

---

# Robust Linear Regression: Gradient-descent, Early-stopping, and Beyond

---

Meyer Scetbon

Facebook AI Research

Elvis Dohmatob

## Abstract

In this work we study the robustness to adversarial attacks, of early-stopping strategies on gradient-descent (GD) methods for linear regression. More precisely, we show that early-stopped GD is optimally robust (up to an absolute constant) against Euclidean-norm adversarial attacks. However, we show that this strategy can be arbitrarily sub-optimal in the case of general Mahalanobis attacks. This observation is compatible with recent findings in the case of classification Vardi et al. (2022) that show that GD provably converges to non-robust models. To alleviate this issue, we propose to apply instead a GD scheme on a transformation of the data adapted to the attack. This data transformation amounts to apply feature-depending learning rates and we show that this modified GD is able to handle any Mahalanobis attack, as well as more general attacks under some conditions. Unfortunately, choosing such adapted transformations can be hard for general attacks. To the rescue, we design a simple and tractable estimator whose adversarial risk is optimal up to within a multiplicative constant of 1.1124 in the population regime, and works for any norm.

## 1 Introduction

Machine learning models are highly sensitive to small perturbations known as *adversarial examples* (Szegedy et al., 2013), which are often imperceptible by humans. While various strategies such as adversarial training (Madry et al., 2018) can mitigate this vulnerability empirically, the situation remains highly problematic for many safety-critical applications like autonomous vehicles or health, and motivates a better theoretical understanding of what mechanisms may be causing this.

---

Proceedings of the 26<sup>th</sup> International Conference on Artificial Intelligence and Statistics (AISTATS) 2023, Valencia, Spain. PMLR: Volume 206. Copyright 2023 by the author(s).

From a theoretical perspective, the case of classification is rather well-understood. Indeed, the hardness of classification under test-time adversarial attacks has been crisply characterized (Bhagoji et al., 2019; Bubeck et al., 2018). In the special case of linear classification, explicit lower-bounds have been obtained (Schmidt et al., 2018; Bhattacharjee et al., 2021). However, the case of regression is relatively understudied. Recently, Xing et al. (2021) have initiated a theoretical study of linear regression under Euclidean attacks, where an adversary is allowed to attack the input data point at test time. The authors proposed a two-stage estimator and proved its consistency. The optimal estimator obtained in (Xing et al., 2021) corresponds to a ridge shrinkage (i.e  $\ell_2$  penalization).

In this paper, we consider linear regression under adversarial test-time attacks w.r.t arbitrary norms (not just Euclidean /  $\ell_2$ -norms as in (Xing et al., 2021)), and analyze the robustness of gradient-descent (GD) along the entire optimization path. By doing so we observe that GD might fail to capture a robust predictor along its path especially in the case of non-Euclidean attacks. We propose a variant of the GD scheme where an adapted transformation on the data is performed before applying a GD scheme. This allows us to understand the effect on robustness of early-stopping strategies (not training till the end) for general attacks. Finally we design a generic algorithm able to produce a robust predictor against any norm-attacks.

### 1.1 Summary of main contributions

Our main contributions are summarized as follows.

– **Case of gradient-descent (GD).** In Proposition 4.2, we show that early-stopped GD achieves near-optimal adversarial risk in case of Euclidean attacks. Early-stopping is crucial because the predictor obtained by running GD till the end can be arbitrarily sub-optimal in terms of adversarial risk (Proposition 4.1). Contrasting with Proposition 4.2, we show in Proposition 4.4 that early-stopped GD can be arbitrarily sub-optimal in the non-Euclidean case, e.g when the attacker’s norm is a Mahalanobis norm. Thus, GD, along its entire optimization path, can fail to find robust model in general.

– **An Adapted GD scheme (GD+).** We propose a mod-

ified version of GD, termed GD+, in which the dynamics are forced to be non-uniform across different features and exhibit different regimes where early-stopped GD+ can achieve near-optimal adversarial risk. More precisely, (i) we show that it achieves near-optimal adversarial risk in the case of Mahalanobis norm attacks (Proposition 5.1), (ii) we also prove that it is near-optimally robust under  $\ell_p$ -norm attacks as soon as the features are uncorrelated (Theorem 5.2) and (iii) we study the robustness along the entire optimization path of GD+ in the case of general norm-attacks. In particular, we provide a sufficient condition on the model such that early-stopped GD+ achieves near-optimal adversarial risk under general norm-attacks in Theorem 5.1 and we show that when this condition is not satisfied, GD+ can be arbitrarily sub-optimal (Proposition 5.3).

– **A two-stage estimator for the general case.** Finally, we propose a simple two-stage algorithm (Algorithm 1) which works for arbitrary norm-attacks, and achieves optimal adversarial risk up to within a multiplicative constant factor in the population regime. Consistency and statistical guarantees for the proposed estimator are also provided.

## 1.2 Related Work

The theoretical understanding of adversarial examples is now an active area of research. Below is a list of works which are most relevant to our current letter.

Tsipras et al. (2019) considers a specific data distribution where good accuracy implies poor robustness. (Shafahi et al., 2018; Mahloujifar et al., 2018; Gilmer et al., 2018; Dohmatob, 2019) show that for high-dimensional data distributions which have concentration property (e.g., multivariate Gaussians, distributions satisfying log-Sobolev inequalities, etc.), an imperfect classifier will admit adversarial examples. Dobriban et al. (2020) studies tradeoffs in Gaussian mixture classification problems, highlighting the impact of class imbalance. On the other hand, Yang et al. (2020) observed empirically that natural images are well-separated, and so locally-lipschitz classifiers shouldn't suffer any kind of test error vs robustness tradeoff.

In the context of linear classification (Schmidt et al., 2018; Bubeck et al., 2018; Khim and Loh, 2018; Yin et al., 2019; Bhattacharjee et al., 2021; Min et al., 2021a,b), established results show a clear gap between learning in ordinary and adversarial settings. Li et al. (2020) studies the dynamics of linear classification on separable data, with exponential-tail losses. The authors show that GD converges to separator of the dataset, which is minimal w.r.t to a norm which is an interpolation between the  $\ell_2$  norm (reminiscent of normal learning), and  $\ell_q$ -norm, where  $q$  is the harmonic conjugate of the attacker's norm. Vardi et al. (2022) showed that on two-layer neural networks, gradient-descent with exponential-tail loss function converges to weights which are vulnerable to adversarial examples.

Javanmard et al. (2020) study tradeoffs between ordinary and adversarial risk in linear regression, and computed exact Pareto optimal curves. Javanmard and Mehrabi (2021) also revisit this tradeoff for latent models and show that this tradeoff is mitigated when the data enjoys a low-dimensional structure. Dohmatob (2021); Hassani and Javanmard (2022) study the tradeoffs between interpolation, normal risk, and adversarial risk, for finite-width overparameterized networks with linear target functions. Javanmard and Soltanolkotabi (2022) investigate the effect of adversarial training on the standard and adversarial risks and derive a precise characterization of them for a class of minimax adversarially trained models.

The work most related to ours is (Xing et al., 2021) which studied minimax estimation of linear models under adversarial attacks in Euclidean norm. They showed that the optimal robust linear model is a ridge estimator whose regularization parameter is a function of the population covariance matrix  $\Sigma$  and the generative linear model  $w_0$ . Since neither  $w_0$  nor  $\Sigma$  is known in practice, the authors proposed a two-stage estimator in which the first stage is consistent estimators for  $w_0$  and  $\Sigma$  and the second stage is solving the ridge problem. In (Xing et al., 2021), the authors cover only the case of Euclidean attacks while here, we extend the study of the adversarial risk in linear regression under general norm-attacks. More precisely, here we are interested in understanding the robustness of GD, for general attacks and along the entire optimization path. As a separate contribution, we also propose a new consistent two-stage estimator based on a "dualization" of the adversarial problem that can be applied for general attacks.

## 1.3 Outline of Manuscript

In Section 2, we present the problem setup, main definitions, and some preliminary computations. The adversarial risk of (early-stopped) GD in the infinite-sample regime is analyzed in Section 4; cases of optimality and sub-optimality of this scheme are characterized. In Section 5, GD+ (an improved version of GD) is proposed, and its adversarial risk in the population regime is studied. In Section 6, we consider the finite samples regime and we propose a simple two-stage estimator, which works for all attacker norms. Its adversarial risk is shown to be optimal up to within a multiplicative factor, and additive statistical estimation error due to finite samples.

## 2 Preliminaries

**Notations.** Let us introduce some basic notations. Additional technical notations are provided in the appendix.

$[d]$  denotes the set of integers from 1 to  $d$  inclusive. The maximum (resp. minimum) of two real numbers  $a$  and  $b$  will be denoted  $a \vee b$  (resp.  $a \wedge b$ ). The operator norm of

a matrix  $A$  is denoted  $\|A\|_{op}$  and corresponds to the positive square-root of the largest eigenvalue of  $AA^\top$ . Given a positive-definite (p.d.) matrix  $S \in \mathbb{R}^{d \times d}$  and a vector  $w \in \mathbb{R}^d$ , the Mahalanobis norm of  $w$  induced by  $S$  is defined by  $\|w\|_S := \sqrt{w^\top S w}$ . We denote respectively  $M_d(\mathbb{R})$  and  $S_d^{++}(\mathbb{R})$  the set of  $d \times d$  matrices and positive-definite matrices. Given  $p \in [1, \infty]$ , its harmonic conjugate, denoted  $q(p)$  is the unique  $q \in [1, \infty]$  such that  $1/p + 1/q = 1$ . For a given norm  $\|\cdot\|$ , we denote  $\|\cdot\|_*$  its dual norm, defined by

$$\|w\|_* := \sup\{x^\top w \mid x \in \mathbb{R}^d, \|x\| \leq 1\}, \quad (1)$$

Note that if the norm  $\|\cdot\|$  is an  $\ell_p$ -norm for some  $p \in [1, \infty]$ , then the dual norm is the  $\ell_q$ -norm, where  $q$  is the harmonic conjugate of  $p$ . For example, the dual of the euclidean norm (corresponding to  $p = 2$ ) is the euclidean norm itself, while the dual of the usual  $\ell_\infty$ -norm is the  $\ell_1$ -norm.

The unit-ball for a norm  $\|\cdot\|$  is denoted  $B_{\|\cdot\|}^d$ , and defined by  $B_{\|\cdot\|}^d := \{x \in \mathbb{R}^d \mid \|x\| \leq 1\}$ . Given any  $s \geq 0$ , set of  $s$ -sparse vectors in  $\mathbb{R}^d$  is denoted  $B_0^d$  and defined by

$$B_0^d(s) := \{w \in \mathbb{R}^d \mid \|w\|_0 \leq s\}. \quad (2)$$

Finally, define absolute constants  $c_0 := \sqrt{2/\pi}$ ,  $\alpha := 2/(1 + c_0) \approx 1.1124$  and  $\beta = 1.6862$ .

## 2.1 Problem Setup

Fix a vector  $w_0 \in \mathbb{R}^d$ , a positive definite matrix  $\Sigma$  of size  $d$ , and consider an i.i.d. dataset  $\mathcal{D}_n = \{(x_1, y_1), \dots, (x_n, y_n)\}$  of size  $n$ , given by

$$y_i = x_i^\top w_0 + \epsilon_i, \text{ for all } i \in [n] \quad (3)$$

where  $x_1, \dots, x_n$  i.i.d  $\sim N(0, \Sigma)$  and  $\epsilon_1, \dots, \epsilon_n$  i.i.d  $\sim N(0, \sigma_\epsilon^2)$  independent of the  $x_i$ 's. Thus, the distribution of the features is a centered multivariate Gaussian distribution with covariance matrix  $\Sigma$ , while  $w_0$  is the generative model.  $\sigma_\epsilon \geq 0$  measures the size of the noise. These assumptions on the model will be assumed along all the formal claims of the paper. We also refer to  $n$  as the *sample size*, and to  $d$  as the *input-dimension*.

In most of our analysis, except otherwise explicitly stated, we will consider the case of infinite-data regime  $n = \infty$  –or more generally  $n \gg d$ , which allows us to focus on the effects inherent to the data distribution (controlled by feature covariance matrix  $\Sigma$ ) and the inductive bias of the norm w.r.t which the attack is measured, while side-stepping issues due to finite samples and label noise. Also note that in this infinite-data setting, label noise provably has no influence on the learned model.

