{\lambda,0} = \hat{\boldsymbol{\theta}}^{\lambda+\rho,\epsilon} + \mathbf{h}$ which is the minimal point of the function $P_{\lambda,0}(\boldsymbol{\theta})$. As such we have

$$\mathbf{X}^T(\mathbf{z} - \mathbf{X}\mathbf{h}) = \mathbf{X}^T(\mathbf{y} - \mathbf{X}\hat{\boldsymbol{\theta}}^{\lambda,0}) \in \lambda\partial\|\hat{\boldsymbol{\theta}}^{\lambda,0}\|_1 \quad (286)$$

Therefore

$$\|\mathbf{X}_{U^c}^T(\mathbf{z} - \mathbf{X}\mathbf{h})\|_\infty \leq \lambda \Rightarrow -\mathbf{h}_{U^c}^T \mathbf{X}_{U^c}^T(\mathbf{z} - \mathbf{X}\mathbf{h}) \geq -\lambda\|\mathbf{h}_{U^c}\|_1 \quad (287)$$

From which we obtain

$$-\mathbf{h}_{U^c}^T \mathbf{X}_{U^c}^T(\mathbf{z} - \mathbf{X}_U \mathbf{h}_U) \geq -\lambda\|\mathbf{h}_{U^c}\|_1 - \|\mathbf{X}_{U^c} \mathbf{h}_{U^c}\|_2^2 \quad (288)$$

and as such

$$\frac{1}{2}\|\mathbf{z} - \mathbf{X}\mathbf{h}\|_2^2 = \frac{1}{2}\|\mathbf{z} - \mathbf{X}_U \mathbf{h}_U\|_2^2 - \mathbf{h}_{U^c}^T \mathbf{X}_{U^c}^T + \frac{1}{2}\|\mathbf{A}_{U^c} \mathbf{h}_{U^c}\|_2^2 \geq \frac{1}{2}\|\mathbf{z} - \mathbf{X}_U \mathbf{h}_U\|_2^2 - \lambda\|\mathbf{h}_{U^c}\|_1 + \frac{1}{2}\|\mathbf{X}_{U^c} \mathbf{h}_{U^c}\|_2^2 \quad (289)$$

From which we obtain

$$\begin{aligned} mP_{\lambda+\rho,\epsilon}(\hat{\boldsymbol{\theta}}^{\lambda,0}) &\geq \frac{1}{2}\|\mathbf{z} - \mathbf{X}_U \mathbf{h}_U\|_2^2 - \lambda\|\mathbf{h}_{U^c}\|_1 + \frac{1}{2}\|\mathbf{X}_{U^c} \mathbf{h}_{U^c}\|_2^2 + (\lambda + \rho)\|\hat{\boldsymbol{\theta}}_U^{\lambda+\rho,\epsilon} + \mathbf{h}_U\|_1 + \\ &\quad (\lambda + \rho)\|\hat{\boldsymbol{\theta}}_{U^c}^{\lambda+\rho,\epsilon} + \mathbf{h}_{U^c}\|_1 + \frac{\epsilon}{2}\|\hat{\boldsymbol{\theta}}_U^{\lambda+\rho,\epsilon} + \mathbf{h}_U\|_2^2 + \frac{\epsilon}{2}\|\hat{\boldsymbol{\theta}}_{U^c}^{\lambda+\rho,\epsilon} + \mathbf{h}_{U^c}\|_2^2 \end{aligned} \quad (290)$$

We note that $\mathbf{w} = 0$ is the minimum point of the function

$$\frac{1}{2}\|\mathbf{z} - \mathbf{X}_U \mathbf{w}\|_2^2 + (\lambda + \rho)\|\hat{\boldsymbol{\theta}}^{\lambda+\rho,\epsilon} + \mathbf{w}\|_1 + \frac{\epsilon}{2}\|\hat{\boldsymbol{\theta}}_U^{\lambda+\rho,\epsilon} + \mathbf{w}\|_2^2 \quad (291)$$

Therefore from lemma 16 we get that

$$\begin{aligned} &\frac{1}{2}\|\mathbf{z} - \mathbf{X}_U \mathbf{h}_U\|_2^2 + (\lambda + \rho)\|\hat{\boldsymbol{\theta}}^{\lambda+\rho,\epsilon} + \mathbf{h}_U\|_1 + \frac{\epsilon}{2}\|\hat{\boldsymbol{\theta}}_U^{\lambda+\rho,\epsilon} + \mathbf{h}_U\|_2^2 \\ &\geq \frac{\sigma_{max}^2(\mathbf{X}_U)}{2}\|\mathbf{h}_U\|_2^2 + \frac{1}{2}\|\mathbf{z}\|_2^2 + (\lambda + \rho)\|\hat{\boldsymbol{\theta}}^{\lambda+\rho,\epsilon}\|_1 + \frac{\epsilon}{2}\|\hat{\boldsymbol{\theta}}^{\lambda+\rho,\epsilon}\|_2^2 \end{aligned} \quad (292)$$

