
Regularization properties of adversarially-trained linear regression

Antônio H. Ribeiro
Uppsala University
antonio.horta.ribeiro@it.uu.se

Dave Zachariah
Uppsala University
dave.zachariah@it.uu.se

Francis Bach
PSL Research University, INRIA
francis.bach@inria.fr

Thomas B. Schön
Uppsala University
thomas.schon@it.uu.se

Abstract

State-of-the-art machine learning models can be vulnerable to very small input perturbations that are adversarially constructed. Adversarial training is an effective approach to defend against it. Formulated as a min-max problem, it searches for the best solution when the training data were corrupted by the worst-case attacks. Linear models are among the simple models where vulnerabilities can be observed and are the focus of our study. In this case, adversarial training leads to a convex optimization problem which can be formulated as the minimization of a finite sum. We provide a comparative analysis between the solution of adversarial training in linear regression and other regularization methods. Our main findings are that: (A) Adversarial training yields the minimum-norm interpolating solution in the overparameterized regime (more parameters than data), as long as the maximum disturbance radius is smaller than a threshold. And, conversely, the minimum-norm interpolator is the solution to adversarial training with a given radius. (B) Adversarial training can be equivalent to parameter shrinking methods (ridge regression and Lasso). This happens in the underparametrized region, for an appropriate choice of adversarial radius and zero-mean symmetrically distributed covariates. (C) For ℓ_∞ -adversarial training—as in square-root Lasso—the choice of adversarial radius for optimal bounds does not depend on the additive noise variance. We confirm our theoretical findings with numerical examples.

1 Introduction

Adversarial attacks generated striking examples of the brittleness of modern machine learning. The framework considers inputs contaminated with disturbances deliberately chosen to maximize the model error and, even for small disturbances, can cause a substantial performance drop in otherwise state-of-the-art models [1]–[6]. Adversarial training [7] is one of the most effective approaches for deep learning models to defend against adversarial attacks [7]–[9]. It considers training models on samples that have been modified by an adversary, with the goal of obtaining a model that will be more robust when faced with new adversarially perturbed samples. The training procedure is formulated as a min-max problem, searching for the best solution to the worst-case attacks.

Despite its success in producing state-of-the-art results in adversarial defense benchmarks [10], there are still major challenges in using adversarial training. The min-max problem is a hard optimization problem to solve and it is still an open question how to solve it efficiently. Moreover, these methods still produce large errors in new adversarially disturbed test points, often much larger than the adversarial error obtained during training.




Figure 1: **Regularization paths** estimated in the Diabetes dataset [18]. Regularization paths are plots showing the coefficient estimates for varying regularization parameter λ or, in adversarial training, perturbation radius δ . This type of plot is commonly used in analysing Lasso and variants, i.e. [18], [19]. On the horizontal axis, we give the ℓ_1 -norm of the estimated parameter. On the vertical axis, the coefficients obtained for each method. The dataset has $p = 10$ baseline variables (age, sex, body mass index, average blood pressure, and six blood serum measurements), which were obtained for $n = 442$ diabetes patients. The model output is a quantitative measure of the disease progression. See Appendix B.5 for the relationship between $\|\hat{\beta}\|_1$ and δ and λ .

To get insight into adversarial training and try to tackle these challenges, a growing body of work [11]–[15] studies fundamental properties of adversarial attacks and adversarial training in linear models. Linear models allow for analytical analysis while still reproducing phenomena observed in state-of-the-art models. Consider a training dataset $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$ consisting of n data points of dimension $\mathbb{R}^p \times \mathbb{R}$, adversarial training in linear regression corresponds to minimizing (in $\beta \in \mathbb{R}^p$)

$$R^{\text{adv}}(\beta; \delta, \|\cdot\|) = \frac{1}{n} \sum_{i=1}^n \max_{\|\Delta \mathbf{x}_i\| \leq \delta} |y_i - (\mathbf{x}_i + \Delta \mathbf{x}_i)^\top \beta|^2. \quad (1)$$

Restricting the analysis to linear models allows for a simplified analysis: the problem is convex and the next proposition allows us to express R^{adv} in terms of the dual norm, $\|\beta\|_* = \sup_{\|\mathbf{x}\| \leq 1} |\beta^\top \mathbf{x}|$.

Proposition 1 (Dual formulation). Let $\|\cdot\|_*$ be the dual norm of $\|\cdot\|$, then

$$R^{\text{adv}}(\beta; \delta, \|\cdot\|) = \frac{1}{n} \sum_{i=1}^n (|y_i - \mathbf{x}_i^\top \beta| + \delta \|\beta\|_*)^2. \quad (2)$$

This simple reformulation removes one of the major challenges: it removes the inner optimization problem and yields a model that is more amenable to analysis. Similar reformulations can be obtained for classification. This simplification made the analysis of adversarial attacks in linear models fruitful to get insights into the properties of adversarial training and examples. From giving a counterexample to the idea that deep neural networks’ vulnerabilities to adversarial attacks were caused by their nonlinearity [2], [5], [16], to help explaining how implicitly regularized overparametrized models can be robust, see [17].

In this paper, we further explore this reformulation and *provide a thorough characterization of adversarial training in linear regression problems*. We do this in a comparative fashion, establishing when the solution of adversarially-trained linear models coincides with the solution of traditional regularization methods. For instance, a simple inspection of (2) hints at the similarities with other parameter shrinking methods. This similarity is confirmed in Figure 1,¹ which illustrates how the solutions of ℓ_∞ -adversarial training can be almost indistinguishable from Lasso. We explain the similarities and differences with other methods. Our contributions are:

- A. In the overparametrized region, we prove that the *minimum-norm interpolator is the solution to adversarial training for δ smaller than a certain threshold* (Section 4).
- B. We establish *conditions under which the solution coincides with Lasso and ridge regression*, respectively (Section 5).
- C. We show that adversarial training can be framed in the *robust regression* framework, and use it to establish connections with the *square-root Lasso*. We show that adversarial training (like the *square-root Lasso*) can obtain bounds on the prediction error that do not require the noise variance to be known (Section 6).
- D. We prove a more *general version of Proposition 1*, valid for general lower-semicontinuous and convex loss functions (Section 8).

¹Code for reproducing the figures is available in: <https://github.com/antonior92/advtrain-linreg>

2 Related work

Generalization of minimum-norm interpolators. The study of minimum-norm interpolators played a key role in explaining why overparametrized models generalize—an important open question for which traditional theory failed to explain empirical results [20]—and for which significant progress has been made over the past few years [21]. These estimates provide a simple scenario where we can interpolate noisy data and still generalize well. Minimum-norm interpolators have indeed been quite a fruitful setting to study the phenomenon of *benign overfitting*, i.e., when the model interpolates the training data but still generalizes well to new samples. In [22] the authors use these estimates to prove consistency, a development that later had several extensions [23]. From another angle, in a series of insightful papers, Belkin *et al.* [24], [25] explore the phenomena of *double-descent*: where a double-descent curve subsumes the textbook U-shaped bias–variance trade-off curve, with a second decrease in the error occurring beyond the point where the model has reached the capacity of interpolating the training data. Interestingly, minimum-norm interpolators are also basic scenarios for observing this phenomena [24]–[26]. *Our research connects the research on robustness to adversarial attacks to the study of generalization of minimum-norm interpolators.*