## 2.2 Adversarial Robustness Risk

Given a linear model  $w \in \mathbb{R}^d$ , an attacker is allowed to swap a clean test point  $x \sim N(0, \Sigma)$  with a corrupted ver-

sion  $x' = x + \delta$  thereof. The perturbation  $\delta = \delta(x) \in \mathbb{R}^d$  is constrained to be small: this is enforced by demanding that  $\|\delta\| \leq r$ , where  $\|\cdot\|$  is a specified norm and  $r \geq 0$  is the attack budget. One way to measure the performance of a linear model  $w \in \mathbb{R}^d$  under such attacks of size  $r$ , is via its so-called adversarial risk (Madry et al., 2018; Xing et al., 2021).

**Definition 2.1.** For any  $w \in \mathbb{R}^d$  and  $r \geq 0$ , define the adversarial risk of  $w$  at level  $r \geq 0$  as follows

$$E^{\|\cdot\|}(w, w_0, r) := \mathbb{E}_x \left[ \sup_{\|\delta\| \leq r} ((x + \delta)^\top w - x^\top w_0)^2 \right], \quad (4)$$

where  $x \sim N(0, \Sigma)$  is a random test point.

It is clear that  $r \mapsto E^{\|\cdot\|}(w, w_0, r)$  is a non-decreasing function and  $E^{\|\cdot\|}(w, w_0, 0)$  corresponds to the ordinary risk of  $w$ , namely

$$E(w, w_0) := \mathbb{E}_x [(x^\top w - x^\top w_0)^2] = \|w - w_0\|_\Sigma^2. \quad (5)$$

In classical regression setting, the aim is to find  $w$  which minimizes  $E(w, w_0)$ . In the adversarial setting studied here, the aim is to minimize  $E^{\|\cdot\|}(w, w_0, r)$  for any  $r \geq 0$ .

We will henceforth denote by  $E_{opt}^{\|\cdot\|}(w_0, r)$  the smallest possible adversarial risk of a linear model for  $\|\cdot\|$ -attacks of magnitude  $r$ , that is

$$E_{opt}^{\|\cdot\|}(w_0, r) := \inf_{w \in \mathbb{R}^d} E^{\|\cdot\|}(w, w_0, r). \quad (6)$$

We start with the following well-known elementary but useful lemma which proved in the supplemental. Also see Xing et al. (2021); Javanmard and Soltanolkotabi (2022) for the special case of Euclidean-norm attacks.

**Lemma 2.1.** Recall that  $c_0 = \sqrt{2/\pi}$ , then for any  $w \in \mathbb{R}^d$  and  $r \geq 0$ , it holds that

$$E^{\|\cdot\|}(w, w_0, r) = \|w - w_0\|_\Sigma^2 + r^2 \|w\|_*^2 + 2c_0 r \|w - w_0\|_\Sigma \|w\|_*. \quad (7)$$

The mysterious constant  $c_0 = \sqrt{2/\pi}$  in Lemma 2.1 corresponds to the expected absolute value of a standard Gaussian random variable. In order to obtain a robust predictor to adversarial attacks, one aims at minimizing the adversarial risk introduced in (4). However the objective function of the problem, even in the linear setting (7), is rather complicated to optimize due to its non-convexity.

## 3 A Proxy for Adversarial Risk

The following lemma will be one of the main workhorses in subsequent results, as it allows us to replace the adversarial risk functional  $E$  with a more tractable proxy  $\tilde{E}$ .

**Lemma 3.1.** For any  $w \in \mathbb{R}^d$  and  $r \geq 0$ , it holds that

$$E^{\|\cdot\|}(w, w_0, r) \leq \tilde{E}^{\|\cdot\|}(w, w_0, r) \leq \alpha \cdot E^{\|\cdot\|}(w, w_0, r), \quad (8)$$

where  $\alpha := 2/(1 + \sqrt{2/\pi}) \approx 1.1124$  and

$$\tilde{E}^{\|\cdot\|}(w, w_0, r) := (\|w - w_0\|_{\Sigma} + r\|w\|_{\star})^2. \quad (9)$$

The result is proved in the appendix. Since  $\alpha \approx 1.1124$ , the above approximation would allow us to get roughly 90%-optimality in the adversarial risk by minimizing the (much simpler) proxy function  $w \mapsto \tilde{E}^{\|\cdot\|}(w, w_0, r)$  instead. This will precisely be the focus of the next sections.

We also denote  $\tilde{E}_{opt}^{\|\cdot\|}(w_0, r)$  the smallest possible value of the adversarial risk proxy  $\tilde{E}^{\|\cdot\|}(\cdot, w_0, r)$ .

## 4 Gradient-Descent for Linear Regression

While the optimization of adversarial risk can be complex, the minimization of ordinary risk can be obtained by a simple gradient-descent. When applying a vanilla gradient-descent (GD) scheme with a step-size  $\eta > 0$ , starting at  $w_{\eta}(0) := 0_d$  to the ordinary risk defined in (5), one obtains at each iteration  $t \geq 1$  the following updates:

$$\begin{aligned} w_{\eta}(t, w_0) &:= w_{\eta}(t-1) - \eta \nabla_w E(w_{\eta}(t-1), w_0) \\ &= (I_d - (I_d - \eta \Sigma)^t) w_0 \end{aligned} \quad (10)$$

This scheme can be seen as a discrete approximation of the gradient flow induced by the following ODE:

$$\begin{cases} \dot{w}(t) = -\Sigma(w(t) - w_0) \\ w(0) = 0_d \end{cases}$$

which has a closed-form solution given by

$$w(t, w_0) := (I_d - \exp(-t\Sigma))w_0. \quad (11)$$

Our goal is to evaluate the robustness of the predictors obtained along the gradient descent path. We consider the continuous-time GD, as the analysis of discrete-time GD is analogous due to the absence of noise: the former is an infinitely small step-size  $\eta$  limit of the latter (Ali et al., 2019, 2020). In the following we mean by early-stopped GD, any predictor (indexed by training time,  $t$ ) obtained along the path of the gradient flow. Observe that for such predictors one has always that

$$E^{\|\cdot\|^2}(w(t, w_0), w_0, r) \geq E_{opt}^{\|\cdot\|^2}(w_0, r)$$

and so for any  $t \geq 0$ . In particular we will focus on the one that minimizes the adversarial risk at test time, that is  $\inf_{t \geq 0} E^{\|\cdot\|^2}(w(t, w_0), w_0, r)$ .

### 4.1 Euclidean Attacks: Almost-optimal Robustness

Here we consider Euclidean attacks, meaning that the attacker's norm is  $\|\cdot\| = \|\cdot\|_{\star} = \|\cdot\|_2$ . In the following proposition, we first characterize the non-robustness of the generative model  $w_0$ .

**Proposition 4.1.** If  $r \leq \sqrt{2/\pi} \frac{\|w_0\|_2}{\|w_0\|_{\Sigma^{-1}}}$ , then we have that

$$E^{\|\cdot\|^2}(w_0, w_0, r) = E_{opt}^{\|\cdot\|^2}(w_0, r).$$

and as soon as  $r > \sqrt{2/\pi} \frac{\|w_0\|_2}{\|w_0\|_{\Sigma^{-1}}}$ , we have

$$E^{\|\cdot\|^2}(w_0, w_0, r) / E_{opt}^{\|\cdot\|^2}(w_0, r) \geq \frac{r^2 \|w_0\|_2^2}{\|w_0\|_{\Sigma}^2}.$$

It is important to notice that the generative model  $w_0$  can be optimal w.r.t both the standard risk  $E(\cdot, w_0)$  and the adversarial risk  $E^{\|\cdot\|^2}(\cdot, w_0, r)$  for Euclidean attacks as soon as  $r$  is sufficiently small. However, as  $r$  increases, its adversarial risk becomes arbitrarily large. Therefore, applying a GD scheme until convergence may lead to predictors which are not robust to adversarial attacks even in the Euclidean setting. In the next proposition we investigate the robustness of the predictors obtained along the path of the GD scheme  $(w(t, w_0))_{t \geq 0}$  and we show that for any attack  $r \geq 0$ , this path contains an optimally robust predictor (up to an absolute constant  $\beta := 1.6862$ ).

**Proposition 4.2.** The following hold:

– If  $r \leq \sqrt{2/\pi} \frac{\|w_0\|_2}{\|w_0\|_{\Sigma^{-1}}}$  or  $r \geq \sqrt{\pi/2} \frac{\|w_0\|_{\Sigma^2}}{\|w_0\|_{\Sigma}}$ , then

$$\inf_{t \geq 0} E^{\|\cdot\|^2}(w(t, w_0), w_0, r) = E_{opt}^{\|\cdot\|^2}(w_0, r). \quad (12)$$

– If  $\sqrt{2/\pi} \frac{\|w_0\|_2}{\|w_0\|_{\Sigma^{-1}}} < r < \sqrt{\pi/2} \frac{\|w_0\|_{\Sigma^2}}{\|w_0\|_{\Sigma}}$ , we have that

$$\inf_{t \geq 0} E^{\|\cdot\|^2}(w(t, w_0), w_0, r) \leq \beta E_{opt}^{\|\cdot\|^2}(w_0, r) \quad (13)$$

Therefore the early-stopped vanilla GD scheme is able to capture an almost-optimally robust predictor for any Euclidean attack of radius  $r$  (see Figure 1 for an illustration).

In our proof of Proposition 4.2, we show that GD early-stopped at time  $t$  has the same adversarial risk (up to multiplicative constant) as a ridge estimator with regularization parameter  $\lambda \propto 1/t$ . The result then follows from (Xing et al., 2021), where it was shown the minimizer of the adversarial risk under Euclidean attacks is a ridge estimator.

In the isotropic case, i.e. when  $\Sigma = I_d$ , early-stopped GD even achieves the exact optimal adversarial risk.

**Proposition 4.3.** Assume that  $\Sigma = I_d$ , then for all  $r \geq 0$ , we have  $\inf_{t \geq 0} E^{\|\cdot\|^2}(w(t, w_0), w_0, r) = E_{opt}^{\|\cdot\|^2}(w_0, r)$ .




Figure 1: We consider a  $d = 2$  dimensional case where  $w_0$  and  $\Sigma$  are sampled according to a Gaussian and an Exponential distributions respectively. We plot the adversarial risk under Euclidean attacks of the optimal predictor along the GD path (**GD: adv. risk**) as well as its standard risk (**GD: standard risk**) and we compare them to the risks of the optimal predictor (**Opt**) solving adversarial problem when varying the attack strength  $r$ .

However, such results are only possible when the attacks are Euclidean. In the following section, we show that vanilla GD, along its whole optimization path, can be arbitrarily sub-optimal in terms of adversarial risk for Mahalanobis Attacks.

## 4.2 Mahalanobis Attacks: Sub-optimality of Gradient-Descent

Let us consider attacks w.r.t the Mahalanobis norm induced by a symmetric positive definite matrix  $B$ , i.e. we consider the case where

$$\|\cdot\| = \|\cdot\|_B := \|B^{1/2} \cdot\|_2.$$

In the next proposition, we present a simple case where GD fails to be adversarially robust under such attacks.

**Proposition 4.4.** *Let  $d = 2$ ,  $\Sigma = I_d$  and for any integer  $m \geq 1$ , let us consider the following positive-definite matrix*

$$B = B(m) = \begin{pmatrix} 1/m & 0 \\ 0 & m \end{pmatrix}. \quad (14)$$

Also, consider the following choice of generative model  $w_0 = w_0(m) = (1/\sqrt{m}, 1)$ . Then, for any fixed  $r > 0$ ,

$$\lim_{m \rightarrow +\infty} \frac{\inf_{t \geq 0} E^{\|\cdot\|_B}(w(t, w_0), w_0, r)}{E_{opt}^{\|\cdot\|_B}(w_0, r)} = +\infty.$$

Therefore under Mahalanobis attacks, any predictor obtained along the path of GD can be arbitrarily sub-optimal. To alleviate this issue, we propose in the next section a modified version of the vanilla GD scheme which can handle such attacks and even more general ones.

## 5 An Adapted Gradient-Descent: GD+

In fact, it is possible to obtain an almost optimally robust predictor for Mahalanobis attacks using a modified version of the GD scheme, termed GD+. Let  $M \in \mathcal{M}_d(\mathbb{R})$  be an arbitrary invertible matrix. In order to build an almost-optimally robust predictor for such attacks, we propose to apply a GD scheme on a transformed version of the data. More precisely, we propose to apply a GD scheme to the following objective function:

$$E_M(w, w_0) := \mathbb{E}_{(x,y) \sim P_{xy}} (w^\top Mx - y)^2$$

which leads to the following optimization dynamics:

$$w^M(t, w_0) := (I_d - e^{-tM\Sigma M^T})(M^{-1})^\top w_0. \quad (15)$$

This transformation amounts to apply feature-dependent gradient steps determined by  $M$  to the classical GD scheme. In the following proposition, we show that when  $M$  is adapted to the attack, early-stopped GD+ is optimally robust (up to an absolute constant).