Substituting this in above we get that

$$\begin{aligned} &nP_{\lambda+\rho,\epsilon}(\hat{\boldsymbol{\theta}}^{\lambda,0}) - nP_{\lambda+\rho,\epsilon}(\hat{\boldsymbol{\theta}}^{\lambda+\rho,\epsilon}) \\ &\geq \frac{\sigma_{max}^2(\mathbf{X}_U)}{2}\|\mathbf{h}_U\|_2^2 - \lambda\|\mathbf{h}_{U^c}\|_1 - \frac{1}{2}\|\mathbf{A}_{U^c} \mathbf{h}_{U^c}\|_2^2 + (\lambda + \rho)\|\hat{\boldsymbol{\theta}}_{U^c}^{\lambda+\rho,\epsilon} + \mathbf{h}_{U^c}\|_1 - (\lambda + \rho)\|\hat{\boldsymbol{\theta}}_{U^c}^{\lambda+\rho,\epsilon}\|_1 \\ &\quad - \frac{\epsilon}{2}\|\hat{\boldsymbol{\theta}}_{U^c}^{\lambda+\rho,\epsilon}\|_2^2 + \frac{\epsilon}{2}\|\hat{\boldsymbol{\theta}}_{U^c}^{\lambda+\rho,\epsilon} + \mathbf{h}_{U^c}\|_2^2 \\ &\geq \frac{\sigma_{max}^2(\mathbf{X}_U)}{2}\|\mathbf{h}_U\|_2^2 + \rho\|\mathbf{h}_{U^c}\|_1 - \frac{1}{2}\|\mathbf{X}_{U^c} \mathbf{h}_{U^c}\|_2^2 - 2(\lambda + \rho)\|\hat{\boldsymbol{\theta}}_{U^c}^{\lambda+\rho,\epsilon}\|_1 - 2\|\hat{\boldsymbol{\theta}}_{U^c}^{\lambda+\rho,\epsilon}\|_2\|\mathbf{h}_{U^c}\|_2 \\ &\geq \frac{\sigma_{max}^2(\mathbf{X}_U)}{2}\|\mathbf{h}_U\|_2^2 + \delta\sqrt{n}\|\mathbf{h}_U\|_2 - \frac{1}{2}\|\mathbf{X}_{U^c} \mathbf{h}_{U^c}\|_2^2 - 2(\lambda + \rho)n\mu - 2\sqrt{n}\mu\|\mathbf{h}_{U^c}\|_2 \end{aligned} \quad (293)$$

Where we have made use of the fact that

$$\rho\|\mathbf{h}_{U^c}\|_1 \geq -\delta\|\mathbf{h}_{U^c}\|_1 \geq -\delta\|\mathbf{h}_U\|_1 \geq -\delta\sqrt{n}\|\mathbf{h}_U\|_2 \quad (294)$$

and in (Candes et al., 2006)(equation 11) it is proven that

$$\|\mathbf{h}_{U^c}\|_1 \leq \frac{|S|}{|L|}\|\mathbf{h}_U\|_2^2 = \frac{1}{k}\|\mathbf{h}_U\|_2^2 \quad (295)$$

Also in (Candes et al., 2006) (equation 12) it is shown that

$$\|\mathbf{X}_{U^c} \mathbf{h}_{U^c}\|_2 \leq \sqrt{1 + \delta_{k|S|}(\mathbf{X})} \sqrt{\frac{|S|}{|T|}}\|\mathbf{h}_U\|_2 = \sqrt{\frac{1 + \delta_{k|S|}(\mathbf{X})}{k}}\|\mathbf{h}_U\|_2 \quad (296)$$

As such

$$\begin{aligned} & nP_{\lambda+\rho,\epsilon}(\hat{\boldsymbol{\theta}}^{\lambda,0}) - nP_{\lambda+\rho,\epsilon}(\hat{\boldsymbol{\theta}}^{\lambda+\rho,\epsilon}) \\ & \geq \frac{1 - \delta_{(1+k)|S|}(\mathbf{X}) - \frac{1 - \delta_{k|S|}(\mathbf{X})}{k}}{2} \|\mathbf{h}_U\|_2^2 - \left(1 - \frac{1}{\sqrt{k}}\right) \delta \sqrt{n} \|\mathbf{h}_U\|_2 - (\lambda + \rho)n\delta \end{aligned} \quad (297)$$