Adversarial training in linear models. The generalization of adversarial attacks in linear models is well-studied. Tsipras *et al.* [16] and Ilyas *et al.* [5] use linear models to explain the conflict between robustness and high-performance models observed in neural networks; Ribeiro *et al.* [17] use these models to show how overparameterization affects robustness to perturbations; Taheri *et al.* [11] derive asymptotics for adversarial training in binary classification. Dan *et al.* [27] and Dobriban *et al.* [28] study adversarial robustness in Gaussian classification problems. Javanmard *et al.* [12] provide asymptotics for adversarial training in linear regression, Javanmard *et al.* [29] in classification settings and Hassani *et al.* [13] for random feature regressions. Min *et al.* [14] investigate how the dataset size affects adversarial performance. Yin *et al.* [15] provide an analysis of ℓ_∞ -attack on linear classifiers based on Rademacher complexity. These works, however, focus on the generalization properties. *We provide a fresh perspective on the problem by clarifying the connection of adversarial training to other regularization methods.*

Robust regression and square-root Lasso. The properties of Lasso [30] for recovering sparse parameters in noisy settings have been extensively studied, and bounds on the parameter estimation error are well-known [31]. However, these results rely on choices of regularization parameters that depend on the variance of additive noise. Square-root Lasso [32] was proposed to circumvent this difficulty. We show that for ℓ_∞ -adversarial, *similarly to the square-root Lasso*, bounds can be obtained with the adversarial radius set without knowledge about the variance of the additive noise. We also show that both methods fit into the robust regression framework [33]. The work on robust classification [34] is also connected to our work, there they show that certain robust support vector machines are equivalent to SVM. Our results for classification in Section 8 can be viewed as a generalization of their Theorem 3.

Dual formulation. Variations of Proposition 1 are presented in [12], [17], [35]. Equivalent results in the context of classification are presented in [2], [15]. We use the reformulation extensively in our developments and also present a generalized statement for it.

3 Background

Different estimators will be relevant to our developments and will be compared in this paper.

Adversarially-trained linear regression. Our main object of study is the minimization of

$$R^{\text{adv}}(\beta; \delta, \|\cdot\|) = \frac{1}{n} \sum_{i=1}^n \max_{\|\Delta \mathbf{x}_i\| \leq \delta} |y_i - (\mathbf{x}_i + \Delta \mathbf{x}_i)^\top \beta|^2 \stackrel{(a)}{=} \frac{1}{n} \sum_{i=1}^n (|y_i - \mathbf{x}_i^\top \beta| + \delta \|\beta\|_*)^2,$$

where equality (a) follows from Proposition 1. We use $\|\beta\|_* = \sup_{\|\mathbf{x}\| \leq 1} |\beta^\top \mathbf{x}|$ to denote the dual norm of $\|\cdot\|$. We highlight that the ℓ_2 -norm, $\|\beta\|_2 = \sqrt{\sum_i |\beta_i|^2}$ is dual to itself. The ℓ_1 -norm, $\|\beta\|_1 = \sum_i |\beta_i|$, is the dual norm of the ℓ_∞ -norm, $\|\beta\|_\infty = \max_i |\beta_i|$. When the adversarial disturbances are constrained to the ℓ_p ball: $\{\Delta \mathbf{x} : \|\Delta \mathbf{x}\|_p \leq \delta\}$ we will call it ℓ_p -adversarial training. Our focus will be on ℓ_∞ and ℓ_2 -adversarial training.

Parameter-shrinking methods. Parameter-shrinking methods explicitly penalize large values of β . Regression methods that involve shrinkage include Lasso [30] and ridge regression, which minimize,

respectively

$$R^{\text{lasso}}(\boldsymbol{\beta}; \lambda) = \frac{1}{n} \sum_{i=1}^n |y_i - \mathbf{x}_i^\top \boldsymbol{\beta}|^2 + \lambda \|\boldsymbol{\beta}\|_1 \quad \text{and} \quad R^{\text{ridge}}(\boldsymbol{\beta}; \lambda) = \frac{1}{n} \sum_{i=1}^n |y_i - \mathbf{x}_i^\top \boldsymbol{\beta}|^2 + \lambda \|\boldsymbol{\beta}\|_2^2.$$

Square-root Lasso: Setting (with guarantees) the Lasso regularization parameters requires knowledge about the variance of the noise. Square-root Lasso [32] circumvent this difficulty by minimizing:

$$R^{\sqrt{\text{lasso}}}(\boldsymbol{\beta}, \lambda) = \sqrt{\frac{1}{n} \sum_{i=1}^n |y_i - \mathbf{x}_i^\top \boldsymbol{\beta}|^2} + \lambda \|\boldsymbol{\beta}\|_1. \quad (3)$$

Minimum-norm interpolator: Let $\mathbf{X} \in \mathbb{R}^{n \times p}$ denote the matrix of stacked vectors \mathbf{x}_i and $\mathbf{y} \in \mathbb{R}^n$ is the vector of stacked outputs. When matrix \mathbf{X} has full row rank, the linear system $\mathbf{X}\boldsymbol{\beta} = \mathbf{y}$ has multiple solutions. In this case the minimum $\|\cdot\|_*$ -norm interpolator is the solution of

$$\min_{\boldsymbol{\beta}} \|\boldsymbol{\beta}\|_* \quad \text{subject to} \quad \mathbf{X}\boldsymbol{\beta} = \mathbf{y}. \quad (4)$$

4 Adversarial training in the overparametrized regime

Our first contribution is to provide conditions for when the minimum-norm interpolator is equivalent to adversarial training in the overparametrized case (we will assume that $\text{rank}(\mathbf{X}) = n$ and $n < p$). On the one hand, our result gives further insight into adversarial training. On the other hand, it allows us to see minimum-norm interpolators as a solution of the adversarial training problem.

Theorem 1. Assume the matrix $\mathbf{X} \in \mathbb{R}^{n \times p}$ to have full row rank, let $\hat{\boldsymbol{\alpha}}$ denote the solution of the dual problem $\max_{\|\boldsymbol{\alpha}^\top \mathbf{X}\| \leq 1} \boldsymbol{\alpha}^\top \mathbf{y}$, and $\bar{\delta}$ denote the threshold

$$\bar{\delta} = (n \|\hat{\boldsymbol{\alpha}}\|_\infty)^{-1}. \quad (5)$$

The minimum $\|\cdot\|_*$ -norm interpolator minimizes the adversarial risk $R^{\text{adv}}(\boldsymbol{\beta}, \delta, \|\cdot\|)$ if and only if $\delta \in (0, \bar{\delta}]$.

Remark 1 (Bounds on $\bar{\delta}$). The theorem allows us to directly compute $\bar{\delta}$, but it does require solving the dual problem to the minimum-norm solution. In Appendix A.3, we provide bounds on $\bar{\delta}$ that avoid solving the optimization problem. For instance, for ℓ_∞ -adversarial attacks, we have the following bounds depending on the singular values of \mathbf{X} : $\frac{1}{\sqrt{p}} \sigma_n(\mathbf{X}) \leq n\bar{\delta} \leq \sqrt{p} \sigma_1(\mathbf{X})$. Where σ_1 and σ_n denote the largest and smallest *positive* singular values of \mathbf{X} , respectively.