**Proposition 5.1.** *For any  $B \in \mathcal{S}_d^{++}(\mathbb{R})$  and  $r \geq 0$ ,*

$$\inf_{t \geq 0} E^{\|\cdot\|_B}(B^{1/2}w^{B^{1/2}}(t, w_0), r) \leq \beta E_{opt}^{\|\cdot\|_B}(r, w_0).$$

Therefore by choosing  $M = B^{1/2}$ , GD+ is able to obtain near-optimality under  $\|\cdot\|_B$ -norm attack. See Figure 2 for an illustration. Note that when  $B = B^{1/2} = I_d$ , then  $w^{I_d}(t, w_0) = w(t, w_0)$ ,  $\|\cdot\|_{I_d} = \|\cdot\|_2$ , and we recover as a special case our result obtained in Proposition 4.2.




Figure 2: Here we plot the example studied in Proposition 4.4 for a fixed radius  $r = 1$  when varying  $m$ . **GD+** represents the modified GD scheme with  $M = B^{1/2}$  where  $B$  is defined in Eq. (14), **GD** represent the vanilla GD and **Opt** is the optimal predictor. Observe that the optimal adversarial risk goes to 0 as  $m$  goes to  $+\infty$  while the adversarial risk of the vanilla GD converges towards a constant.

In the following section, we investigate the robustness of GD+ scheme on more general attacks.

## 5.1 Robustness Under General Attacks

Our goal here is to provide a simple control of the adversarial risk of GD+ under general norm-attack  $\|\cdot\|$ . For that purpose, define the *condition number*  $\kappa^{\|\cdot\|}(M)$  of any matrix  $M$ , w.r.t to the attacker's norm  $\|\cdot\|$  as

$$\begin{aligned}\kappa^{\|\cdot\|}(M) &:= \lambda_{\max}^{\|\cdot\|}(M) / \lambda_{\min}^{\|\cdot\|}(M), \text{ where} \\ \lambda_{\max}^{\|\cdot\|}(M) &:= \sup_{w \neq 0} \frac{\|Mw\|_2}{\|w\|_*}, \text{ and} \\ \lambda_{\min}^{\|\cdot\|}(M) &:= \inf_{w \neq 0} \frac{\|Mw\|_2}{\|w\|_*}\end{aligned}$$

where dual norm  $\|\cdot\|_*$  is defined with respect to the attacker-norm  $\|\cdot\|$ . We are now ready to show a general upper-bound of our modified GD scheme.

**Proposition 5.2.** *For any  $r \geq 0$ , and invertible matrix  $M \in \mathcal{M}_d(\mathbb{R})$ , it holds that*

$$\frac{\inf_{t \geq 0} E^{\|\cdot\|}(M^\top w^M(t, w_0), w_0, r)}{E_{opt}^{\|\cdot\|}(r, w_0)} \leq \beta \kappa^{\|\cdot\|}((M^\top)^{-1})^2.$$

Therefore along the path of GD+ induced by  $M$ , one can find a  $\beta \kappa^{\|\cdot\|}((M^\top)^{-1})^2$ -optimally robust predictor against  $\|\cdot\|$ -attack. In particular, observe that when  $M = B^{1/2}$ , and  $\|\cdot\| = \|\cdot\|_B$ , we obtain that  $\kappa^{\|\cdot\|}((M^\top)^{-1}) = 1$  and we recover as a special case the result of Proposition 5.1.

**Remark 5.1.** *It is important to notice that  $M$  can be chosen arbitrarily, and therefore adapted to the norm of the attacks such that  $M \rightarrow \kappa^{\|\cdot\|}((M^\top)^{-1})$  is minimized. However, minimizing this quantity in general is hard due to the arbitrary choice of the norm  $\|\cdot\|$ .*

In the next section, we study a specific case of GD+ and provide a sufficient condition on the model  $w_0$  such that this scheme is near-optimal under general norm-attacks.

## 5.2 A Sufficient Condition for Optimality

We consider the general case of an arbitrary norm-attack  $\|\cdot\|$  with dual  $\|\cdot\|_*$  and we focus on a very specific path induced GD+, which is when  $M = \Sigma^{-1/2}$ . In that case, data are normalized and the path drawn by  $(Mw^M(t, w_0))_{t \geq 0}$  is in fact a uniform shrinkage of the generative model. More precisely, the predictors obtained along such a path are exactly the one in the chord  $[0, w_0] := \{\gamma w_0 \mid \gamma \in [0, 1]\}$ . In particular, the optimal adversarial risk achieved by this modified GD scheme is given by

$$\begin{aligned}\inf_{t \geq 0} E^{\|\cdot\|}(\Sigma^{-1/2} w^{\Sigma^{-1/2}}(t, w_0), w_0, r) \\ = \inf_{\gamma \in [0, 1]} E^{\|\cdot\|}(\gamma w_0, w_0, r)\end{aligned}\tag{16}$$

Let  $g(w_0) \in \mathbb{R}^d$  be a subgradient of  $\|\cdot\|_*$  at  $w_0$ . For example, in the case of  $\ell_\infty$ -norm-attacks, one may take

$g(w_0) = (\text{sign}(w_{0,1}), \dots, \text{sign}(w_{0,d}))$ , with  $\text{sign}(0) := 0$ . In the case of a Mahalanobis attack where  $\|\cdot\| = \|\cdot\|_B$  for some positive-definite matrix  $B$ , one can take  $g(w_0) = B^{1/2}w_0/\|w_0\|_B$  with  $g(0) = 0$ . We can now state our sufficient condition for near-optimality of GD+.

**Condition 5.1.** *The subgradient  $g(w_0) \in \mathbb{R}^d$  can be chosen such that*

$$\frac{\|g(w_0)\| \|w_0\|_*}{\|g(w_0)\|_{\Sigma^{-1}} \|w_0\|_{\Sigma}} \geq c,$$

where  $c$  is a positive absolute constant.

The above condition is sufficient in order to obtain near-optimality of GD+ as we show in the next proposition (see Figure 3 for an illustration).

**Theorem 5.1.** *Suppose Condition 5.1 is in order. Then, for any positive  $r$ , it holds for  $M = \Sigma^{-1}$  that*

$$\frac{\inf_{t \geq 0} E^{\|\cdot\|}(M^\top w^M(t, w_0), w_0, r)}{E_{opt}^{\|\cdot\|}(r, w_0)} \leq (1 \vee 1/c^2)\alpha.$$

In particular, for the case of  $\ell_\infty$ -norm-attacks, we have the following corollary.

**Corollary 5.1.** *Consider the case where  $\ell_\infty$ -norm-attacks. If there exists an absolute constant  $c > 0$  such that  $\|w_0\|_1 \geq c\sqrt{d}\|w_0\|_2$ , then with  $M = \Sigma^{-1}$  it holds that*

$$\frac{\inf_{t \geq 0} E^{\|\cdot\|_\infty}(M^\top w^M(t, w_0), w_0, r)}{E_{opt}^{\|\cdot\|_\infty}(r, w_0)} \leq \left(1 \vee \frac{\kappa^{\|\cdot\|_2}(\Sigma)}{c^2}\right)\alpha.$$

For example, when  $w_0 = (1, \dots, 1)$  –or equivalently, random  $w_0 \sim N(0, I_d)$ , one can take  $c = 1$ , and observe that  $\|w_0\|_1 \gtrsim d$ ,  $\|w_0\|_2 \asymp \sqrt{d}$ , and so the bound in Corollary 5.1 holds.

The Condition 5.1 that ensures optimality of GD+ in Theorem 5.1 cannot be removed. Indeed, we exhibit a simple case where the uniform shrinkage strategy fails miserably to find robust models even when they exist.

**Proposition 5.3.** *Let  $\Sigma = I_d$ , then it is possible to construct  $w_0 \in \mathbb{R}^d$  and  $r > 0$  such that in the limit  $d \rightarrow \infty$ , it holds that  $r \rightarrow 0$ ,  $r\sqrt{d} \rightarrow +\infty$ , and*

$$\frac{\inf_{t \geq 0} E^{\|\cdot\|_\infty}(w(t, w_0), w_0, r)}{E_{opt}^{\|\cdot\|_\infty}(r, w_0)} \rightarrow +\infty.$$

Therefore the uniform shrinkage strategy induced by GD+ (which here reduces to vanilla GD since  $\Sigma = I_d$ ) is not adapted for all scenarios; it may fail to find a robust model even when one exists.

In the next section, we restrict ourselves to the case of  $\ell_p$ -norm-attacks for  $p \in [1, +\infty]$  and show that GD+ is able to reach optimal robustness as soon as the data has uncorrelated features.




Figure 3: We consider the case exhibited in the Corollary 5.1 when  $d = 2$ ,  $\Sigma = I_d$  and  $w_0 = (1, 1)$  under  $\ell_\infty$ -attacks. We plot the adversarial as well as the standard risks for both the optimal shrinkage predictor as well as the optimal predictor of the adversarial risk when varying  $r$ .

### 5.3 Optimality Under $\ell_p$ -norm Attacks

Now, let  $p \in [1, +\infty]$  and consider attacks w.r.t the  $\ell_p$ -norm, i.e the attack strength is measured w.r.t the norm  $\|\cdot\| = \|\cdot\|_p$ , with dual norm  $\|\cdot\|_* = \|\cdot\|_q$ , where  $q \in [1, +\infty]$  is the harmonic conjugate of  $p$ . Popular examples in the literature are  $p = 2$  (corresponding to Euclidean attacks, considered in Section 4.1) and  $p = \infty$ . In this section we assume that  $\Sigma$  is a diagonal positive-definite matrix. This assumption translates the fact that, both norm  $\|\cdot\|_\Sigma$  and  $\|\cdot\|_q$  act on the same coordinates system. When these two norms are aligned, we show in the next theorem that the the minimiser of the proxy introduced in Eq. (9) is in fact a non-uniform shrinkage of  $w_0$  which can be recovered by GD+. An illustration of the result is provided in Figure 4.

**Theorem 5.2.** *Let  $\Sigma$  be any definite positive diagonal matrix and  $p \in [1, +\infty]$ , then we have*

$$\inf_{M \in \mathcal{M}_d(\mathbb{R}), t \geq 0} \frac{E^{\|\cdot\|_p}(M^\top w^M(t, w_0), w_0, r)}{E_{opt}^{\|\cdot\|_p}(r, w_0)} \leq \alpha.$$

Therefore GD+ is able to reach near-optimality in term of adversarial risk, and so for any  $\ell_p$ -attacks with  $p \in [1, +\infty]$  as soon as  $\|\cdot\|_p$  and  $\|\cdot\|_\Sigma$  are aligned.

Applying a GD+ scheme in practice might be difficult for general attacks, as the choice of the transformation  $M$  must be adapted accordingly. To alleviate this issue, we propose in the next section a simple and tractable two-stage estimator able to reach near-optimality and so for general attack.




Figure 4: We consider the case where  $d = 2$ ,  $\Sigma$  is diagonal and  $w_0$  and the diagonal coefficients of  $\Sigma$  are sampled according to a Gaussian and an exponential distribution respectively. The choice of  $M$  is obtained by fine-tuning the GD+. We compare the adversarial and standard risks under  $\ell_\infty$  of the optimal predictor in the GD+ path with those of the optimally robust predictor when varying the radius  $r$ .

## 6 Efficient Algorithms for Attacks in General Norms

We propose a simple tractable estimator whose adversarial risk is optimal up to within a multiplicative constant of 1.1124. Here, we drop the assumption of infinite training data  $n = \infty$ . Thus, the estimators are functions of the finite-training dataset  $D_n := \{(x_1, y_1), \dots, (x_n, y_n)\}$ , generated according to (3).

### 6.1 A Two-Stage Estimator and its Statistical Analysis

Consider any vector  $\hat{w}$  which minimizes the adversarial risk proxy  $w \mapsto \tilde{E}^{\|\cdot\|}(w, w_0, r)$  defined in (9). Note apart from its clear dependence on the generative model  $w_0$ ,  $\tilde{E}^{\|\cdot\|}$  also depends on the feature covariance matrix  $\Sigma$ . However, we don't assume that  $w_0$  nor  $\Sigma$  are known before hand; it has to be estimated from the finite training dataset  $\mathcal{D}_n$ . Thus, we propose a two-stage estimator described below in Algorithm 1.