Noting that $|S| < Kn$. By equation 270,

$$\alpha_1 = 1 - \delta_{(1+k)|S|}(\mathbf{X}) - \frac{1 + \delta_{k|S|}(\mathbf{X})}{k} \geq 1 - \delta_{n(1+k)K}(\mathbf{X}) - \frac{1 + \delta_{nkK}(\mathbf{X})}{k} \geq 1 - \delta_l(\mathbf{X}) - \frac{1 + \delta_l(\mathbf{X})}{k} \geq \alpha \quad (298)$$

from which we obtain

$$nP_{\lambda+\rho,\epsilon}(\hat{\boldsymbol{\theta}}^{\lambda,0}) - nP_{\lambda+\rho,\epsilon}(\hat{\boldsymbol{\theta}}^{\lambda+\rho,\epsilon}) \geq \frac{\alpha}{2} \|\mathbf{h}_U\|_2^2 - \left(1 - \frac{1}{\sqrt{k}}\right) \delta \sqrt{n} \|\mathbf{h}_U\|_2 - (\lambda + \delta)n\delta \quad (299)$$

combining this with equation X above

$$\frac{\alpha}{2} \|\mathbf{h}_U\|_2^2 - \left(1 - \frac{1}{\sqrt{k}}\right) \delta \sqrt{n} \|\mathbf{h}_U\|_2 - (\lambda + \rho)n\delta \leq \frac{\delta}{2} \|\mathbf{h}_U\|_2^2 + M\delta \sqrt{n} \|\mathbf{h}_U\|_2 + n\delta^2 \quad (300)$$

We see that

$$\|\mathbf{h}_U\|_2^2 \leq \left(1 + \frac{1}{k}\right) \|\mathbf{h}\|_2^2 \quad (301)$$

From this we see that

$$\frac{\alpha}{2(1 + \frac{1}{k})} \|\mathbf{h}\|_2^2 - \frac{1 + \frac{1}{\sqrt{k}}}{\sqrt{1 + \frac{1}{k}}} \delta \sqrt{n} \|\mathbf{h}\|_2 - (\lambda + \delta)n\delta \leq \frac{\delta}{2} \|\mathbf{h}\|_2^2 + M\delta \sqrt{n} \|\mathbf{h}\|_2 + n\delta^2 \quad (302)$$

Since we know that $k > 1$, we can see that for any choice of $\eta > 0$ the value of δ can be made sufficiently small to ensure that equation 302 implies that the lemma holds. \square

We can now show the universality of the generalization error, we note that for the two cases, the term:

$$\begin{aligned} (\boldsymbol{\theta}^{\lambda,0} - \boldsymbol{\theta}^*)^T \mathbf{R}(\boldsymbol{\theta}^{\lambda,0} - \boldsymbol{\theta}^*) &= (\boldsymbol{\theta}^{\lambda,0} - \boldsymbol{\theta}^{\lambda+\rho,\epsilon} + \boldsymbol{\theta}^{\lambda+\rho,\epsilon} - \boldsymbol{\theta}^*)^T \mathbf{R}(\boldsymbol{\theta}^{\lambda,0} - \boldsymbol{\theta}^{\lambda+\rho,\epsilon} + \boldsymbol{\theta}^{\lambda+\rho,\epsilon} - \boldsymbol{\theta}^*) \\ &= (\boldsymbol{\theta}^{\lambda,0} - \boldsymbol{\theta}^{\lambda+\rho,\epsilon})^T \mathbf{R}(\boldsymbol{\theta}^{\lambda,0} - \boldsymbol{\theta}^{\lambda+\rho,\epsilon}) + 2(\boldsymbol{\theta}^{\lambda,0} - \boldsymbol{\theta}^{\lambda+\rho,\epsilon})^T \mathbf{R}(\boldsymbol{\theta}^{\lambda+\rho,\epsilon} - \boldsymbol{\theta}^*) \\ &\quad + (\boldsymbol{\theta}^{\lambda+\rho,\epsilon} - \boldsymbol{\theta}^*)^T \mathbf{R}(\boldsymbol{\theta}^{\lambda+\rho,\epsilon} - \boldsymbol{\theta}^*) \\ &\leq \|\boldsymbol{\theta}^{\lambda,0} - \boldsymbol{\theta}^{\lambda+\rho,\epsilon}\|_2^2 \|\mathbf{R}\|_2 + 2\|(\boldsymbol{\theta}^{\lambda,0} - \boldsymbol{\theta}^{\lambda+\rho,\epsilon})^T\|_2 \|\mathbf{R}\|_2 \|\boldsymbol{\theta}^{\lambda+\rho,\epsilon} - \boldsymbol{\theta}^*\|_2 \\ &\quad + (\boldsymbol{\theta}^{\lambda+\rho,\epsilon} - \boldsymbol{\theta}^*)^T \mathbf{R}(\boldsymbol{\theta}^{\lambda+\rho,\epsilon} - \boldsymbol{\theta}^*) \\ &= (\boldsymbol{\theta}^{\lambda+\rho,\epsilon} - \boldsymbol{\theta}^*)^T \mathbf{R}(\boldsymbol{\theta}^{\lambda+\rho,\epsilon} - \boldsymbol{\theta}^*) \end{aligned} \quad (303)$$