Proof of Theorem 1. Let $\hat{\boldsymbol{\beta}}$ and $\hat{\boldsymbol{\alpha}}$ be the minimum $\|\cdot\|_*$ -norm solution and the solution of the associated dual problem. Throughout the proof we denote $R^{\text{adv}}(\boldsymbol{\beta}) = R^{\text{adv}}(\boldsymbol{\beta}; \delta, \|\cdot\|)$, dropping the last two arguments. The subgradient of $R^{\text{adv}}(\boldsymbol{\beta})$ evaluated at $\hat{\boldsymbol{\beta}}$ is given by²

$$\partial R^{\text{adv}}(\hat{\boldsymbol{\beta}}) = \frac{2\delta}{n} \sum_{i=1}^n \|\hat{\boldsymbol{\beta}}\|_* (\rho_i \mathbf{x}_i + \delta \partial \|\hat{\boldsymbol{\beta}}\|_*), \quad (6)$$

for any $\boldsymbol{\rho} = (\rho_1, \dots, \rho_n) \in \mathbb{R}^n$ that satisfies $\|\boldsymbol{\rho}\|_\infty \leq 1$. We have that any element of $\partial \|\hat{\boldsymbol{\beta}}\|_*$ can be written as $\mathbf{X}^\top \hat{\boldsymbol{\alpha}}$ (we prove that in Lemma 1 in the Appendix). Hence, we can rewrite

$$\partial R^{\text{adv}}(\hat{\boldsymbol{\beta}}) = \frac{2\delta}{n} \sum_{i=1}^n \|\hat{\boldsymbol{\beta}}\|_* (\rho_i \mathbf{x}_i + \delta \mathbf{X}^\top \hat{\boldsymbol{\alpha}}) = 2\delta \|\hat{\boldsymbol{\beta}}\|_* \left(\frac{\mathbf{X}^\top \boldsymbol{\rho}}{n} + \delta \mathbf{X}^\top \hat{\boldsymbol{\alpha}} \right).$$

If $\delta \leq \bar{\delta} = 1/(n \|\hat{\boldsymbol{\alpha}}\|_\infty)$, we can take $\boldsymbol{\rho} = -n\delta \hat{\boldsymbol{\alpha}}$ and the subderivative contains zero. On the other hand, if $\delta > 1/(n \|\hat{\boldsymbol{\alpha}}\|_\infty)$ then $(\frac{\mathbf{X}^\top \boldsymbol{\rho}}{n} + \delta \mathbf{X}^\top \hat{\boldsymbol{\alpha}})$ is not zero for $\|\boldsymbol{\rho}\|_\infty \leq 1$. \square

The minimum ℓ_1 -norm interpolator is well studied in the context of ‘basis pursuit’ and allows the recovery of low-dimensional representations in sparse signals [39]. The interest in the minimum ℓ_2 -norm is more recent. It played an important role in studying the interplay between interpolation and generalization, being used in many recent papers where the double-descent [24], [26] and the benign overfitting phenomena [22] were observed.

²The subgradient of a function $\omega : \mathbb{R}^p \rightarrow \mathbb{R}$ is the set:

$$\partial \omega(\boldsymbol{\beta}_0) = \{\mathbf{v} \in \mathbb{R}^p : \omega(\boldsymbol{\beta}) - \omega(\boldsymbol{\beta}_0) \geq \mathbf{v}(\boldsymbol{\beta} - \boldsymbol{\beta}_0), \forall \boldsymbol{\beta} \in \mathbb{R}^p\}.$$

See [36]–[38] for properties. We use $\partial \|\hat{\boldsymbol{\beta}}\|_*$ to denote the subgradient of $\omega(\boldsymbol{\beta}) = \|\boldsymbol{\beta}\|_*$ evaluated at $\hat{\boldsymbol{\beta}}$.

Theorem 1 gives a new perspective on the robustness of the minimum-norm interpolator, allowing us to see it as a solution to adversarial training with adversarial radius $\bar{\delta}$. Figure 2 shows that the radius $\bar{\delta}$ of the adversarial training problem corresponding to the minimum-norm interpolator increases with the ratio p/n . It is natural to expect that training with a larger radius would yield more adversarially robust models on new test points. The next result formalizes this intuition by establishing an upper bound on the robustness gap: the difference between the square-root of the expected adversarial squared error $\mathcal{R}_*^{\text{adv}}(\beta; \delta_{\text{test}}, \|\cdot\|) = \mathbb{E}_{y_0, \mathbf{x}_0} [\max_{\|\Delta \mathbf{x}_i\| \leq \delta_{\text{test}}} (y_0 - (\mathbf{x}_0 + \Delta \mathbf{x}_0)^\top \beta)^2]$ and the expected squared error $\mathcal{R}_*(\beta) = \mathbb{E}_{y_0, \mathbf{x}_0} [(y_0 - \mathbf{x}_0^\top \beta)^2]$. The upper bound depends on the in-training adversarial error and on the ratio $\frac{\delta_{\text{test}}}{\delta}$ which are quantities that you can directly compute.




Figure 2: **Threshold $\bar{\delta}$ vs. number of features.** We fix $n = 60$ and show the value of $\bar{\delta}$ (defined in (5)) as a function of the number of features p . The dotted lines give the reference $\delta_{\text{test}} = 0.01 \mathbb{E}[\|\mathbf{x}\|]$ for comparison.

Proposition 2. Assume \mathbf{X} to have full row rank and let $\hat{\beta}$ be the minimum $\|\cdot\|_*$ -norm interpolator, then

$$\sqrt{\mathcal{R}_*^{\text{adv}}(\hat{\beta}; \delta_{\text{test}}, \|\cdot\|)} - \sqrt{\mathcal{R}_*(\hat{\beta})} \leq \frac{\delta_{\text{test}}}{\bar{\delta}} \sqrt{R^{\text{adv}}(\hat{\beta}; \bar{\delta}, \|\cdot\|)}. \quad (7)$$

Remark 2. The example in Figure 2 is one case where adding more features makes the minimum-norm interpolator more robust. Examples showing the opposite and illustrating that adding more features to linear models can make them less adversarially robust are abundant in the literature [16], [17]. Indeed, examples that consider the minimum ℓ_2 -interpolator subject to ℓ_∞ -adversarial attacks during test-time result in this scenario, as discussed in [17]. Proposition 7 in the Appendix is the equivalent of Proposition 2 for this case (mismatched norms in train and test). There, the upper bound grows with \sqrt{p} , which explains the vulnerability in this scenario.

Remark 3. A pitfall of analyzing the minimum-norm solution properties in linear models is that it requires the analysis of nested problems: different choices of p require different covariates \mathbf{x} . The random projection model proposed by Bach [40] avoids this pitfall by considering the input \mathbf{x} is fixed, but only a projected version of it $\mathbf{S}\mathbf{x}$ is observed, with the number of parameters being estimated changing with the number of projections observed. Adversarially training this model consists in finding a parameter $\hat{\beta}$ that minimizes: $\frac{1}{n} \sum_{i=1}^n \max_{\|\Delta \mathbf{x}_i\| \leq \delta} |y_i - (\mathbf{x}_i + \Delta \mathbf{x}_i)^\top \mathbf{S}^\top \beta|^2$. We generalize our results so that we can also study this case (See Appendix A.4).

5 Adversarial training and parameter-shrinking methods

Theorem 1 stated an equivalence with the minimum-norm solution when δ is small. The next result applies to the other extreme of the spectrum, characterizing the case when δ is large.