---

**Algorithm 1** Proposed Two-Stage Estimator.

- 1: **Stage 1:** Compute consistent estimators  $\hat{w}_0$  and  $\hat{\Sigma}$  from  $w_0$  and  $\Sigma$  respectively, from the data  $\mathcal{D}_n$ .
  - 2: **Stage 2:** Compute  $\hat{w}$  which minimizes the adversarial risk proxy  $w \mapsto \tilde{E}^{\|\cdot\|}(w, w_0, r)$  defined in (9). See Algorithms 2 and 3 for implementations of this step.
- 

**Stage 1** of Algorithm 1 can be implemented using off-the-shelf estimators which adapt to the structural assumptions on  $\Sigma$  and  $w_0$  and  $\Sigma$  (sparsity, etc.). Later, we will pro-




Figure 5: Experiments with our two-stage estimator Algorithm 1. Here, we focus on  $\ell_\infty$ -norm-attacks. "1+2" means **Stage 2** of our Algorithm 1 is computed via the primal-dual algorithm (Algorithm 2), while "1+3" means it is computed via the Algorithm 3. The experimental setting here is: input-dimension  $d = 200$ ; covariance matrix  $\Sigma$  of the features = heterogeneous diagonal, with entries from  $\text{Exp}(1)$  distribution; the generative  $w_0$  is  $s$ -sparse vector (with  $s = 10$ ), normalized so that  $\|w_0\|_\Sigma^2 = 1$ . Notice how the adversarial risk of the estimator improves with number of samples  $n$ , as expected. See supplemental for details.

vide simple tractable algorithms for implementing **Stage 2**. Note that **Stage 2** implicitly requires the knowledge of the attacker-norm as it aims at minimizing  $\widehat{E}^{\|\cdot\|}$ .

## 6.2 Consistency of Proposed Two-Stage Estimator

We now establish the consistency of our proposed two-stage estimator  $\widehat{w}$  computed by Algorithm 1. Let  $\widehat{\Sigma}$  be an operator-norm consistent data-driven minimax estimator for  $\Sigma$ , and let  $\widehat{w}_0$  be a consistent estimator of the generative model  $w_0$ . Define error terms

$$e_1 := \|\widehat{w}_0 - w_0\|_2, \quad e_2 := \|\widehat{\Sigma} - \Sigma\|_{op}. \quad (17)$$

We are now ready to state the adversarial risk-consistency result for our proposed two-stage estimator (Algorithm 1).

**Theorem 6.1.** *For all  $r \geq 0$ , it holds that*

$$\begin{aligned} \widehat{E}_{opt}^{\|\cdot\|}(w_0, r) &\leq \widehat{E}^{\|\cdot\|}(\widehat{w}, w_0, r) \leq \widehat{E}_{opt}^{\|\cdot\|}(w_0, r) + \Delta, \\ E_{opt}^{\|\cdot\|}(w_0, r) &\leq E^{\|\cdot\|}(\widehat{w}, w_0, r) \leq \alpha E_{opt}^{\|\cdot\|}(w_0, r) + \Delta, \end{aligned}$$

where  $\alpha := 2/(1 + c_0) \approx 1.1124$ ,  $\Delta = O(e_1^2 + e_2^2)$ , where the hidden constant in the big- $O$  is of order  $\max(\|\Sigma\|_{op}^2, \|w_0\|_\Sigma^2)$ .

Thus, our proposed two-stage estimator is robust-optimal up to within a multiplicative factor  $2/(1 + c_0) \approx 1.1124$ , and an additive term  $O(e_1^2 + e_2^2)$  which is due to estimation error of  $w_0$  and  $\Sigma$  from training data. Note that if we assume that the covariance matrix  $\Sigma$  of the features is known, or equivalently that we have access to an unlimited supply of unlabelled data from  $N(0, \Sigma)$ , then we effectively have  $e_2 = 0$ . In this case, the statistical error of our proposed estimator  $\widehat{w}$  is dominated by the error  $e_1^2$  associated with estimating the generative model  $w_0$ . Under sparsity assumptions, this error term is further bounded by  $\frac{\sigma_\epsilon^2 s \log(ed/s)}{n}$

(thanks to the following well-known result), which tends to zero if the input dimension  $d$  does not grow much faster than the sample size  $n$ .

**Proposition 6.1** ((Bickel et al., 2009; Bellec et al., 2018)). *If  $1 \leq s \leq d/2$ , then under some mild technical conditions, it holds w.h.p that*

$$\inf_{\widehat{w}_0} \sup_{w_0 \in B_0^d(s)} \underbrace{\|\widehat{w}_0 - w_0\|_2}_{e_1} \asymp \sigma_\epsilon \sqrt{\frac{s \log(ed/s)}{n}}. \quad (18)$$

where  $B_0^d(s)$  is defined in Eq. (2). Moreover, the above minimax bound is attained by the square-root Lasso estimator with tuning parameter  $\lambda$  given by  $\lambda \asymp \sqrt{\frac{\log(2d/s)}{n}}$ .

**Remark 6.1.** *Note that, in the special case of Euclidean-norm attacks, our Theorem 6.1 (which works for all attack norms) recovers the adversarial risk-consistency result established in (Xing et al., 2021) as a special case.*

## 6.3 Algorithm 1: Primal-Dual Algorithm

We now devise a simple primal-dual algorithm for computing the second stage of our proposed estimator (Algorithm 1). The algorithm works for any covariance matrix  $\Sigma$  and norm-attack with tractable proximal operator.

Let  $\widehat{w}_0$  and  $\widehat{\Sigma}$  be the estimates computed in the **Stage 1** of Algorithm 1. Define and  $K := \widehat{\Sigma}^{1/2}$ ,  $a := K\widehat{w}_0$ ,  $f(z) := \|z - a\|_2$ ,  $g(w) := r\|w\|_*$ . Recall now that  $\sqrt{\widehat{E}^{\|\cdot\|}}(w, w_0, r) = \|w - w_0\|_\Sigma + r\|w\|_*$  and so for any model  $w$ . Then by "dualizing", we get that

$$\begin{aligned} \inf_{w \in \mathbb{R}^d} \|w - w_0\|_\Sigma + r\|w\|_* &= \inf_{w \in \mathbb{R}^d} f(Kw) + g(w) \\ &= \inf_{w \in \mathbb{R}^d} g(w) + \sup_{z \in \mathbb{R}^d} z^\top Kw - f^*(z) \\ &= \inf_{w \in \mathbb{R}^d} \sup_{z \in \mathbb{R}^d} \underbrace{z^\top Kw - f^*(z)}_{H(w, z)} + g(w), \end{aligned} \quad (19)$$

where  $f^*$  is the Fenchel-Legendre transform of  $f$ . Consider the following so-called Chambolle-Pock algorithm (Chambolle and Pock, 2010) for computing a saddle-point for the function  $H$ .

**Algorithm 2** Primal-Dual algorithm which implements **Stage 2** of Algorithm 1. Only one iteration is shown here.

**Inputs:**  $\widehat{w}_0, \widehat{\Sigma}, \eta_1, \eta_2, z^{(0)} = w^{(0)} = \mathbf{0}_d$ .

- 1:  $z^{(t+1)} \leftarrow \text{proj}_{B_{\|\cdot\|_2}^d} (z^{(t)} + \eta_2 \widehat{\Sigma}^{1/2} (u^{(t)} - \widehat{w}_0))$
- 2:  $w^{(t+1)} \leftarrow \text{prox}_{\eta_1 r \|\cdot\|_*} (w^{(t)} - \eta_1 \widehat{\Sigma}^{1/2} z^{(t+1)})$ ,
- 3:  $u^{(t+1)} \leftarrow 2w^{(t+1)} - w^{(t)}$

Here, the  $\eta_k$ 's are stepsizes chosen such that  $\eta_1 \eta_2 \|\widehat{\Sigma}\|_{op}^{1/2} < 1$ . The nice thing here is that the projection onto the  $\ell_2$ -ball






Figure 6: We compare our proposed two-stage estimator given in Alg. 1 with the one proposed in Xing et al. (2021), in the specific setting of Euclidean attacks when varying sample size  $n$  and attack strength  $r$ . Here, **Stage 2** of our two-stage estimator is computed via Alg. 2. Dashed and plain lines respectively represent the standard and adversarial risks. We see that though more general, our algorithm is able to recover similar performances as Xing et al. (2021) in the special case of Euclidean attacks.

(line 1) admits a simple analytic formula. In the case of  $\ell_\infty$ -norm-attacks, the second line corresponds to the well-known soft-thresholding operator; etc. Refer to Figures 5 & 6 for empirical illustrations of the algorithm. The following convergence result follows directly from (Chambolle and Pock, 2010).

**Proposition 6.2.** *Algorithm 2 converges to a stationary point  $(w^{(\infty)}, z^{(\infty)})$  of the  $H$  at an ergodic rate  $O(1/t)$ .*

#### 6.4 Algorithm 2: Simple Thresholding-Based Algorithm in the Case of $\ell_\infty$ -Norm Attacks

Now, suppose the covariance matrix of the features  $\Sigma$  is diagonal, i.e  $\Sigma = \text{diag}(\lambda_1, \dots, \lambda_d)$ . In the case of  $\ell_\infty$ -attacks, we provide a much simpler and faster algorithm for computing the second stage of our proposed estimator. Indeed, for  $\ell_\infty$ -attack we obtain an explicit form of the optimal solution minimizing Eq. (9) as shown in the next Proposition. We deduce the following result.

**Proposition 6.3.** *Let  $c := \max_{1 \leq j \leq d} |(w_0)_j| \lambda_j$ . There exists  $t \in [0, c]$  such that  $w(t) \in \mathbb{R}^d$  is a minimizer of the convex function  $w \mapsto \|w - w_0\|_\Sigma + r\|w\|_1$ , where*

$$w(t)_j = \text{ST}((w_0)_j; rt/\lambda_j), \text{ for all } j \in [d]. \quad (20)$$

and  $\text{ST}(\cdot; s)$  is the soft-thresholding operator at level  $s$ .

An inspection of (20) reveals a kind of feature-selection. Indeed, if the  $j$ th component of the ground-truth model  $w_0$  is small in the sense that  $|(w_0)_j| \leq rt/\lambda_j$ , then  $w(t)_j = 0$ , i.e the  $j$ th component of  $w_0$  should be suppressed. On the other hand, if  $|(w_0)_j| \geq rt/\lambda_j$ , then  $(w_0)_j$  should be replaced by the translated version  $(w_0)_j - \text{sign}((w_0)_j)rt/\lambda_j$ . That is weak components of  $w_0$  are suppressed, while strong components are boosted.

The result is Algorithm 3, a simple method for computing **Stage 2** of our proposed two-stage estimator (Algorithm 1). Refer to Figure 5 for an empirical illustration.

---

**Algorithm 3** Non-uniform soft-thresholding which implements **Stage 2** of Algorithm 1 for diagonal covariance matrix under  $\ell_\infty$ -attacks.

---

**Inputs:**  $\hat{w}_0, \hat{\Sigma}$

- 1: **Compute**  $\hat{c} = \max_{1 \leq j \leq d} |(\hat{w}_0)_j| \hat{\lambda}_j$ .
  - 2: **For each**  $t$  in a finite grid of values between 0 and  $\hat{c}$ , use held-out to retain the value of  $t$  for which the adversarial risk of  $\hat{w}(t)$  is minimal, where each component of  $\hat{w}(t)$  is given by  $\hat{w}(t)_j = \text{ST}((\hat{w}_0)_j; rt/\hat{\lambda}_j)$ .
  - 3: **Return**  $\hat{w}(t)$ .
- 

## 7 Concluding Remarks

In our work, we have undertaken a study of the robustness of gradient-descent (GD), under test-time attacks in the context of linear regression. Our work provides a clear characterization of when GD and a modified version (GD+) –with feature-dependent learning rates– can succeed to achieve the optimal adversarial risk (up to an absolute constant). This characterization highlights the effect over the covariance structure of the features, and also, of the norm used to measure the strength of the attack.

Finally, our paper proposes a statistically consistent and simple two-stage estimator which achieves optimal adversarial risk in the population regime, up to within a constant factor. Our proposed estimator adapts to attacks w.r.t general norms, and to the covariance structure of the features.