Where the final step is by the lemma above showing the asymptotic equivalence of the two norm. By symmetry the argument may be repeated to show that

$$(\boldsymbol{\theta}^{\lambda+\rho,\epsilon} - \boldsymbol{\theta}^*)^T \mathbf{R}(\boldsymbol{\theta}^{\lambda+\rho,\epsilon} - \boldsymbol{\theta}^*) \leq (\boldsymbol{\theta}^{\lambda,0} - \boldsymbol{\theta}^*)^T \mathbf{R}(\boldsymbol{\theta}^{\lambda,0} - \boldsymbol{\theta}^*) \quad (304)$$

In the asymptotic limit. This fact, in conjunction with Theorem 10 proves the universality of the generalization error. Finally we show that the universality of the test functions $h(\boldsymbol{\theta})$.

Lemma 14. For a function h have that with high probability that

$$\lim_{n \rightarrow \infty} \left| \frac{h(\hat{\boldsymbol{\theta}}^{\lambda,0})}{n} - M_{\lambda|x|,h} \right| = 0 \quad (305)$$

in which M is the function given in (271).

Proof. We denote $\hat{\theta}^{\lambda,\epsilon}$ as the minimal solution of

$$P_{1\lambda,\epsilon} = \frac{1}{2n} \|\mathbf{y} - \mathbf{X}\theta\|_2^2 + \lambda \|\theta\|_1 + \frac{\epsilon}{2} \|\theta\|_2^2 \quad (306)$$

Because of the results of the lemma C.6 above we know that

$$\begin{aligned} \frac{h(\hat{\theta}^{\lambda,0})}{n} - M_{\lambda|x|,h} &= \frac{\sum_i^n h(\hat{\theta}_i^{\lambda,0})}{n} - M_{\lambda|x|,h} = \left(\frac{\sum_i^n h(\hat{\theta}_i^{\lambda,0})}{n} - \frac{\sum_i^n h(\hat{\theta}_i^{\lambda+\rho,\epsilon})}{n} \right) \\ &+ \left(\frac{\sum_i^n h(\hat{\theta}_i^{\lambda+\rho,\epsilon})}{n} - M_{(\lambda+\rho)|x|+\epsilon x^2/2,h} \right) + (M_{(\lambda+\rho)|x|+\epsilon x^2/2,g} - M_{\lambda|x|,h}) \end{aligned} \quad (307)$$

Letting $\hat{\theta}^{\lambda,0} - \mathbf{p} = \hat{\theta}^{\lambda+\rho,\epsilon}$. Then a Taylor expansion gives us

$$\frac{\sum_i^n h(\hat{\theta}_i^{\lambda,0})}{n} - \frac{\sum_i^n h(\hat{\theta}_i^{\lambda+\rho,\epsilon})}{n} = \frac{\sum_i^n h'(\hat{\theta}_i^{\lambda+\rho,\epsilon}) p_i + h''(\eta_i) p_i^2/2}{n} \quad (308)$$

for some η . Using the Cauchy-Schwartz inequality and using the fact that $h'' < L$ for some value of L we get that

$$\left| \frac{\sum_i^n h(\hat{\theta}_i^{\lambda,0})}{n} - \frac{\sum_i^n h(\hat{\theta}_i^{\lambda+\rho,\epsilon})}{n} \right| \leq \sqrt{\frac{\sum_i^n (h')^2(\hat{\theta}_i^{\lambda+\rho,\epsilon})}{n}} \sqrt{\frac{\sum_i^n p_i^2}{n}} + \frac{L}{2} \frac{\sum_i^n p_i^2}{n} \quad (309)$$

As $h'' < L$ we note that $|h'(x)| < L|x| + C$ for some constant C_2 . As such

$$\frac{\sum_i^n (h')^2(\hat{\theta}_i^{\lambda+\rho,\epsilon})}{n} \leq 2C_2^2 \frac{\sum_i^n (\hat{\theta}_i^{\lambda+\rho,\epsilon})^2}{n} + 2C_2^2 \quad (310)$$

Where C_2 is another positive constant. From theorem 4 the term $\frac{\sum_i^n (\hat{\theta}_i^{\lambda+\rho,\epsilon})^2}{n}$ converges in probability to some value. As such there exists a constant $R > 0$ such that

$$\mathbb{P} \left(\frac{\sum_i^n (h')^2(\hat{\theta}_i^{\lambda+\rho,\epsilon})}{n} \geq R^2 \right) \rightarrow 0 \quad (311)$$