Proposition 3 (Zero solution of adversarial training). The zero solution $\hat{\beta} = \mathbf{0}$ minimizes the adversarial training if and only if $\delta \geq \frac{\|\mathbf{X}^\top \mathbf{y}\|}{\|\mathbf{y}\|_1}$.

In the remaining of this section, we will characterize the solution between these two extremes. We compare adversarial training to Lasso and ridge regression.

5.1 Relation to Lasso and ridge regression

Proposition 1 makes it clear by inspection that adversarial training is similar to other well-known parameter-shrinking regularization methods. Indeed, for ℓ_∞ -adversarial attacks, the cost function $R^{\text{adv}}(\beta; \delta, \|\cdot\|_\infty)$ in its dual form is remarkably similar to Lasso [30] and $R^{\text{adv}}(\beta; \delta, \|\cdot\|_2)$, to ridge regression (the cost functions were presented in Section 3). In Figures 1 and 3, we show the regularization paths of these methods (the dataset is described in [18]). We observe that *ℓ_∞ -adversarial training produces sparse solutions* and that Lasso and ℓ_∞ -adversarial training have extremely similar regularization paths. There are also striking similarities between ridge regression and ℓ_2 -adversarial training, with the notable difference that for large δ the solution of ℓ_2 -adversarial




Figure 3: **Regularization paths** in the Diabetes dataset [18] (see dataset description in Figure 1). On the horizontal axis, we give the inverse of the regularization parameter (in log scale). On the vertical axis, the coefficients.

training is zero (which is explained by Proposition 3). The next proposition justifies why we can observe such similarities in part of the regularization path. It applies to data that has been normalized and for positive responses y (as it is the case in our example).

Proposition 4. Assume the output is positive and the data is normalized: $\mathbf{y} \geq 0$ and $\mathbf{X}^\top \mathbf{1} = \mathbf{0}$. The solution of adversarial training $\hat{\beta}$, for $\|\hat{\beta}\|_* \leq \min_i \frac{|y_i|}{\|x_i\|}$ is also the solution of

$$\min_{\beta} \frac{1}{n} \sum_{i=1}^n |y_i - \mathbf{x}_i^\top \beta|^2 + \left(\delta \|\beta\|_* + \frac{1}{n} \|\mathbf{y}\|_1 \right)^2.$$

The proposition is proved in Appendix B. It establishes the equivalence between ℓ_∞ -adversarial training and Lasso (even though there is not a closed-formula expression for the map between δ and λ'). Indeed, *under the assumptions of the propositions, $\hat{\beta}$ is the solution of ℓ_∞ -adversarial problem only if it is the solution of*

$$\min_{\beta} \frac{1}{n} \sum_{i=1}^n (y_i - \mathbf{x}_i^\top \beta) + \lambda' \|\beta\|_1, \quad (8)$$

for some λ' . We can prove this statement using the proposition: under the appropriate assumptions, the ℓ_∞ -adversarial problem solution is also a solution to the problem

$$\min_{\beta} \frac{1}{n} \sum_{i=1}^n (y_i - \mathbf{x}_i^\top \beta)^2 + \left(\delta \|\beta\|_1 + \frac{1}{n} \|\mathbf{y}\|_1 \right)^2,$$

which in turn, is the Lagrangian formulation of the following constrained optimization problem

$$\min_{\beta} \frac{1}{n} \sum_{i=1}^n (y_i - \mathbf{x}_i^\top \beta)^2 \text{ subject to } \left(\delta \|\beta\|_1 + \frac{1}{n} \|\mathbf{y}\|_1 \right)^2 \leq \Delta,$$

for some $\Delta \geq \frac{1}{n} \|\mathbf{y}\|_1$. The constraint can be rewritten as $\delta \|\beta\|_1 \leq \sqrt{\Delta} - \frac{1}{n} \|\mathbf{y}\|_1$. And, the result follows because (8) is the Lagrangian formulation of this modified problem. A similar reasoning could be used to connect the result for ℓ_2 -adversarial training with ridge regression.

More generally, even when the condition on $\hat{\beta}$ is not satisfied and if \mathbf{y} is negative, we show in Appendix B that for zero-mean and symmetrically distributed covariates, i.e., $\mathbb{E}[\mathbf{x}] = 0$ and $(\mathbf{x} \sim -\mathbf{x})$, adversarial training has approximately the same solution as

$$\min_{\beta} \frac{1}{n} \sum_{i=1}^n |y_i - \mathbf{x}_i^\top \beta|^2 + \left(\delta \|\beta\|_* + \frac{1}{n} \mathbf{s}^\top \mathbf{y} \right)^2,$$

for $\mathbf{s} = \text{sign}(\mathbf{y} - \mathbf{X}^\top \hat{\beta})$. This explains the similarities in the regularization paths.

5.2 Transition into the interpolation regime

The discussion above hints at the similarities between adversarial training and parameter-shrinking methods. Interestingly, Lasso and ridge regression are also connected to minimum-norm interpolators. The ridge regression solution converges to the minimum-norm solution as the parameter vanishes




Figure 4: **Training mean squared error vs inverse adversarial radius.** *Left:* for ridge and ℓ_2 -adversarial training. *Right:* for Lasso and ℓ_∞ -adversarial training. The error bars give the median and the 0.25 and 0.75 quantiles from 5 realizations. The vertical black lines show $\bar{\delta}$ in (5). Figure S.9 (appendix) shows the test MSE. For ridge regression or Lasso (see Section 3), the x -axis should be read as $1/\lambda$, rather than $1/\delta$. We generate the data synthetically using an isotropic Gaussian feature model (see Section 7) with $n = 60$ training data points and $p = 200$ features.

i.e., $\hat{\beta}^{\text{ridge}}(\lambda) \rightarrow \hat{\beta}^{\text{min-}\ell_2}$ as $\lambda \rightarrow 0^+$. Similarly, there is a relation between the minimum ℓ_1 -norm solution and Lasso. The relation requires additional constraints because, for the overparameterized case, Lasso does not necessarily have a unique solution. Nonetheless, it is proved in [41, Lemma 7] that the Lasso solution by the LARS algorithm satisfies $\hat{\beta}^{\text{lasso}}(\lambda) \rightarrow \hat{\beta}^{\text{min-}\ell_1}$ as $\lambda \rightarrow 0^+$.

For a sufficiently small δ , the solution of ℓ_2 -adversarial training equals the minimum ℓ_2 -norm solution; and, the solution of ℓ_∞ -adversarial training equals the minimum ℓ_1 -norm solution. There is a notable difference though: while this happens *only in the limit* for ridge regression and Lasso, for adversarial training this happens *for all values δ smaller than the threshold $\bar{\delta}$* . We illustrate this phenomenon next. Unlike ridge regression and Lasso which converge towards the interpolation solution, adversarial training goes through abrupt transitions and suddenly starts to interpolate the data. Figure 4 illustrates this phenomenon in synthetically generated data (with isotropic Gaussian feature, see Section 7).