## References

- Ali, A., Dobriban, E., and Tibshirani, R. (2020). The implicit regularization of stochastic gradient flow for least squares. In III, H. D. and Singh, A., editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 233–244. PMLR.
- Ali, A., Kolter, J. Z., and Tibshirani, R. J. (2019). A continuous-time view of early stopping for least squares regression. In *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, volume 89 of *Proceedings of Machine Learning Research*, pages 1370–1378. PMLR.
- Bellec, P. C., Lecué, G., and Tsybakov, A. B. (2018). Slope meets Lasso: Improved oracle bounds and optimality. *The Annals of Statistics*, 46(6B):3603 – 3642.
- Bhagoji, A. N., Cullina, D., and Mittal, P. (2019). Lower bounds on adversarial robustness from optimal transport.
- Bhattacharjee, R., Jha, S., and Chaudhuri, K. (2021). Sample complexity of robust linear classification on separated data. In *Proceedings of the 38th International Conference on Machine Learning (ICML)*. PMLR.
- Bickel, P. J., Ritov, Y., and Tsybakov, A. B. (2009). Simultaneous analysis of Lasso and Dantzig selector. *The Annals of Statistics*, 37(4):1705 – 1732.
- Bubeck, S., Price, E., and Razenshteyn, I. P. (2018). Adversarial examples from computational constraints. *CoRR*, abs/1805.10204.
- Chambolle, A. and Pock, T. (2010). A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of Mathematical Imaging and Vision*, 40.
- Dobriban, E., Hassani, H., Hong, D., and Robey, A. (2020). Provable tradeoffs in adversarially robust classification. *arXiv preprint arXiv:2006.05161*.
- Dohmatob, E. (2019). Generalized no free lunch theorem for adversarial robustness. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, volume 97 of *Proceedings of Machine Learning Research*. PMLR.
- Dohmatob, E. (2021). Fundamental tradeoffs between memorization and robustness in random features and neural tangent regimes. *arXiv preprint arXiv:2106.02630*.
- Gilmer, J., Metz, L., Faghri, F., Schoenholz, S. S., Raghu, M., Wattenberg, M., and Goodfellow, I. J. (2018). Adversarial spheres. *CoRR*, abs/1801.02774.
- Hassani, H. and Javanmard, A. (2022). The curse of overparametrization in adversarial training: Precise analysis of robust generalization for random features regression. *arXiv preprint arXiv:2201.05149*.
- Javanmard, A. and Mehrabi, M. (2021). Adversarial robustness for latent models: Revisiting the robust-standard accuracies tradeoff. *arXiv preprint arXiv:2110.11950*.
- Javanmard, A. and Soltanolkotabi, M. (2022). Precise statistical analysis of classification accuracies for adversarial training. *The Annals of Statistics*, 50(4):2127–2156.
- Javanmard, A., Soltanolkotabi, M., and Hassani, H. (2020). Precise tradeoffs in adversarial training for linear regression. In *COLT*.
- Khim, J. and Loh, P.-L. (2018). Adversarial risk bounds via function transformation. *arXiv preprint arXiv:1810.09519*.
- Li, Y., Fang, E. X., Xu, H., and Zhao, T. (2020). Implicit bias of gradient descent based adversarial training on separable data. In *ICLR 2020*.
- Madry, A., Makelov, A., Schmidt, L., Tsipras, D., and Vladu, A. (2018). Towards deep learning models resistant to adversarial attacks.
- Mahloujifar, S., Diochnos, D. I., and Mahmood, M. (2018). The curse of concentration in robust learning: Evasion and poisoning attacks from concentration of measure. *CoRR*, abs/1809.03063.
- Min, Y., Chen, L., and Karbasi, A. (2021a). The curious case of adversarially robust models: More data can help, double descend, or hurt generalization. In *Uncertainty in Artificial Intelligence (UAI)*.
- Min, Y., Chen, L., and Karbasi, A. (2021b). The curious case of adversarially robust models: More data can help, double descend, or hurt generalization. In *Proceedings of the Thirty-Seventh Conference on Uncertainty in Artificial Intelligence*. PMLR.
- Rockafellar, R. T. (1970). On the maximal monotonicity of subdifferential mappings. *PACIFIC JOURNAL OF MATHEMATICS*, 33(1).
- Schmidt, L., Santurkar, S., Tsipras, D., Talwar, K., and Madry, A. (2018). Adversarially robust generalization requires more data. *CoRR*, abs/1804.11285.
- Shafahi, A., Huang, W. R., Studer, C., Feizi, S., and Goldstein, T. (2018). Are adversarial examples inevitable? *CoRR*, abs/1809.02104.
- Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., and Fergus, R. (2013). Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*.
- Tsipras, D., Santurkar, S., Engstrom, L., Turner, A., and Madry, A. (2019). Robustness may be at odds with accuracy. In *International Conference on Learning Representations (ICLR)*, volume abs/1805.12152.
- Vardi, G., Yehudai, G., and Shamir, O. (2022). Gradient methods provably converge to non-robust networks. *NeurIPS*.

- Wei, Y., Fang, B., and Wainwright, M. J. (2020). From Gauss to Kolmogorov: Localized measures of complexity for ellipses. *Electronic Journal of Statistics*, 14(2):2988 – 3031.
- Xing, Y., Zhang, R., and Cheng, G. (2021). Adversarially robust estimate and risk analysis in linear regression. In *AISTATS*, volume 130. PMLR.
- Yang, Y.-Y., Rashtchian, C., Zhang, H., Salakhutdinov, R. R., and Chaudhuri, K. (2020). A closer look at accuracy vs. robustness. In *Advances in Neural Information Processing Systems*, volume 33, pages 8588–8601. Curran Associates, Inc.
- Yin, D., Kannan, R., and Bartlett, P. (2019). Rademacher complexity for adversarially robust generalization. In *International conference on machine learning*.

# Supplementary Material

## Robust Linear Regression: Gradient-descent, Early-stopping, and Beyond

---

### Contents

<b>A Further Details of Experiments</b>	<b>12</b>
A.1 Case of Optimality of GD under Euclidean Attacks . . . . .	13
A.2 Suboptimality of GD under General Mahalanobis Attacks . . . . .	13
A.3 Sufficient Condition for Optimality . . . . .	13
A.4 Case of $\ell_p$ -Norm Attacks . . . . .	13
A.5 Experiments for Algorithm 1, the Two-Stage Estimator Proposed in Section 6 . . . . .	13
<b>B Preliminaries</b>	<b>14</b>
B.1 Additional Notations . . . . .	14
B.2 Proof of Lemma 2.1: Analytic Formula for Adversarial Risk . . . . .	14
B.3 Proof of Lemma 3.1: Proxy for Adversarial Risk . . . . .	15
<b>C Proof of the Results of Section 4: Gradient-descent for Linear Regression</b>	<b>15</b>
C.1 Proof of Proposition 4.1: Sub-optimality of the Generative Model w.r.t Adversarial Risk . . . . .	15
C.2 Proof of Proposition 4.2: Optimality of Gradient-Descent (GD) in Case of Euclidean Attacks . . . . .	16
C.3 Uniform Shrinkage of $w_0$ is Sub-optimal in General . . . . .	17
C.4 Proof of Proposition 4.4: Sub-optimality of GD in Case of Mahalanobis Attacks . . . . .	18
<b>D Proofs of Results of Section 5: An Adapted Gradient-Descent (GD+)</b>	<b>19</b>
D.1 Proof of Proposition 5.1: Optimality of GD+ in Case of Mahalanobis Norm Attacks . . . . .	19
D.2 Proof of Proposition 5.2: Case of Optimality of GD+ in Case of General Norm Attacks . . . . .	20
D.3 Proof of Theorem 5.1: A Case of Optimality of GD for Non-Euclidean / General Norm Attacks . . . . .	20
D.4 Proof of Proposition 5.3: A Case of Sub-optimality of GD to $\ell_\infty$ -Norm Attacks . . . . .	21
D.5 Proof of Theorem 5.2: Case of Optimality of GD+ to $\ell_p$ -Norm Attacks . . . . .	22
<b>E Proof of Theorem 6.1: Consistency of Proposed Two-Stage Estimator (Algorithm 1)</b>	<b>25</b>

### A Further Details of Experiments

In this section we detail the experiments presented in the main paper, to empirically verify our theoretical results. In the Figures 1, 2, 3 and 4, we consider only  $d = 2$ -dimensional cases as we do not have access to the optimal value of the adversarial risk and propose to compute it using a grid-search.

### A.1 Case of Optimality of GD under Euclidean Attacks

In Figure 1, we plot the optimal value of the adversarial risk under Euclidean attacks obtained along the path of the GD scheme and compare it with the true optimal adversarial risk when varying the attack strength  $r$ . We observe that we recover our statement proved in Proposition 4.2, that is that GD is optimally robust (up to an absolute constant) and in fact we observe that we have exact equality between these two risks. We also compare their standard risks and observe that they are also equal. It is important to notice that for  $r$  sufficiently small, the optimally robust predictor has a null standard risk and therefore is also optimal, which is the result proved in Proposition 4.1.

### A.2 Suboptimality of GD under General Mahalanobis Attacks

In Figure 2, we consider the example proposed in Proposition 4.4 and we plot the adversarial risks of the vanilla GD presented in Eq. (11), as well as the one obtained by our modified GD+ presented in Eq. (15) and compare them to the optimal adversarial risk when varying  $m$ . We show that vanilla GD can be arbitrarily sub-optimal as  $m$  goes to infinity while our modified GD+, when selecting the adapted transformation  $B(m)$ , is able to reach the optimal adversarial risk.

### A.3 Sufficient Condition for Optimality

In Figure 3, we consider a simple example where the Condition 5.1 is satisfied and illustrate that in that case the uniform shrinkage strategy, obtained by applying a GD scheme on the normalized data (i.e. by considering GD+ with  $M = \Sigma^{-1/2}$ ), is optimally robust (as shown in Theorem 5.1). For that purpose we consider a dense generative model  $w_0$  under  $\ell_\infty$ -attacks and a covariance matrix  $\Sigma$  with a small condition number (in fact we consider  $\Sigma = I_d$  for simplicity) and we show that the adversarial risk of the uniform shrinkage strategy matches the optimal one.

### A.4 Case of $\ell_p$ -Norm Attacks

In Figure 4, we illustrate the result obtained in Theorem 5.2. For that purpose we consider a general case where the covariance matrix is diagonal and sampled according to an Exponential distribution and the generative model is also chosen at random according to a Gaussian distribution. We consider the case of  $\ell_\infty$ -attacks as it is the most used setting in practice. In order to select  $M$  in GD+, we propose to solve the following system in  $D$  (with  $D$  diagonal and positive definite)

$$(I_d - \exp(-tD))w_0 = w_{opt},$$

where  $w_{opt}$  is obtained by solving the adversarial problem using a grid-search and  $w_0$  is the generative model. Then we consider our GD+ scheme using  $M = (D\Sigma^{-1})^{1/2}$ . Indeed, in the proof of Theorem 5.2, we propose a constructive approach in order to show the optimality of GD+ which simply requires to find a diagonal matrix  $D$  solving the above system. We compare the adversarial risk of our GD+ with this specific choice of  $M$  against the true optimal adversarial risk and we observe that the two curves coincide when varying the radius  $r$ .

### A.5 Experiments for Algorithm 1, the Two-Stage Estimator Proposed in Section 6

In Figure 5, we plot the adversarial risk computed by Algorithm 1, in the case of  $\ell_\infty$ -norm attacks. the experimental setting here is: input-dimension  $d = 200$ ; covariance matrix  $\Sigma$  of the features is heterogeneous diagonal, with diagonal entries drawn iid from an exponential distribution with mean  $\lambda = 1$ ; the generative  $w_0$  is  $s$ -sparse vector with  $s = 10$ , normalized so that  $\|w_0\|_\Sigma^2 = 1$ . As explained in the figure legends, the acronym "1+3" means it is computed via the Algorithm 3 (which works for diagonal covariance matrix and  $\ell_\infty$ -norm attack), while "1+2" means **Stage 2** of our Algorithm 1 is computed via the primal-dual algorithm (Algorithm 2, which works for every feature covariance matrix  $\Sigma$ ).

Notice how the adversarial risk of the estimator improves with number of samples  $n$ , as expected (Theorem 6.1). Also, unlike "1+3", for small values of the attack strength  $r$  the adversarial risk achieved by "1+2" is dominated by optimization error incurred in **Stage 2** (Algorithm 2). Thanks to Proposition 6.2, this can be alleviated by running more iterations of Algorithm 2.

In Figure 6, we also compare our proposed algorithm with the Two-Stage Estimator Proposed in Xing et al. (2021). Here, the attack norm is Euclidean, and the covariance matrix of the features  $\Sigma$  is as in the above paragraph.