For an arbitrary choice of $\delta > 0$. We choose $\eta_1 > 0$ such that $R\sqrt{\eta_1} + c_1\eta_1/2 < \delta/3$. Furthermore we can verify that we can choose an η_2 such that for every $0 < \epsilon < \eta_2, |\rho| < \eta_2$ that

$$|M_{(\lambda+\rho)|x|+\epsilon x^2/2,h} - M_{\lambda|x|,h}| \leq \frac{\delta}{3} \quad (312)$$

letting $\eta = \min(\eta_1, \eta_2)$. Assume that Lemma 16 holds with a proper choice of ϵ and ρ for this η . This leads to the following holding true with high probability

$$\frac{\sum_i^n p_i^2}{n} < \eta \leq \eta_1 \quad (313)$$

From which we find

$$\left| \frac{\sum_i^n h(\hat{\theta}_i^{\lambda,0})}{n} - \frac{\sum_i^n h(\hat{\theta}_i^{\lambda+\rho,\epsilon})}{n} \right| \leq R\sqrt{\eta_1} + c_1\eta_1/2 < \frac{\delta}{3} \quad (314)$$

Finally we note from theorem 4 that

$$\mathbb{P} \left(\left| \frac{\sum_i^n h(\hat{\theta}_i^{\lambda+\rho,\epsilon})}{n} - M_{(\lambda+\rho)|x|+\epsilon x^2/2,h} \right| > \frac{\delta}{3} \right) \rightarrow 0 \quad (315)$$

Combining all of the bounds we get with high probability that

$$\left| \frac{h(\hat{\theta}^{\lambda,0})}{n} - M_{\lambda|x|,h} \right| \leq \delta \quad (316)$$

Since we can choose delta to be arbitrarily small this leads to the desired results. \square

C.7 Auxiliary lemmas for proving Theorem 6

Lemma 15 (Panahi and Hassibi (2017) lemma 9). *Consider the function $\rho(\mathbf{e}) = \frac{1}{2}\|\mathbf{h} + \mathbf{A}\mathbf{e}\|_2^2 + \lambda\|\mathbf{e} + \boldsymbol{\theta}^*\|_1 + \mathbf{a}^T\mathbf{e}$ and suppose that it is minimized at \mathbf{e}^* . At an arbitrary point \mathbf{e} and $\mathbf{q} \in \partial\rho(\mathbf{e})$, then*

$$\|\mathbf{e} - \mathbf{e}^*\|_2 \leq \frac{1}{\sigma_{\min}^2(\mathbf{A})}\|\mathbf{q}\|_2 \quad (317)$$

Lemma 16 (Panahi and Hassibi (2017) lemma 10). *Consider the function $\rho(\mathbf{e}) = \frac{1}{2}\|\mathbf{h} + \mathbf{P}\mathbf{e}\|_2^2 + \lambda\|\boldsymbol{\theta}^* + \mathbf{e}\|_1 + \frac{\epsilon}{2}\|\mathbf{e}\|_2^2$ and suppose that it is minimized at \mathbf{e}^* . Let \mathbf{e} be an arbitrary point, then*

$$\rho(\mathbf{e}) - \rho(\mathbf{e}^*) \geq \frac{\sigma_{\min}(\mathbf{P})}{2}\|\mathbf{e} - \mathbf{e}^*\|_2^2 \quad (318)$$

Proof. Let $\mathbf{w} = \frac{\mathbf{e} - \mathbf{e}^*}{\|\mathbf{e} - \mathbf{e}^*\|_2}$ and $f(\nu) = \rho(\mathbf{e}^* + \nu\mathbf{w})$. Notice that $\rho(\mathbf{e}) = f(\|\mathbf{e} - \mathbf{e}^*\|_2)$ and f is minimized at 0. A direct calculation shows that f can be written as $f = \frac{1}{2}\alpha\nu^2 + g(\nu)$, where g is convex and $\alpha = \|\mathbf{P}\mathbf{w}\|_2^2 + \epsilon/2 \geq \sigma_{\min}(\mathbf{P})^2$. Then by lemma 17 this results in

$$\rho(\mathbf{e}) - \rho(\mathbf{e}^*) = f(\|\mathbf{e} - \mathbf{e}^*\|_2) - f(0) \geq \frac{\alpha}{2}\|\mathbf{e} - \mathbf{e}^*\|_2^2 \geq \frac{\sigma_{\min}(\mathbf{P})}{2}\|\mathbf{e} - \mathbf{e}^*\|_2^2 \quad (319)$$