Intuitive explanation for abrupt transitions. To get insight into why the abrupt transition occurs we present the case where $n = 1$, i.e., the dataset has a single data point (\mathbf{x}, y) , where $y = 1$ and $\|\mathbf{x}\|_2 = 1$. We can write $R^{\text{adv}}(\beta; \delta, \|\cdot\|_2) = (\mathcal{L}(\beta))^2$, where

$$\mathcal{L}(\beta) = |y - \mathbf{x}^\top \beta| + \delta \|\beta\|_2.$$

The function \mathcal{L} is necessarily minimized along the subspace spanned by the vector \mathbf{x} . Now, along this line, the function \mathcal{L} is piecewise linear with three segments, see Figure 5. One of the three segments changes the slope sign as δ decreases and the minimum of the function changes abruptly, thus explaining why abrupt transitions occur.




Figure 5: **Function $\mathcal{L}(\beta)$ for $\beta = \alpha \mathbf{x}$.**

6 Relation to robust regression and square-root Lasso

In this section we establish that both adversarial training and square-root Lasso can be integrated into the robust regression framework. We further investigate the similarities between the methods by analyzing the prediction error upper bounds.

Robust regression framework. Robust linear regression considers the minimization of the following cost function

$$R^{\text{robust}}(\beta; \mathcal{S}) = \max_{\Delta \in \mathcal{S}} \|\mathbf{y} - (\mathbf{X} + \Delta)\beta\|_2, \quad (9)$$

where the disturbance matrix Δ is constrained to belong to the ‘disturbance set’ \mathcal{S} . We connect robust regression, square-root Lasso and adversarial training. The next proposition gives the equivalence between robust regression for a row-bounded disturbance sets $\mathcal{R}_{p,\delta}$ and ℓ_p -adversarial training, see the appendix for the proof.

Proposition 5. For a disturbance set with the rows bounded by δ :

$$\mathcal{R}_{p,\delta} = \left\{ \begin{bmatrix} \text{---} & \Delta \mathbf{x}_1 & \text{---} \\ & \vdots & \\ \text{---} & \Delta \mathbf{x}_n & \text{---} \end{bmatrix} : \|\Delta \mathbf{x}_i\| \leq \delta, \forall i \right\},$$

we have that $\arg \min_{\beta} R^{\text{robust}}(\beta, \mathcal{R}_{p,\delta}) = \arg \min_{\beta} R^{\text{adv}}(\beta; \delta, \|\cdot\|)$.

On the other hand, there is an equivalence between square-root Lasso and robust regression for column-bounded disturbance sets $\mathcal{C}_{2,\delta}$. This was established by Xu *et al.* [33, Theorem 1] and we repeat it below.

Proposition 6. [33] For a disturbance set with columns bounded by δ :

$$\mathcal{C}_{2,\delta} = \left\{ \begin{bmatrix} \left| \right| & & \left| \right| \\ \zeta_1 & \cdots & \zeta_m \\ \left| \right| & & \left| \right| \end{bmatrix} : \|\zeta_i\|_2 \leq \delta, \forall i \right\},$$

we have that $\arg \min_{\beta} R^{\text{robust}}(\beta; \mathcal{C}_{2,\delta}) = \arg \min_{\beta} R^{\sqrt{\text{lasso}}}(\beta, \delta)$.

The above discussion hints at the similarities between adversarial training and square-root Lasso, as instances of robust regression under different constraints. Square-root Lasso minimizes $R^{\sqrt{\text{lasso}}}(\beta, \lambda) = n^{-1/2} \|\mathbf{y} - \mathbf{X}\beta\|_2 + \lambda \|\beta\|_1$ (See Section 3). The main motivation for the square-root Lasso is that it is a ‘pivotal’ method for sparse recovery: that is, it attains near-oracle performance without knowledge of the variance levels to set the regularization parameter [32]. As we will show in the next section, a similar property applies to ℓ_∞ -adversarial training.

Fixed-design analysis and similarities with square-root Lasso. In this section, we assume that the data was generated as: $y_i = \mathbf{x}_i^\top \beta^* + \varepsilon_i$ where β^* is the parameter vector used to generate the data. Under these assumptions, we can derive an upper bound for the (in-sample) prediction error:

Theorem 2. Let $\delta > \delta^* = 3 \frac{\|\mathbf{X}^\top \varepsilon\|_\infty}{\|\varepsilon\|_1}$, the prediction error of ℓ_∞ -adversarial training satisfies the bound:

$$\frac{1}{n} \|\mathbf{X}(\hat{\beta} - \beta^*)\|_2^2 \leq 8\delta \|\beta^*\|_1 \left(\frac{1}{n} \|\varepsilon\|_1 + 10\delta \|\beta^*\|_1 \right). \quad (10)$$

For comparison, we also provide the result for Lasso: (Adapted from [31, Thm. 7.13, p. 210])

Theorem 3. [31] Let $\lambda > \lambda^* = 3 \|\frac{\mathbf{X}^\top \varepsilon}{n}\|_\infty$, the prediction error of Lasso satisfies the bound:

$$\frac{1}{n} \|\mathbf{X}(\hat{\beta} - \beta^*)\|_2^2 \leq 8\lambda \|\beta^*\|_1. \quad (11)$$

Without additional constraints (see Remark 5) it can be shown that for Lasso it is not possible to improve the above bound. To satisfy this bound, however, Lasso requires knowledge of the magnitude of the noise ε . In Theorem 3, $\lambda^* = \frac{3}{n} \|\mathbf{X}^\top \varepsilon\|_\infty$. Hence, if we rescale ε (i.e., $\varepsilon \rightarrow \eta\varepsilon$) then a correspondent change in magnitude follows in λ^* (i.e., $\lambda^* \rightarrow \eta\lambda^*$). Square-root Lasso [32] avoids this problem and achieves a similar rate even without knowing the variance: i.e., it is a ‘pivotal’ method. Our method has similar properties and allows us to set δ without estimating the variance. This can be seen in Theorem 2, where re-scaling ε does not alter the value of $\delta^* = 3 \|\mathbf{X}^\top \varepsilon\|_\infty / \|\varepsilon\|_1$ because it affects the numerator and denominator simultaneously.

For instance, if we assume ε has i.i.d. $\mathcal{N}(0, \sigma^2)$ entries and that the matrix \mathbf{X} is fixed with $\max_{j=1, \dots, m} \|\mathbf{x}_j\|_\infty \leq M$. For $\lambda \propto M\sigma\sqrt{(\log p)/n}$, we (with high-probability) satisfy the condition in Theorem 3, obtaining: $\frac{1}{n} \|\mathbf{X}(\hat{\beta} - \beta^*)\|_2^2 \lesssim M\sigma\sqrt{(\log p)/n}$. For adversarial training, we can set: $\delta \propto M\sqrt{(\log p)/n}$, and (with high-probability) satisfy the theorem condition, obtaining the same bound. Notice that the choice of δ is not dependent on σ . We provide a full analysis in Appendix C.3.

Remark 4 (On the relation between $\bar{\delta}$ and δ^*). For sufficiently large n , we have $\delta^* > \bar{\delta}$ in the scenario above—the noise ε has i.i.d. normal entries $\mathcal{N}(0, \sigma^2)$ and the matrix \mathbf{X} is fixed. We can prove it by contradiction: if $\delta^* \leq \bar{\delta}$ then we could apply the bound from Theorem 2 to the minimum ℓ_1 -norm

interpolator (due to Theorem 1). Hence, $\sigma^2 \approx \frac{1}{n} \|\varepsilon\|_2^2 = \frac{1}{n} \|\mathbf{X}(\hat{\beta} - \beta^*)\|_2^2 \lesssim M\sigma\sqrt{(\log p)/n}$ and we can always choose a sufficiently large value of n for which the inequality is false.