## B Preliminaries

### B.1 Additional Notations

The asymptotic notation  $F(d) = O(G(d))$  (also written  $F(d) \lesssim G(d)$ ) means there exists a constant  $c$  such that  $F(d) \leq c \cdot G(d)$  for sufficiently large  $d$ , while  $F(d) = \Omega(G(d))$  means  $G(d) = O(F(d))$ , and  $F(d) = \Theta(G(d))$  or  $F(d) \asymp G(d)$  means  $F(d) \lesssim G(d) \lesssim F(d)$ . Finally,  $F(d) = o(G(d))$  means  $F(d)/G(d) \rightarrow 0$  as  $d \rightarrow \infty$ . Thus,  $F(d) = O(1)$  means that  $F(d)$  is bounded for sufficiently large  $d$ ;  $F(d) = \Omega(1)$  means that  $F(d)$  is bounded away from zero for sufficiently large  $d$ ;  $F(d) = o(1)$  means that  $F(d) \rightarrow 0$ . The acronym w.h.p is used to indicate that a statement is true except on an event of probability  $o(1)$ .

Also recall all the notations introduced in the beginning of Section 2.

### B.2 Proof of Lemma 2.1: Analytic Formula for Adversarial Risk

Recall the adversarial risk functional  $E$  defined in (7).

**Lemma 2.1.** *Recall that  $c_0 = \sqrt{2/\pi}$ , then for any  $w \in \mathbb{R}^d$  and  $r \geq 0$ , it holds that*

$$E^{\|\cdot\|}(w, w_0, r) = \|w - w_0\|_{\Sigma}^2 + r^2 \|w\|_{\star}^2 + 2c_0 r \|w - w_0\|_{\Sigma} \|w\|_{\star}. \quad (7)$$

For the proof, we will need the following auxiliary lemma.

**Lemma B.1.** *For any  $x, w \in \mathbb{R}^d$ ,  $r \geq 0$  and  $y \in \mathbb{R}$ , the following identity holds*

$$\sup_{\|\delta\| \leq r} ((x + \delta)^{\top} w - y)^2 = (|x^{\top} w - y| + r \|w\|_{\star})^2. \quad (21)$$

*Proof.* Note that  $h(x, \delta)/2 = \eta(x)^2/2 + g(x, \delta)/2$ , where  $g(x, \delta) := w(\delta)^2 - 2\eta(x)w(\delta)$ , and  $\eta(x) := w(x) - y$ , and  $w(x) := x^{\top} w$ . Now, because the function  $z \rightarrow z^2/2$  is its own Fenchel-Legendre conjugate, we can "dualize" our problem as follows

$$\begin{aligned} \sup_{\|\delta\| \leq r} g(x, \delta)/2 &= \sup_{\|\delta\|_{\star} \leq r} -\eta(x)w(\delta) + \sup_{z \in \mathbb{R}} zw(\delta) - z^2/2 \\ &= \sup_{z \in \mathbb{R}} -z^2/2 + \sup_{\|\delta\| \leq r} (z - \eta(x))w(\delta) \\ &= \sup_{z \in \mathbb{R}} r \|w\|_{\star} |z - \eta(x)| - z^2/2 \\ &= \sup_{s \in \{\pm 1\}} \sup_{z \in \mathbb{R}} r s (z - \eta(x)) - z^2/2 \\ &= \sup_{s \in \{\pm 1\}} -r \|w\|_{\star} s \eta(x) + \sup_{z \in \mathbb{R}} r \|w\|_{\star} s z - z^2/2 \\ &= \sup_{s \in \{\pm 1\}} -r \|w\|_{\star} s \eta(x) + r^2 \|w\|_{\star}^2 / 2 \\ &= r \|w\|_{\star} |\eta(x)| + r^2 \|w\|_{\star}^2 / 2. \end{aligned}$$

We deduce that

$$\sup_{\|\delta\| \leq r} h(x, \delta)/2 = \eta(x)^2/2 + r \|w\|_{\star} |\eta(x)| + r^2 \|w\|_{\star}^2 / 2 = (|\eta(x)| + r \|w\|_{\star})^2 / 2,$$

as claimed.  $\square$

*Proof of Lemma 2.1.* Indeed, thanks to Lemma B.1 with  $y = x^{\top} w_0$ , one has

$$E^{\|\cdot\|}(w, w_0, r) := \mathbb{E}_{x \sim N(0, \Sigma)} \sup_{\|\delta\| \leq r} h(x, \delta) = \mathbb{E}_{x \sim N(0, \Sigma)} [(\eta(x) + r \|w\|_{\star})^2],$$

where the functions  $h$  and  $\eta$  are as in the proof of Lemma B.1. The result then follows upon noting that

$$\begin{aligned}\mathbb{E}_{x \sim N(0, \Sigma)}[\eta(x)^2] &= \mathbb{E}_{x \sim N(0, \Sigma)}[(x^\top w - x^\top w_0)^2] = \|w - w_0\|_\Sigma^2, \\ \mathbb{E}_{x \sim N(0, \Sigma)}|\eta(x)| &= \mathbb{E}_{x \sim N(0, \Sigma)}|x^\top w - x^\top w_0| = c_0 \|w - w_0\|_\Sigma,\end{aligned}$$

where  $c_0 := \sqrt{2/\pi}$ . □

### B.3 Proof of Lemma 3.1: Proxy for Adversarial Risk

Recall the definition of the adversarial risk proxy  $\tilde{E}$  from (9).

**Lemma 3.1.** *For any  $w \in \mathbb{R}^d$  and  $r \geq 0$ , it holds that*

$$E^{\|\cdot\|}(w, w_0, r) \leq \tilde{E}^{\|\cdot\|}(w, w_0, r) \leq \alpha \cdot E^{\|\cdot\|}(w, w_0, r), \quad (8)$$

where  $\alpha := 2/(1 + \sqrt{2/\pi}) \approx 1.1124$  and

$$\tilde{E}^{\|\cdot\|}(w, w_0, r) := (\|w - w_0\|_\Sigma + r\|w\|_\star)^2. \quad (9)$$

*Proof.* Let us first show the following useful Lemma.

**Lemma B.2.** *For any  $a, b, c \geq 0$  with  $c \leq 1$ , it holds that*

$$(a + b)^2 \geq a^2 + b^2 + 2abc \geq \frac{1+c}{2}(a + b)^2, \quad (22)$$

with equality if  $c = 1$ .

*Proof.* Let  $h(a, b, c) := a^2 + b^2 + 2abc$ . For the LHS, it suffices to observe that  $h(a, b, c) \leq h(a, b, 1) = (a + b)^2$ . For the RHS, WLOG assume that  $a \neq 0$ , and set  $t := b/a \geq 0$ . Observe  $\frac{h(a, b, c)}{(a + b)^2} = \frac{1 + t^2 + 2ct}{(1 + t)^2}$ , which is minimized when  $t = 1$ , because  $0 \leq c \leq 1$  by assumption. We deduce that  $h(a, b, c)/(a + b)^2 \geq (1 + 1 + 2c)/(1 + 1)^2 = (1 + c)/2$ , as claimed. □

Now, recall from Lemma 2.1 that  $E^{\|\cdot\|}(w, w_0, r) = \|w - w_0\|_\Sigma^2 + r^2\|w\|_\star^2 + 2c_0r\|w - w_0\|_\Sigma\|w\|_\star$  with  $c_0 := \sqrt{\pi/2}$ , and by denoting  $a = \|w - w_0\|_\Sigma$ ,  $b = r\|w\|_\star$  and  $c = c_0$  the result follows directly from Lemma B.2. □

## C Proof of the Results of Section 4: Gradient-descent for Linear Regression

### C.1 Proof of Proposition 4.1: Sub-optimality of the Generative Model w.r.t Adversarial Risk

**Proposition 4.1.** *If  $r \leq \sqrt{2/\pi} \frac{\|w_0\|_2}{\|w_0\|_{\Sigma^{-1}}}$ , then we have that*

$$E^{\|\cdot\|^2}(w_0, w_0, r) = E_{opt}^{\|\cdot\|^2}(w_0, r).$$

and as soon as  $r > \sqrt{2/\pi} \frac{\|w_0\|_2}{\|w_0\|_{\Sigma^{-1}}}$ , we have

$$E^{\|\cdot\|^2}(w_0, w_0, r)/E_{opt}^{\|\cdot\|^2}(w_0, r) \geq \frac{r^2\|w_0\|_2^2}{\|w_0\|_\Sigma^2}.$$

*Proof.* In (Xing et al., 2021, Proposition 1), the authors already show the optimality of  $w_0$  for  $r \leq \sqrt{2/\pi} \frac{\|w_0\|_2}{\|w_0\|_{\Sigma^{-1}}}$ . Now remark that for  $w = 0$ , one has

$$E_{opt}^{\|\cdot\|^2}(w_0, r) \leq E^{\|\cdot\|^2}(0, w_0, r) \leq \|w_0\|_\Sigma^2 \quad (23)$$

Moreover, we have that for all  $r \geq 0$ ,  $E^{\|\cdot\|^2}(w_0, w_0, r) = r^2\|w_0\|_2^2$ . Combining it with Eq. (23) gives desired result. □

## C.2 Proof of Proposition 4.2: Optimality of Gradient-Descent (GD) in Case of Euclidean Attacks

Recall the GD dynamics  $t \mapsto w(t, w_0)$  given in (11). We restate the result here for convenience.

**Proposition 4.2.** *The following hold:*

– If  $r \leq \sqrt{2/\pi} \frac{\|w_0\|_2}{\|w_0\|_{\Sigma^{-1}}}$  or  $r \geq \sqrt{\pi/2} \frac{\|w_0\|_{\Sigma^2}}{\|w_0\|_{\Sigma}}$ , then

$$\inf_{t \geq 0} E^{\|\cdot\|_2}(w(t, w_0), w_0, r) = E_{opt}^{\|\cdot\|_2}(w_0, r). \quad (12)$$

– If  $\sqrt{2/\pi} \frac{\|w_0\|_2}{\|w_0\|_{\Sigma^{-1}}} < r < \sqrt{\pi/2} \frac{\|w_0\|_{\Sigma^2}}{\|w_0\|_{\Sigma}}$ , we have that

$$\inf_{t \geq 0} E^{\|\cdot\|_2}(w(t, w_0), w_0, r) \leq \beta E_{opt}^{\|\cdot\|_2}(w_0, r) \quad (13)$$

For the proof, we will need the following lemma.

**Lemma C.1.** *For any  $\beta \geq 0$ , let us denote  $w_\beta(w_0) = (\beta I_d + \Sigma)^{-1} \Sigma w_0$  the ridge solution. Then for all  $r \geq 0$  and  $t \geq 0$ , we have*

$$E^{\|\cdot\|_2}(w(t, w_0), w_0, r) \leq E^{\|\cdot\|_2}(w_{1/t}(w_0), w_0, 1.2985r).$$

*Proof.* In the Euclidean case we obtain a simple expression of  $E^{\|\cdot\|_2}(\cdot, w_0, r)$ . Let us denote  $\{v_1, \dots, v_d\}$  an orthonormal basis where  $\Sigma$  can be diagonalized,  $\lambda_1 \geq \dots \geq \lambda_d$  the eigenvalues of  $\Sigma$  and set  $c_j := w_0^\top v_j$  for all  $j \in [d]$ . The adversarial risk associated to  $w(t)$  can be written as

$$E^{\|\cdot\|_2}(w(t, w_0), w_0, r) = E_1(t, w_0) + r^2 E_2(t, w_0) + 2r \sqrt{2/\pi} \sqrt{E_1(t, w_0) E_2(t, w_0)}, \text{ where}$$

$$E_1(t, w_0) := \sum_{j=1}^d (1 - \alpha \lambda_j)^{2t} \lambda_j c_j^2, \text{ and } E_2(t, w_0) := \sum_{j=1}^d (1 - (1 - \alpha \lambda_j)^t)^2 c_j^2.$$

By using the elementary fact that  $1 - x \leq \exp(-x) \leq 1/(1+x)$  for all  $x \geq 0$ , we obtain first that

$$E_1(t, w_0) = \sum_{j=1}^d \exp(-\lambda_j 2t) \lambda_j c_j^2 \leq \sum_{j=1}^d \frac{\lambda_j}{(1 + t \lambda_j)^2} c_j^2 = \|w_{1/t}(w_0) - w_0\|_{\Sigma}^2 = E_1(w_{1/t}(w_0)).$$

In addition, using the elementary fact<sup>1</sup> that  $1 - \exp(-x) \leq 1.2985x/(1+x)$  for all  $x \geq 0$ , it follows that

$$E_2(t, w_0) = \sum_{j=1}^d (1 - \exp(-\lambda_j t))^2 c_j^2 \leq 1.6862 \sum_{j=1}^d \left( \frac{\lambda_j t}{1 + \lambda_j t} \right)^2 c_j^2 = 1.2985^2 E_2(w_{1/t}(w_0))$$

and the result follows.  $\square$

Finally, we will need the following elementary lemma.