□

Lemma 17 (Panahi and Hassibi (2017) lemma 11). *Suppose $g(\nu)$ is a convex function on \mathbb{R} and ν^* is a minimum point of the function $f(\nu) = \frac{\alpha}{2}\nu^2 + g(\nu)$. Then for any $\nu \in \mathbb{R}$,*

$$f(\nu) - f(\nu^*) \geq \frac{\alpha}{2}(\nu - \nu^*)^2 \quad (320)$$

Proof. From the optimality of ν^* , we have that $-\alpha\nu^* \in \partial g(\nu^*)$. Therefore,

$$g(\nu) \geq g(\nu^*) - \alpha\nu^*(\nu - \nu^*) \quad (321)$$

Hence,

$$f(\nu) - f(\nu^*) = \alpha\nu^*(\nu - \nu^*) + \frac{\alpha}{2}(\nu - \nu^*)^2 + g(\nu) - g(\nu^*) \geq \frac{\alpha}{2}(\nu - \nu^*)^2$$

□

D Example Case : Elastic Net Regularization

We consider the case of Elastic Net Regularization, the case that

$$r(\boldsymbol{\theta}) = \lambda\|\boldsymbol{\theta}\|_1 + \frac{\alpha}{2}\|\boldsymbol{\theta}\|_2^2 \quad (322)$$

A simple computation gives us that

$$\hat{\boldsymbol{\theta}}_i = (\text{prox}_{\frac{\lambda}{2c_1}}(\boldsymbol{\theta}^* - \frac{c_2\sqrt{\gamma}}{2c_1}\boldsymbol{\phi}))_i = \begin{cases} \frac{2c_1\boldsymbol{\theta}_i^*}{2c_1+\alpha} - \frac{c_2\sqrt{\gamma}}{2c_1+\alpha}\phi_i - \frac{\lambda}{2c_1+\alpha} & \frac{2c_1\boldsymbol{\theta}_i^*}{2c_1+\alpha} + \frac{c_2\sqrt{\gamma}}{2c_1+\alpha}\phi_i > \frac{\lambda}{2c_1+\alpha} \\ \frac{2c_1\boldsymbol{\theta}_i^*}{2c_1+\alpha} + \frac{c_2\sqrt{\gamma}}{2c_1+\alpha}\phi_i & \frac{2c_1\boldsymbol{\theta}_i^*}{2c_1+\alpha} + \frac{c_2\sqrt{\gamma}}{2c_1+\alpha}\phi_i < -\frac{\lambda}{2c_1+\alpha} \\ 0 & \|\frac{2c_1\boldsymbol{\theta}_i^*}{2c_1+\alpha} + \frac{c_2\sqrt{\gamma}}{2c_1+\alpha}\phi_i\| \leq \frac{\lambda}{2c_1+\alpha} \end{cases} \quad (323)$$

We note that this can equivalently be expressed in the form of a soft thresholding operator

$$(\hat{\boldsymbol{\theta}}_3)_i = \mathcal{T}_{\frac{\lambda}{2c_1+\alpha}}\left(\frac{2c_1\boldsymbol{\theta}_i^*}{2c_1+\alpha} - \frac{c_2\sqrt{\gamma}}{2c_1+\alpha}\phi_i\right) \quad (324)$$

Substituting the value of the proximal operator in to the Moreau envelope we find that

$$\begin{aligned}
 & \frac{1}{m} \sum_i^m \mathcal{M}_{\frac{1}{2c_1}r}(\boldsymbol{\theta}^* - \frac{c_2\sqrt{\gamma}}{2c_1}\boldsymbol{\phi}) = \\
 & \frac{1}{2c_1 + \alpha} \left(\alpha c_1 (\theta_i^*)^2 + 2c_1 \lambda \theta_i^* - \frac{1}{2} \lambda^2 + \frac{c_2^2 \gamma}{4c_1} \phi_i^2 + \alpha c_2 \sqrt{\gamma} \phi \theta_i^* + \lambda c_2 \sqrt{\gamma} \phi \right) \mathbf{1}_{\{\phi_i < -\zeta_{1i}\}} \\
 & \frac{1}{2c_1 + \alpha} \left(\alpha c_1 (\theta_i^*)^2 - 2c_1 \lambda \theta_i^* - \frac{1}{2} \lambda^2 + \frac{\alpha c_2^2 \gamma}{4c_1} \phi_i^2 - \alpha c_2 \sqrt{\gamma} \phi \theta_i^* + \lambda c_2 \sqrt{\gamma} \phi \right) \mathbf{1}_{\{\phi_i > \zeta_{2i}\}} \\
 & \left(c_1 (\theta_i^*)^2 + \frac{c_2^2 \gamma}{4c_1} \phi_i^2 - c_2 \sqrt{\gamma} \theta_i^* \phi_i \right) \mathbf{1}_{\{\zeta_{1i} \leq \phi_i \leq \zeta_{2i}\}} \quad (325)
 \end{aligned}$$

in which $\mathbf{1}_A$ is the characteristic function on the set A , and

$$\zeta_{1i} = \frac{(\lambda - 2\hat{c}_1 \theta_i^*)}{\sqrt{\gamma \hat{c}_2}} \quad \zeta_{2i} = \frac{(\lambda + 2\hat{c}_1 \theta_i^*)}{\sqrt{\gamma \hat{c}_2}} \quad (326)$$