Remark 5. In the original square-root Lasso paper [32], a bound is obtained for the ℓ_2 parameter distance: $\|\hat{\beta} - \beta^*\|_2^2$ for the case \mathbf{X} satisfy the restricted eigenvalue condition and β^* is sparse. Under these more strict assumptions (restricted eigenvalue condition) we can also obtain a faster convergence rate for the prediction error. Here we focus on the fixed-design prediction error without additional assumptions on \mathbf{X} . For other predictors, these other proofs follow similar steps and yield similar requirements on λ^* , see [31, Chapter 7]. However, we leave these and other analysis (such as variable selection consistency) of adversarial training for future work.

7 Numerical Experiments

We study five different examples. The main goal is to experimentally confirm our theoretical findings. For each scenario, we compute and plot: train and test MSE for different choices of p , n and δ . For comparison purposes, we also compute and plot train and test MSE for Lasso and ridge regression and minimum-norm interpolators. Finally, in line with our discussion in Section 4, we compute and plot $\bar{\delta}$ as a function of p/n . In all the numerical examples the adversarial training solution is implemented by minimizing (2) using CVXPY [42]. The scenarios under consideration are described below (see Appendix D for additional details)

1. **Isotropic Gaussian feature model.** The output is a linear combination of the features plus additive noise: $y_i = \mathbf{x}_i^\top \beta + \epsilon_i$, for Gaussian noise and covariates: $\epsilon_i \sim \mathcal{N}(0, \sigma^2)$ and $\mathbf{x}_i \sim \mathcal{N}(0, \mathbf{I}_p)$.
2. **Latent-space feature model** [26, Section 5.4]. The features \mathbf{x} are noisy observations of a lower-dimensional subspace of dimension d . A vector in this *latent space* is represented by $\mathbf{z} \in \mathbb{R}^d$. $\mathbf{x} = \mathbf{W}\mathbf{z} + \mathbf{u}$. The output is a linear combination of the latent space plus noise.
3. **Random Fourier features model** [43]. We apply a random Fourier feature map to inputs of the Diabetes dataset [18]. Random fourier features are obtained by the transformation $\mathbf{x}_i = \sqrt{2/m} \cos(\mathbf{W}\mathbf{z}_i + \mathbf{b})$ where the entries of \mathbf{W} and \mathbf{b} are independently sampled. It can be seen as a one-layer untrained neural network and approximates the Gaussian kernel feature map [43].
4. **Random projection model** [40]. The data is generated as in the isotropic Gaussian feature model, but we only observe $\mathbf{S}\mathbf{x}_i$, i.e., a projection of the inputs.
5. **Phenotype prediction from genotype.** We illustrate our method on the Diverse MAGIC wheat dataset [44] from the National Institute for Applied Botany. We use a subset of the genotype to predict one of the continuous phenotypes.

Results. We experimentally corroborate Theorem 1 in all scenarios: In Figures 4 and S.8 (Appendix) we show the train MSE as we change the adversarial radius δ , confirming the abrupt transitions into the interpolation regime and also that $\bar{\delta}$ is the transition point. In Figures S.9 and S.10 we show the corresponding test MSE without and in the presence of an adversary, respectively. Interestingly, the adversarial radius that yields the best results is not always equal to the radius δ_{test} the model will be evaluated on. Figures 2 and S.5 (Appendix) display $\bar{\delta}$ as a function of the ratio p/n . We observe that $\bar{\delta}/\mathbb{E}[\|\mathbf{x}\|]$ is growing in all examples and we would still expect improved robustness in light of Proposition 2. The Random Fourier features model is the only case where $\bar{\delta}$ seems to decrease (in absolute value). However, $\mathbb{E}[\|\mathbf{x}\|_1]$ is also decreasing (and at a faster rate). Figures S.6 and S.7 give test MSE without and in the presence of an adversary for the minimum-norm interpolator as a function of the ratio p/n , confirming improved adversarial robustness as p grows.

Figure 6 provides a comparison of the test error of the different methods under study. For Lasso, ridge and adversarial training, we use the best δ or λ available for each method (obtained via grid search). We note that while for $p = 1000$ optimally tuned Lasso and ℓ_∞ -adversarial training significantly outperform the corresponding minimum ℓ_1 -norm interpolator. As p increases, the performance of the three different methods becomes quite similar. It is also an example where the minimum ℓ_1 -norm outperforms the minimum ℓ_2 -norm interpolator (for large p). In the same setup, Figure S.12

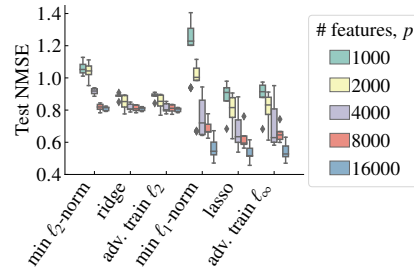


Figure 6: **Test normalized MSE (NMSE)** in the MAGIC dataset.

study a choice of adversarial radius inspired by Theorem 2. We use $\delta \propto \|\mathbf{x}\xi\|_\infty/\|\xi\|_1$ for ξ a vector with zero-mean normal entries. We use a random ξ , since we do not know the true additive noise. Even with this approximation, ℓ_∞ -adversarial training performs comparably with Lasso with the regularization parameter set using 5-fold cross-validation doing a full search in the hyperparameter space. The figure also provides a comparison with square-root Lasso under a similar setting.

8 Results for general loss functions

The following theorem can be used to generalize Proposition 1.

Theorem 4. Let $\mathcal{L} : \mathbb{R} \rightarrow \mathbb{R}$ be a convex and lower-semicontinuous function, for every $\delta \geq 0$,

$$\max_{\|\Delta\mathbf{x}\| \leq \delta} \mathcal{L}((\mathbf{x} + \Delta\mathbf{x})^\top \boldsymbol{\beta}) = \max_{s \in \{-1, 1\}} \mathcal{L}(\mathbf{x}^\top \boldsymbol{\beta} + \delta s \|\boldsymbol{\beta}\|_*). \quad (12)$$

We can find a closed-formula expression for $s_* = \arg \max_{s \in \{-1, 1\}} \mathcal{L}(\mathbf{x}^\top \boldsymbol{\beta} + \delta s \|\boldsymbol{\beta}\|_*)$ in many cases of interest. For **regression problems**, given any non-decreasing function $\ell : \mathbb{R}^+ \rightarrow \mathbb{R}^+$, and $\mathcal{L}(\mathbf{x}^\top \boldsymbol{\beta}) = \ell(|\mathbf{x}^\top \boldsymbol{\beta} - y|)$. If ℓ is lower semicontinuous and convex then so will \mathcal{L} and the result of the theorem holds (the squared loss $\ell(z) = z^2$ is a special case). Then $s_* = -\text{sign}(y - \mathbf{x}^\top \boldsymbol{\beta})$ and

$$\max_{\|\Delta\mathbf{x}\| \leq \delta} \ell(|(\mathbf{x} + \Delta\mathbf{x})^\top \boldsymbol{\beta} - y|) = \ell(|y - \mathbf{x}^\top \boldsymbol{\beta}| + \delta \|\boldsymbol{\beta}\|_*).$$

For **classification**, let $y \in \{-1, 1\}$ and $\mathcal{L}(\mathbf{x}^\top \boldsymbol{\beta}) = \ell(y(\mathbf{x}^\top \boldsymbol{\beta}))$ where ℓ is a non-increasing function. If ℓ is lower semicontinuous and convex then so will \mathcal{L} and the result of the theorem holds. Then $s_* = -y$ and we obtain:

$$\max_{\|\Delta\mathbf{x}\| \leq \delta} \ell(y((\mathbf{x} + \Delta\mathbf{x})^\top \boldsymbol{\beta})) = \ell(y(\mathbf{x}^\top \boldsymbol{\beta}) - \delta \|\boldsymbol{\beta}\|_*).$$

The above result can also be applied to unsupervised learning. In the Appendix, we illustrate how Theorem 4 can be used for **dimensionality reduction**. We consider the problem of finding \mathbf{P} that minimizes the reconstruction error $\|\mathbf{x} + \mathbf{P}\mathbf{P}^\top \mathbf{x}\|_2^2$ (that yields PCA algorithm) and derive an adversarial version of it.