**Lemma C.2.** *for any  $\gamma \geq 1$ ,  $r \geq 0$ , we have  $E^{\|\cdot\|}(w, w_0, \gamma r) \leq \gamma^2 E^{\|\cdot\|}(w, w_0, r)$*

*Proof.* Indeed we have

$$\|w - w_0\|_{\Sigma}^2 \leq \gamma^2 \|w - w_0\|_{\Sigma}^2 \text{ and } 2c_0 \gamma r \|w - w_0\|_{\Sigma} \|w\|_{\star} \leq 2c_0 \gamma^2 r \|w - w_0\|_{\Sigma} \|w\|_{\star}$$

from which follows the result.  $\square$

We are now in place to prove Proposition 4.2.

<sup>1</sup>For example, see (Ali et al., 2019, Lemma 7).



*Proof of Proposition 4.2.* First note that in (Xing et al., 2021, Proposition 1), the authors show that as soon as  $r \leq \sqrt{2/\pi} \frac{\|w_0\|_2}{\|w_0\|_{\Sigma^{-1}}}$ , the robust predictor minimizing Eq. (4) is  $w_0$  which is attained by the GD dynamic when  $t$  goes to infinity. They also show that when  $r \geq \sqrt{\pi/2} \frac{\|w_0\|_{\Sigma^2}}{\|w_0\|_{\Sigma}}$ , the optimal solution is 0 which is reached by the vanilla GD scheme for  $t = 0$ . Let us now show the result for  $\sqrt{2/\pi} \frac{\|w_0\|_2}{\|w_0\|_{\Sigma^{-1}}} < r < \sqrt{\pi/2} \frac{\|w_0\|_{\Sigma^2}}{\|w_0\|_{\Sigma}}$ . Indeed, in (Xing et al., 2021, Proposition 1), they show that for any  $r$  in this range, there exists  $\beta^* > 0$ , such that

$$E^{\|\cdot\|_2}(w_0, r) = E^{\|\cdot\|_2}(w_{\beta^*}(w_0), w_0, r)$$

Then by considering  $t_{opt} := 1/\beta^*$ , we obtain that

$$E^{\|\cdot\|_2}(w(t_{opt}, w_0), w_0, r) \leq E^{\|\cdot\|_2}(w_{1/t_{opt}}(w_0), w_0, 1.2985r) \leq 1.2985^2 E^{\|\cdot\|_2}(w_0, r)$$

where the first inequality follows from Lemma C.1 and the second from Lemma C.2, which gives the desired result.  $\square$

### C.3 Uniform Shrinkage of $w_0$ is Sub-optimal in General

In the case where the isotropic case where  $\Sigma = I_d$  and the attacker's norm is Euclidean, we can show that GD is optimally robust and so for any radius  $r$ . In that case, the GD scheme is simply a uniform shrinkage of the generative model  $w_0$  as we have that

$$w(t, w_0) = (1 - \exp(-t))w_0.$$

Let us now recall our statement for convenience.

**Proposition 4.3.** *Assume that  $\Sigma = I_d$ , then for all  $r \geq 0$ , we have  $\inf_{t \geq 0} E^{\|\cdot\|_2}(w(t, w_0), r) = E_{opt}^{\|\cdot\|_2}(w_0, r)$ .*

*Proof.* First recall, from Proposition 4.2 that for  $r \leq \sqrt{2/\pi}$  or  $r \geq \sqrt{\pi/2}$ , we have already shown the result. Now observe that in the isotropic case, the ridge solution is in fact also a uniform shrinkage of  $w_0$ , as we have  $w_{\beta}(w_0) = (\beta I_d + \Sigma)^{-1} \Sigma w_0 = \frac{1}{\beta+1} w_0$ , and therefore the optimality of the GD follows directly from (Xing et al., 2021, Proposition 1).  $\square$

The next result shows that in the anisotropic case where  $\Sigma \neq I_d$ , uniform shrinkage can achieve arbitrarily bad adversarial risk, while GD continues to be optimal (up to within an absolute multiplicative constant), thanks to Proposition 4.2.

**Proposition C.1.** *For sufficiently large values of the input-dimension  $d$ , it is possible to construct a feature covariance matrix  $\Sigma \in \mathcal{S}_d^{++}(\mathbb{R})$  and generative model  $w_0 \in \mathbb{R}^d$  such that  $\text{tr}(\Sigma), \|w_0\|_{\Sigma} = \Theta(1)$  and*

$$E_{shrink}^{\|\cdot\|_2}(w_0, 1/\sqrt{d}) = \Omega(1), \tag{24}$$

$$E_{opt}^{\|\cdot\|_2}(w_0, 1/\sqrt{d}) = o(1). \tag{25}$$

*Proof.* Fix  $\beta \in (1, \infty)$ , and any integer  $k \geq 1$ , set  $\lambda_k := k^{-\beta}$ . For a large positive integer  $d$ , consider the choice of the feature covariance matrix  $\Sigma = \text{diag}(\lambda_1, \dots, \lambda_d)$ . Thus, the eigenvalues  $\lambda_k$  of  $\Sigma$  decay polynomially as a function of their rank  $k$ , and  $\Sigma$  is very badly conditioned (its condition number is  $d^{\beta}$ , a nontrivial polynomial in the input-dimension  $d$ ). Consider a random choice of the generative model  $w_0 \sim N(0, I_d)$ . We will prove the following stronger statement

**Claim.** *For large values of the input-dimension  $d$  and the above choice of feature covariance matrix  $\Sigma$ , the estimates (24) and (25) hold w.h.p over the choice of  $w_0 \sim N(0, I_d)$ .*

First observe that, by elementary concentration it holds w.p  $1 - e^{-C^d}$  over the choice of  $w_0 \sim N(0, I_d)$  that

$$\begin{aligned} \|w_0\|_2^2 &\asymp \mathbb{E}\|w_0\|_2^2 = d, \\ \|w_0\|_{\Sigma}^2 &\asymp \mathbb{E}\|w_0\|_{\Sigma}^2 = \text{tr}(\Sigma) = \sum_{k=1}^d k^{-\beta} = \Theta(1). \end{aligned}$$

Since  $E_{shrink}^{\|\cdot\|_2}(w_0, r) = \|w_0\|_{\Sigma}^2 \wedge r^2 \|w_0\|_2^2$ , we deduce that for any  $r \geq 0$ , the following holds w.p  $1 - e^{-C^d}$ .

$$E_{shrink}^{\|\cdot\|_2}(w_0, r) \asymp 1 \wedge r^2 d. \tag{26}$$

On the other hand, thanks to Lemma 3.1 one has

$$\sqrt{E_{opt}^{\|\cdot\|_2}(w_0, r)} \asymp \inf_{w \in \mathbb{R}^d} \|w - w_0\|_{\Sigma} + r\|w\|_2 = \sup_{z \in S} z^{\top} w_0, \quad (27)$$

where  $S := B_{\|\cdot\|_{\Sigma^{-1}}}^d \cap rB_2^d = \{z \in \mathbb{R}^d \mid \|z\|_2 \leq r, \|z\|_{\Sigma^{-1}} \leq 1\}$ . In particular, we deduce that  $\mathbb{E}\sqrt{E_{opt}^{\|\cdot\|_2}(w_0, r)} \asymp \omega(S)$ , i.e the *Gaussian width* of  $S$ , defined by

$$\omega(S) = \mathbb{E} \left[ \sup_{z \in S} z^{\top} w_0 \right]. \quad (28)$$

Observe that, the function  $G : a \mapsto \sup_{z \in S} z^{\top} a$  is  $r$ -Lipschitz w.r.t to the Euclidean norm. Indeed, for any  $a, b \in \mathbb{R}^d$ ,

$$|G(a) - G(b)| = \left| \sup_{z \in S} z^{\top} a - \sup_{z \in S} z^{\top} b \right| \leq \sup_{z \in S} |z^{\top} (a - b)| \leq \sup_{z \in rB_2^d} |z^{\top} (a - b)| = r\|a - b\|_2.$$

Therefore, by concentration of Lipschitz functions of Gaussians, we have for any positive  $t$ ,

$$\left| \sup_{z \in S} z^{\top} w_0 - \omega(S) \right| \leq t, \text{ w.p } 1 - 2e^{-t^2 d / (2r^2)}. \quad (29)$$

Subsequently, we will take  $t$  sufficiently small.

We now upper-bound  $\omega(S)$ . Referring to (Wei et al., 2020, Example 4)<sup>2</sup>, one observes that

$$\omega(S) \asymp r^{1-1/\beta}. \quad (30)$$

Combining (26), (30), and (29) applied with  $t = \omega(S) \asymp r^{1-1/\beta}$ , we obtain that if  $r \lesssim 1/\sqrt{d}$ , then it holds w.p  $1 - e^{-Cd} - 2e^{-d/(2r^{2/\beta})} = 1 - o(1)$  over the choice of  $w_0 \sim N(0, I_d)$  that

$$E_{shrink}^{\|\cdot\|_2}(w_0, r) \asymp 1 \wedge r^2 d = r^2 d, \quad (31)$$

$$E_{opt}^{\|\cdot\|_2}(w_0, r) \asymp r^{2(1-1/\beta)}. \quad (32)$$

Plugging  $r = 1/\sqrt{d}$  in the above then completes the proof of the claim.  $\square$

#### C.4 Proof of Proposition 4.4: Sub-optimality of GD in Case of Mahalanobis Attacks

**Proposition 4.4.** *Let  $d = 2$ ,  $\Sigma = I_d$  and for any integer  $m \geq 1$ , let us consider the following positive-definite matrix*

$$B = B(m) = \begin{pmatrix} 1/m & 0 \\ 0 & m \end{pmatrix}. \quad (14)$$

Also, consider the following choice of generative model  $w_0 = w_0(m) = (1/\sqrt{m}, 1)$ . Then, for any fixed  $r > 0$ ,

$$\lim_{m \rightarrow +\infty} \frac{\inf_{t \geq 0} E^{\|\cdot\|_B}(w(t, w_0), w_0, r)}{E_{opt}^{\|\cdot\|_B}(w_0, r)} = +\infty.$$

*Proof.* Let us first compute a lower bound on the optimal adversarial risk along the path of the vanilla GD. Indeed, first note that from Lemma 3.1, one obtains that

$$\inf_{t \geq 0} \alpha E^{\|\cdot\|_B}(w(t, w_0), w_0, r) \geq \inf_{t \geq 0} \tilde{E}^{\|\cdot\|_B}(w(t, w_0), w_0, r).$$

In the particular case where  $\Sigma = I_d$ , we have that  $w(t, w_0) := (1 - \exp(-t))w_0$  and therefore we obtain that

$$\inf_{t \geq 0} \tilde{E}^{\|\cdot\|_B}(w(t, w_0), w_0, r) = \min(\|w_0\|_2^2, r^2 \|w_0\|_{B^{-1}}) = \min(1/m + 1, r^2(1 + 1/m))$$

<sup>2</sup>Alternatively, it is easy to show that  $\omega(S) \asymp \sum_{k=1}^d \min(\lambda_k, r^2)$ , and then proceed from there.