Taking the expectation of the envelope with respect to ϕ and making use of Steins lemma one can obtain

$$\begin{aligned}
 & \frac{1}{m} \mathbb{E} \sum_i^m \mathcal{M}_{\frac{1}{2c_1}r}(\boldsymbol{\theta}^* - \frac{c_2\sqrt{\gamma}}{2c_1}\boldsymbol{\phi}) = \\
 & \frac{1}{2c_1 + \alpha} \left(\alpha c_1 (\theta_i^*)^2 + \frac{\alpha c_2^2 \gamma}{4c_1} + 2c_1 \lambda \theta_i^* - \frac{1}{2} \lambda^2 \right) Q(\zeta_{1i}) + \frac{1}{\sqrt{2\pi}(2c_1 + \alpha)} \left(\frac{c_2^2 \gamma \zeta_{1i}}{4c_1} + \alpha c_2 \sqrt{\gamma} \theta_i^* + \lambda c_2 \sqrt{\gamma} \right) e^{-\zeta_{1i}^2/2} \\
 & \frac{1}{2c_1 + \alpha} \left(\alpha c_1 (\theta_i^*)^2 + \frac{\alpha c_2^2 \gamma}{4c_1} - 2c_1 \lambda \theta_i^* - \frac{1}{2} \lambda^2 \right) Q(\zeta_{2i}) + \frac{1}{\sqrt{2\pi}(2c_1 + \alpha)} \left(\frac{\alpha c_2^2 \gamma \zeta_{2i}}{4c_1} - \alpha c_2 \sqrt{\gamma} \theta_i^* + \lambda c_2 \sqrt{\gamma} \right) e^{-\zeta_{2i}^2/2} \\
 & \left(c_1 (\theta_i^*)^2 + \frac{c_2^2 \gamma}{4c_1} \right) (1 - Q(\zeta_{1i}) - Q(\zeta_{2i})) - \frac{c_2 \sqrt{\gamma} \theta_i^*}{\sqrt{2\pi}} \left(e^{-\zeta_{1i}^2/2} - e^{-\zeta_{2i}^2/2} \right) + \frac{c_2^2 \gamma}{4c_1 \sqrt{2\pi}} \left(-\zeta_{1i} e^{-\zeta_{1i}^2/2} - \zeta_{2i} e^{-\zeta_{2i}^2/2} \right) \quad (327)
 \end{aligned}$$

In which $Q(\cdot)$ is the Q-function. Defined to be

$$Q(x) = \frac{1}{\sqrt{2\pi}} \int_x^\infty e^{-\frac{u^2}{2}} du \quad (328)$$

This expression may be implemented in code and simply evaluated for any choice of the parameters.

D.1 Sparsity

The effect of the ℓ_1 regularization term is to promote sparsity in the solution vector. Let s denote the number of elements of $\hat{\boldsymbol{\theta}}$ that are non-zero. We see that

$$s = \mathbb{E} \sum_i^m \mathbf{1}_{\hat{\theta}_i \neq 0} = \sum_i^m 1 - \mathbb{P}(\hat{\theta}_i = 0) = \sum_i^m 1 - \mathbb{P}(-\zeta_{1i} \leq \phi_i \leq \zeta_{2i}) = \sum_i^m Q(\zeta_{1i}) + Q(\zeta_{2i}) \quad (329)$$

We further consider the term $\frac{1}{m} \hat{\mathbf{e}}^T \boldsymbol{\phi}$ and consider what this concentrates on

$$\frac{1}{m} \mathbb{E}[\hat{\mathbf{e}}^T \boldsymbol{\phi}] = \frac{1}{m} \sum_i^m \mathbb{E}[\hat{e}_i \phi_i] = \frac{1}{m} \sum_i^m \mathbb{E}[\hat{\theta}_i \phi_i] = -\frac{c_2 \sqrt{\gamma}}{2c_1 + \alpha} \frac{1}{m} \mathbb{E} \left[\left(\mathcal{T}'_{\frac{\lambda}{2c_1 + \alpha}} \left(\left(\frac{2c_1 \theta_i^*}{2c_1 + \alpha} + \frac{c_2 \sqrt{\gamma}}{2c_1 + \alpha} \phi_i \right) \right) \right) \right] \quad (330)$$

where in the last inequality we have made use of steins lemma, and \mathcal{T}' is the derivative of the soft thresholding operator defined in (324). We note that the derivative of the soft thresholding operator is the value of s that we are looking for. In symbols

$$\mathbf{1}_{\{\hat{\theta}_i \neq 0\}} \mathcal{T} = \mathcal{T}' \quad (331)$$

From this we note that

$$\frac{1}{m} \mathbb{E}[\mathbf{e}^T \boldsymbol{\phi}] = -\frac{c_2 \sqrt{\gamma}}{2c_1 + \alpha} \frac{s}{m} \quad (332)$$

The value of $\frac{1}{m} \mathbb{E}[\mathbf{e}^T \boldsymbol{\phi}]$ may also be computed directly from definition of $\hat{\boldsymbol{\theta}}$ (equation (323)). Combining these expressions the value of s may be computed.