Remark 6. Over this paper we consider $\mathbf{x}, \boldsymbol{\beta} \in \mathbb{R}^p$, but Theorem 4 holds generally for $\mathbf{x} \in \mathcal{X}$ a vector in a Banach space endowed with norm $\|\cdot\|$ and $\boldsymbol{\beta} \in \mathcal{X}^*$ a continuous linear map $\boldsymbol{\beta} : \mathcal{X} \rightarrow \mathbb{R}$. Just define $\mathbf{x}^\top \boldsymbol{\beta} := \boldsymbol{\beta}(\mathbf{x})$ and take $\|\cdot\|_*$ to be the norm of the dual space \mathcal{X}^* .

9 Conclusion

We study adversarial training in linear regression. Adversarial training allows us to depart from the traditional regularization setting where we can decompose the loss and the regularization terms, minimizing $\frac{1}{n} \sum_{i=0}^n \ell(\mathbf{x}_i^\top \boldsymbol{\beta}, y_i) + \Omega(\boldsymbol{\beta})$ for some penalization function Ω . We show how it provides new insights into the minimum-norm interpolator, by proving a new equivalence between the two methods (Section 4). While adversarially-trained linear regression arrives at similar solutions to parameter shrinking methods in some scenarios (Section 5), it also has some advantages, for instance, the adversarial radius might be set without knowing the noise variance (Section 6). Unlike, squared-root Lasso it achieves this while still minimizing a sum of squared terms. We believe a natural next step is to provide a tailored solver that can allow for efficient solutions for models with many features, rendering adversarially-trained linear regression useful, for instance, in modeling genetics data (see one minimal example of phenotype prediction from the genotype in Section 7).

Another interesting direction is generalizing our results to classification and nonlinear models. Indeed, adversarial training is very often used in the context of neural networks and it is natural to be interested in the behavior of this class of models. Our work allows for the analysis of simplified theoretical models commonly used to study neural networks. Random Fourier feature models can be analyzed using our theory and are they studied in the numerical examples. These models can be seen as a simplified neural network model (one-layer neural network where only the last layer weights are adjusted). In Appendix A.4, we provide adversarial training after linear projections, and could, for instance, be used to analyze deep linear networks, as the ones studied in [45], [46]. Finally, Section 8 provides a result for infinite spaces, and one could attempt to use it to analyze kernel methods and infinitely-wide neural networks [47], [48]. Overall we believe that our results provide an interesting set of tools also for the analysis of nonlinear models, and could provide insight into the empirically observed interplay (see [49]) between robustness and regularization in deep neural networks.

Acknowledgments

The authors would like to thank Carl Nettelblad for very fruitful discussions throughout the work with this research. We also thank Daniel Gedon and Dominik Baumann for their feedback on the first version of the manuscript.

TBS and DZ are financially supported by the Swedish Research Council, with the projects *Deep probabilistic regression – new models and learning algorithms* (contract number: 2021-04301) and *Counterfactual Prediction Methods for Heterogeneous Populations* (contract number:2018- 05040) and by the *Wallenberg AI, Autonomous Systems and Software Program (WASP)* funded by Knut and Alice Wallenberg Foundation. TBS, AHR and DZ by the *Kjell och Märta Beijer Foundation*. FB by the Agence Nationale de la Recherche as part of the “Investissements d’avenir” program, reference ANR-19-P3IA0001 (PRAIRIE 3IA Institute); and by the European Research Council (grant SEQUOIA 724063).

The research was partially conducted during AHR’s research visit to INRIA. The visit was financially supported by the Institut Français de Suède through the SFVE-A mobility program; and, by the European Network of AI Excellence Centres through ELISE mobility program.

References

- [1] J. Bruna, C. Szegedy, I. Sutskever, I. Goodfellow, W. Zaremba, R. Fergus, and D. Erhan, “Intriguing properties of neural networks,” in *International Conference on Learning Representations (ICLR)*, 2014.
- [2] I. J. Goodfellow, J. Shlens, and C. Szegedy, “Explaining and Harnessing Adversarial Examples,” in *International Conference on Learning Representations (ICLR)*, 2015.
- [3] A. Kurakin, I. Goodfellow, S. Bengio, *et al.*, “Adversarial attacks and defences competition,” in *The NIPS ’17 competition: Building Intelligent Systems*, 2018, pp. 195–231.
- [4] A. Fawzi, O. Fawzi, and P. Frossard, “Analysis of classifiers’ robustness to adversarial perturbations,” *Machine Learning*, vol. 107, no. 3, pp. 481–508, 2018.
- [5] A. Ilyas, S. Santurkar, D. Tsipras, L. Engstrom, B. Tran, and A. Madry, “Adversarial Examples Are Not Bugs, They Are Features,” *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [6] X. Yuan, P. He, Q. Zhu, and X. Li, “Adversarial examples: Attacks and defenses for deep learning,” *IEEE Transactions on Neural Networks and Learning Systems*, vol. 30, no. 9, pp. 2805–2824, 2019.
- [7] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, “Towards Deep Learning Models Resistant to Adversarial Attacks,” *International Conference for Learning Representations (ICLR)*, 2018.
- [8] R. Huang, B. Xu, D. Schuurmans, and C. Szepesvari, “Learning with a Strong Adversary,” *arXiv:1511.03034*, 2016.
- [9] T. Bai, J. Luo, J. Zhao, B. Wen, and Q. Wang, “Recent Advances in Adversarial Training for Adversarial Robustness,” *arXiv:2102.01356*, 2021.
- [10] F. Croce, M. Andriushchenko, V. Sehwag, E. Debenedetti, N. Flammarion, M. Chiang, P. Mittal, and M. Hein, “RobustBench: A standardized adversarial robustness benchmark,” in *NeurIPS Datasets and Benchmarks track*, 2021.
- [11] H. Taheri, R. Pedarsani, and C. Thrampoulidis, “Asymptotic Behavior of Adversarial Training in Binary Classification,” *IEEE International Symposium on Information Theory (ISIT)*, vol. 127-132, 2022.
- [12] A. Javanmard, M. Soltanolkotabi, and H. Hassani, “Precise tradeoffs in adversarial training for linear regression,” in *Proceedings of the Conference on Learning Theory*, vol. 125, 2020, pp. 2034–2078.
- [13] H. Hassani and A. Javanmard, “The curse of overparametrization in adversarial training: Precise analysis of robust generalization for random features regression,” *arXiv:2201.05149*, 2022.
- [14] Y. Min, L. Chen, and A. Karbasi, “The Curious Case of Adversarially Robust Models: More Data Can Help, Double Descend, or Hurt Generalization,” *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, vol. 161, pp. 129–139, 2021.
- [15] D. Yin, R. Kannan, and P. Bartlett, “Rademacher Complexity for Adversarially Robust Generalization,” in *Proceeding of the International Conference on Machine Learning*, 2019, pp. 7085–7094.
- [16] D. Tsipras, S. Santurkar, L. Engstrom, A. Turner, and A. Ma, “Robustness May Be At Odds with Accuracy,” *International Conference for Learning Representations*, 2019.
- [17] A. H. Ribeiro and T. B. Schön, “Overparameterized Linear Regression under Adversarial Attacks,” *IEEE Transactions on Signal Processing*, 2023.

- [18] B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani, “Least angle regression,” *The Annals of Statistics*, vol. 32, no. 2, pp. 407–499, 2004.
- [19] T. Hastie, R. Tibshirani, and J. Friedman, *Elements of Statistical Learning*, Second. Springer Science & Business Media, 2009.
- [20] C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals, “Understanding deep learning requires rethinking generalization,” in *International Conference on Learning Representations*, 2017.
- [21] P. L. Bartlett, A. Montanari, and A. Rakhlin, “Deep learning: A statistical viewpoint,” *arXiv:2103.09177*, 2021.
- [22] P. L. Bartlett, P. M. Long, G. Lugosi, and A. Tsigler, “Benign overfitting in linear regression,” *Proceedings of the National Academy of Sciences*, vol. 117, no. 48, pp. 30 063–30 070, 2020.
- [23] F. Koehler, L. Zhou, D. J. Sutherland, and N. Srebro, “Uniform Convergence of Interpolators: Gaussian Width, Norm Bounds and Benign Overfitting,” in *NeurIPS*, 2021.
- [24] M. Belkin, D. Hsu, S. Ma, and S. Mandal, “Reconciling modern machine-learning practice and the classical bias–variance trade-off,” *Proceedings of the National Academy of Sciences*, vol. 116, no. 32, pp. 15 849–15 854, 2019.
- [25] M. Belkin, D. Hsu, and J. Xu, “Two Models of Double Descent for Weak Features,” *SIAM Journal on Mathematics of Data Science*, vol. 2, no. 4, pp. 1167–1180, 2020.
- [26] T. Hastie, A. Montanari, S. Rosset, and R. J. Tibshirani, “Surprises in high-dimensional ridgeless least squares interpolation,” *The Annals of Statistics*, vol. 50, no. 2, pp. 949–986, 2022.
- [27] C. Dan, Y. Wei, and P. Ravikumar, “Sharp statistical guarantees for adversarially robust Gaussian classification,” in *Proceedings of the 37th international conference on machine learning*, ser. Proceedings of machine learning research, vol. 119, PMLR, 2020, pp. 2345–2355.
- [28] E. Dobriban, H. Hassani, D. Hong, and A. Robey, *Provable tradeoffs in adversarially robust classification*, 2022.
- [29] A. Javanmard and M. Soltanolkotabi, “Precise statistical analysis of classification accuracies for adversarial training,” *The Annals of Statistics*, vol. 50, no. 4, pp. 2127–2156, 2022.
- [30] R. Tibshirani, “Regression shrinkage and selection via the LASSO,” *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 267–288, 1996.
- [31] M. J. Wainwright, *High-Dimensional Statistics: a non-Asymptotic Viewpoint*, ser. Cambridge series on statistical and probabilistic mathematics 48. Cambridge University Press, 2019.
- [32] A. Belloni, V. Chernozhukov, and L. Wang, “Square-Root Lasso: Pivotal Recovery of Sparse Signals via Conic Programming,” *Biometrika*, vol. 98, no. 4, pp. 791–806, 2011.
- [33] H. Xu, C. Caramanis, and S. Mannor, “Robust regression and lasso,” *Advances in Neural Information Processing Systems*, vol. 21, 2008.
- [34] —, “Robustness and regularization of support vector machines,” *Journal of Machine Learning Research*, vol. 10, no. 51, pp. 1485–1510, 2009.
- [35] Y. Xing, Q. Song, and G. Cheng, “On the Generalization Properties of Adversarial Training,” in *Proceedings of the International Conference on Artificial Intelligence and Statistics*, 2021, pp. 505–513.
- [36] D. P. Bertsekas, A. Nedi, and A. E. Ozdaglar, *Convex Analysis and Optimization*. 2003.
- [37] F. H. Clarke, *Optimization and Nonsmooth Analysis*. Society for Industrial and Applied Mathematics, 1990.
- [38] S. Boyd, J. Duchi, M. Pilanci, and L. Vandenberghe, “Subgradients (Lecture Notes),” Tech. Rep. Notes for EE364b, Stanford University, Spring 2021-22, 2022.
- [39] S. S. Chen, D. L. Donoho, and M. A. Saunders, “Atomic decomposition by basis pursuit,” *SIAM Journal on Scientific Computing*, vol. 20, no. 1, pp. 33–61, 1998.
- [40] F. Bach, “High-dimensional analysis of double descent for linear regression with random projections,” *arXiv:2303.01372*, 2023.
- [41] R. J. Tibshirani, “The Lasso Problem and Uniqueness,” *Electronic Journal of Statistics*, vol. 7, pp. 1456–1490, 2013.
- [42] S. Diamond and S. Boyd, “CVXPY: A Python-embedded modeling language for convex optimization,” *Journal of Machine Learning Research*, vol. 17, no. 83, pp. 1–5, 2016.
- [43] A. Rahimi and B. Recht, “Random Features for Large-Scale Kernel Machines,” in *Advances in Neural Information Processing Systems 20*, 2008, pp. 1177–1184.
- [44] M. F. Scott, N. Fradgley, A. R. Bentley, *et al.*, “Limited haplotype diversity underlies polygenic trait architecture across 70 years of wheat breeding,” *Genome Biology*, vol. 22, no. 1, p. 137, 2021.

- [45] A. M. Saxe, J. L. McClelland, and S. Ganguli, “Exact solutions to the nonlinear dynamics of learning in deep linear neural networks,” *International Conference on Learning Representations (ICLR)*, 2014.
- [46] T. Laurent and J. Brecht, “Deep Linear Networks with Arbitrary Loss: All Local Minima Are Global,” in *Proceedings of the International Conference on Machine Learning*, 2018, pp. 2902–2907.
- [47] J. Lee, L. Xiao, S. S. Schoenholz, Y. Bahri, R. Novak, J. Sohl-Dickstein, and J. Pennington, “Wide neural networks of any depth evolve as linear models under gradient descent,” *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2020, no. 12, p. 124 002, 2020.
- [48] A. Jacot, F. Gabriel, and C. Hongler, “Neural Tangent Kernel: Convergence and Generalization in Neural Networks,” *Advances in Neural Information Processing Systems 31*, 2018.
- [49] D. Jakubovitz and R. Giryes, “Improving DNN robustness to adversarial attacks using jacobian regularization,” in *European conference on computer vision*, 2018.