E Numerical Simulation Detail

We implement the optimization problem P_3 (14) by making use of the explicitly computed Moreau envelope for the case of elastic net (327). The optimization is solved using a standard iterative approach in which the inner optimizations are solved at constant values of the outer optimizations. This is repeated iteratively until all parameters are determined. Zeroth order gradient methods were attempted, but were highly dependent on the starting choices of the parameters β, q, ξ, t , and frequently failed to converge.

The experimental verification was completed using synthetic data, in which the data points \mathbf{z}_i and the weight matrix \mathbf{W} was drawn from standard normal distributions. The elastic net optimization was solved using the python package `cvxpy`. The values of n and m were chosen such that $n + m = 1000$ and that $m/n \approx \gamma$, for a chosen ratio γ . Each sample was averaged 100 times to account for the randomness in both the input data \mathbf{z} and the weights \mathbf{W} .

E.1 Effective Sparsity

In this section, we plot the effective sparsity s for elastic net as a function of the regularization strength λ for a number of values of $\gamma = \frac{m}{n}$. Recall that s gives the number of nonzero elements in the solution vector $\hat{\boldsymbol{\theta}}$. The plots for the ratios $\frac{s}{m}$ and $\frac{s}{n}$ may be seen in figure 3. The ℓ_2^2 regularization strength was fixed with parameter $\alpha = 0.001$. The solid lines are the theoretical predictions while the dots are determined experimentally. For the experimental values the solution vector $\hat{\boldsymbol{\theta}}$ was determined using a solver, then each element of the solution vector, it was determined to be “zero” (i.e. sparse) if its value was less than $\frac{0.01}{\sqrt{m}}$.

We can see from the figures that for all values of γ the sparsity is similar at both large and small values of regularization. As the number of model parameters increases relative to the number of data points, i.e. as γ grows larger, the regularization strength required to induce a sparse solution drops. Recalling that true solution was half zeros, the value of regularization strength at which $\frac{s}{m} = 0.5$ matches well with the regularization strength that minimizes the generalization error in figure 2(b).

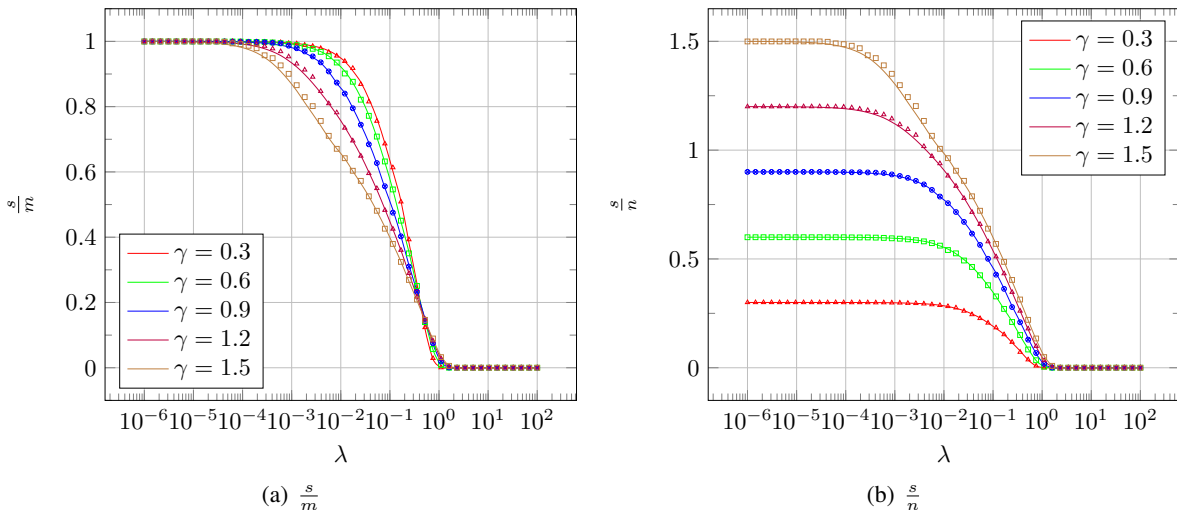


Figure 3: The effective sparsity s as a ratio to the number of model parameters m or the number of data points n for elastic net regularization for varying strengths of the regularization parameter λ . The ℓ_2 regularization term was fixed to 0.001. Multiple values of $\gamma = \frac{m}{n}$ are considered. Solid line is the theoretical prediction, and the dots are experimental values.