Therefore, one has

$$\inf_{t \geq 0} E^{\|\cdot\|_B}(w(t, w_0), w_0, r) \geq \alpha \min(1/m + 1, r^2(1 + 1/m)) \xrightarrow{m \rightarrow \infty} 1/\alpha \min(1, r^2)$$

Let us now compute the value of the adversarial risk proxy  $\tilde{E}(w, w_0, r)$ . To this end, remark that in our specific example, one can write the proxy adversarial risk at  $w = (w_1, w_2)$  as  $\tilde{E}^{\|\cdot\|_B}(w, w_0, r) = (\alpha + r\beta)^2$ , where

$$\alpha = \alpha(w) = \sqrt{(w_1 - 1/\sqrt{m})^2 + (w_2 - 1)^2}, \quad (33)$$

$$\beta = \beta(w) = \sqrt{mw_1^2 + (1/m)w_2^2}, \quad (34)$$

The critical points of  $\tilde{E}^{\|\cdot\|_B}(\cdot, w_0, r)$  are either 0 ( $\beta = 0$ ),  $w_0$  ( $\alpha = 0$ ), or must satisfy  $\nabla_w(\alpha + r\beta) = 0$ , i.e.  $\frac{(w - w_0)}{\alpha} + \frac{rB^{-1}w}{\beta} = 0$ . Let us consider the latter case and let us define  $t := \beta/\alpha$  satisfying the last equation. Then it is equivalent to write  $t(w - w_0) + rB^{-1}w = 0$ . We conclude that the minimum of  $\tilde{E}^{\|\cdot\|_B}(\cdot, w_0, r)$  is of the form

$$w = w(t) := D(t)w_0, \quad (35)$$

where  $D(t) = t(tI_d + rB^{-1})^{-1} = I_d - (tI_d + rB^{-1})^{-1}rB^{-1}$  is the  $d \times d$  diagonal matrix with diagonal entries  $(t/(t + rm), t/(t + r/m))$ . In that case we obtain that

$$\begin{aligned} \tilde{E}_{opt}^{\|\cdot\|_B}(w_0, r) &= \alpha(w(t))^2(1 + rt)^2 \\ &= r^2(1 + rt)^2 \left( \frac{m}{(t + rm)^2} + \frac{1}{m^2(t + r/m)^2} \right) \end{aligned}$$

Under the assumption that: for  $m$  sufficiently large, such  $t = t(m)$  exists and the sequence  $(t(m))_{m \geq 0}$  converges towards an absolute and positive constant  $c$ , we obtain that

$$\tilde{E}_{opt}^{\|\cdot\|_B}(w_0, r) \xrightarrow{m \rightarrow \infty} 0$$

Let us now show that such  $t(m)$  exists and that it does converges towards an absolute constant. Indeed, by definition, showing that  $t$  exists is equivalent to showing that the equation  $\|w(t)\|_{B^{-1}} = t\|w(t) - w_0\|_{\Sigma}$  has a solution, which gives

$$\frac{1}{(t + rm)^2} + \frac{1}{m(t + r/m)^2} = r^2 \left[ \frac{m}{(t + rm)^2} + \frac{1}{m^2(t + r/m)^2} \right] \quad (36)$$

Therefore as soon as  $r^2/m < 1$  and  $r^2m > 1$ , we obtain that such  $t$  exists and satisfies

$$t = \frac{r\sqrt{m - r^2} - r/m\sqrt{r^2m - 1}}{\sqrt{r^2m - 1} - 1/m\sqrt{m - r^2}} \xrightarrow{m \rightarrow \infty} 1$$

and the result follows.  $\square$

## D Proofs of Results of Section 5: An Adapted Gradient-Descent (GD+)

### D.1 Proof of Proposition 5.1: Optimality of GD+ in Case of Mahalanobis Norm Attacks

**Proposition 5.1.** *For any  $B \in \mathcal{S}_d^{++}(\mathbb{R})$  and  $r \geq 0$ ,*

$$\inf_{t \geq 0} E^{\|\cdot\|_B}(B^{1/2}w^{B^{1/2}}(t, w_0), r) \leq \beta E_{opt}^{\|\cdot\|_B}(r, w_0).$$

*Proof.* According to Proposition 4.2, we already know that early-stopped GD scheme on data with covariance  $\Sigma$  generated by a model  $w_0$  achieves adversarial risk which is optimal for Euclidean /  $\ell_2$  attacks, up to within a multiplicative factor of  $\beta$ . Now remark that by considering our transformation on the data  $\tilde{x} := Mx$ , we are defining a GD scheme on data

with covariance  $M\Sigma M^\top$  generated by  $(M^\top)^{-1}w_0$ . We deduce that, early-stopped GD+ achieves the optimal value (up to withing a multiplicative factor of  $\beta$ ) in optimization problem

$$\min_{w \in \mathbb{R}^d} \|w - (M^\top)^{-1}w_0\|_{M\Sigma M^\top}^2 + r^2\|w\|_2^2 + 2rc_0\|w - (M^\top)^{-1}w_0\|_{M\Sigma M^\top}\|w\|_2.$$

Now observe that

$$\|w - (M^\top)^{-1}w_0\|_{M\Sigma M^\top}^2 = \|M^\top w - w_0\|_\Sigma^2$$

and therefore, by applying again a change of variables  $\tilde{w} = M^\top w$ , we obtain that, along the path drawn by  $(M^\top w^M(t, w_0))_{t \geq 0}$ , there exists an optimal solution (up to within a multiplicative factor of  $\beta$ ) for

$$\min_{w \in \mathbb{R}^d} \|w - w_0\|_\Sigma^2 + r^2\|(M^\top)^{-1}w\|_2^2 + 2rc_0\|w - w_0\|_\Sigma\|(M^\top)^{-1}w\|_2.$$

Finally, taking finally  $M = B^{1/2}$  gives the desired result, since the dual norm of  $\|\cdot\|_B$  is  $\|\cdot\|_{B^{-1}}$ .  $\square$

## D.2 Proof of Proposition 5.2: Case of Optimality of GD+ in Case of General Norm Attacks

**Proposition 5.2.** *For any  $r \geq 0$ , and invertible matrix  $M \in \mathcal{M}_d(\mathbb{R})$ , it holds that*

$$\frac{\inf_{t \geq 0} E^{\|\cdot\|}(M^\top w^M(t, w_0), w_0, r)}{E_{opt}^{\|\cdot\|}(r, w_0)} \leq \beta \kappa^{\|\cdot\|} ((M^\top)^{-1})^2.$$

*Proof.* Recall the definition of  $\lambda_{min}^{\|\cdot\|}$  and  $\lambda_{max}^{\|\cdot\|}$  from the beginning of Section 5.1 of the main paper. Observe that for all  $w \in \mathbb{R}^d$

$$\lambda_{min}^{\|\cdot\|}((M^\top)^{-1})\|w\|_* \leq \|(M^\top)^{-1}w\|_2 \leq \lambda_{max}^{\|\cdot\|}((M^\top)^{-1})\|w\|_*. \quad (37)$$

In the following we also define  $\|\cdot\|^M = \|M \cdot\|_2$ . Then, one computes

$$\begin{aligned} \inf_t E^{\|\cdot\|}(Mw^M(t, w_0), w_0, r) &\leq \inf_t E^{\|\cdot\|^M}(Mw^M(t, w_0), w_0, r/\lambda_{min}^{\|\cdot\|}((M^\top)^{-1}))(Mw^M(t, w_0)) \\ &\leq \beta \inf_{w \in \mathbb{R}^d} E^{\|\cdot\|^M}(w, w_0, r/\lambda_{min}^{\|\cdot\|}((M^\top)^{-1})) \\ &\leq \beta \inf_{w \in \mathbb{R}^d} E^{\|\cdot\|}(w, w_0, r\kappa^{\|\cdot\|}((M^\top)^{-1})) \\ &\leq \beta \kappa^{\|\cdot\|}(((M^\top)^{-1})^2) \inf_w E^{\|\cdot\|}(w, w_0, r) \end{aligned}$$

where the first inequality follows from the LHS of Eq. (37), the second from Proposition 5.1, the third from RHS of Eq. (37), and the last one from Lemma C.2.  $\square$

## D.3 Proof of Theorem 5.1: A Case of Optimality of GD for Non-Euclidean / General Norm Attacks

**Theorem 5.1.** *Suppose Condition 5.1 is in order. Then, for any positive  $r$ , it holds for  $M = \Sigma^{-1}$  that*

$$\frac{\inf_{t \geq 0} E^{\|\cdot\|}(M^\top w^M(t, w_0), w_0, r)}{E_{opt}^{\|\cdot\|}(r, w_0)} \leq (1 \vee 1/c^2)\alpha.$$

*Proof.* Let  $g(w_0) \in \mathbb{R}^d$  be a subgradient of  $\|\cdot\|_*$  at  $w_0$ . First note that  $\|g(w_0)\| \leq 1$  and  $g(w_0)^\top w_0 = \|w_0\|_*$  and let us denote

$$\lambda := \frac{\|g(w_0)\|}{\|g(w_0)\|_{\Sigma^{-1}}} \quad (38)$$

Let us now show the following useful Proposition.

**Proposition D.1.** *Let  $w_0 \in \mathbb{R}^d$  and  $r > 0$ . Then we have that*

$$\sqrt{\alpha E_{opt}^{\|\cdot\|}(w_0, r)} \geq \sup_{z \in B_2^d \cap r B_{\|\Sigma^{1/2}\cdot\|}^d} \langle z, \Sigma^{1/2} w_0 \rangle$$

*Proof.* Let us define the support function of a subset  $C$  of  $\mathbb{R}^d$  as the function  $\gamma_C : \mathbb{R}^d \rightarrow (-\infty, \infty]$ , defined by  $\gamma_C(x) := \sup_{w \in \mathbb{R}^d} x^\top w$  and let us introduce the following useful lemma.

**Lemma D.1.** *Let  $A$  and  $B$  be two subsets of topological vector space  $V$ , such that the intersection of the relative interiors of  $A$  and  $B$  is nonempty. Then,  $\gamma_{A \cap B} = \gamma_A \star \gamma_B$ , where  $f \star g : V \rightarrow (-\infty, +\infty]$  is infimal convolution of two proper convex lower-semicontinuous functions  $f, g : V \rightarrow (-\infty, +\infty]$ , defined by  $(f \star g)(v) := \inf_{v' \in V} f(v - v') + g(v')$ .*

Now using Lemma 3.1, recall that

$$\begin{aligned} \sqrt{\alpha E_{opt}^{\|\cdot\|}(w_0, r)} &\geq \sqrt{\tilde{E}_{opt}^{\|\cdot\|}(w_0, r)} = \inf_w \|w - w_0\|_\Sigma + r \|w\|_\star \\ &= (\gamma_{B_{\Sigma^{-1}}}^d \star \gamma_{B_{\|\cdot\|}^d})(w_0) \\ &= \gamma_{S_r}(w_0) \text{ where } S_r := B_{\Sigma^{-1}}^d(1) \cap B_{\|\cdot\|}^d(r) \\ &= \gamma_{\tilde{S}_r}(\Sigma^{1/2} w_0) \text{ where } \tilde{S}_r := B_2^d(1) \cap B_{\|\Sigma^{1/2}\cdot\|}^d(r) \end{aligned}$$

from which the result follows.  $\square$

Define now  $z(w_0) := (r \wedge \lambda) \Sigma^{-1/2} g(w_0) \in \mathbb{R}^d$ , and observe that

$$\begin{aligned} \|\Sigma^{1/2} z(w_0)\| &\leq r \|g(w_0)\| \leq r, \\ \|z(w_0)\|_2 &\leq \lambda \|g(w_0)\|_{\Sigma^{-1}} \leq \|g(w_0)\| \leq 1. \end{aligned}$$

Thus,  $z(w_0) \in B_2^d \cap r B_{\|\Sigma^{1/2}\cdot\|}^d$ . Then, thanks to Proposition D.1, it follows that

$$\begin{aligned} \sqrt{\alpha E_{opt}^{\|\cdot\|}(r, w_0)} &\geq \max_{z \in B_2^d \cap r B_{\|\Sigma^{1/2}\cdot\|}^d} z^\top \Sigma^{1/2} w_0 \\ &\geq z(w_0)^\top \Sigma^{1/2} w_0 = (r \wedge \lambda) g(w_0)^\top w_0 \\ &= (r \wedge \lambda) \|w_0\|_\star, \end{aligned} \tag{39}$$

where  $\alpha := 1/(1 + \sqrt{2/\pi}) \approx 1.1124$ . Let us now define

$$c := \frac{\lambda \|w_0\|_\star}{\|w_0\|_\Sigma}$$

then  $\alpha E_{opt}^{\|\cdot\|}(w_0, r) \geq c^2 \|w_0\|_\Sigma^2 \wedge r^2 \|w_0\|_\star^2 \geq (1 \wedge c^2) (\|w_0\|_\Sigma^2 \wedge r^2 \|w_0\|_\star^2)$  which matches  $\inf_{\gamma \in [0,1]} \tilde{E}^{\|\cdot\|}(\gamma w_0, w_0, r)$  attained by the optimal uniform rescaling of  $w_0$  and the result follows.  $\square$

#### D.4 Proof of Proposition 5.3: A Case of Sub-optimality of GD to $\ell_\infty$ -Norm Attacks

**Proposition 5.3.** *Let  $\Sigma = I_d$ , then it is possible to construct  $w_0 \in \mathbb{R}^d$  and  $r > 0$  such that in the limit  $d \rightarrow \infty$ , it holds that  $r \rightarrow 0$ ,  $r\sqrt{d} \rightarrow +\infty$ , and*

$$\frac{\inf_{t \geq 0} E^{\|\cdot\|_\infty}(w(t, w_0), w_0, r)}{E_{opt}^{\|\cdot\|_\infty}(r, w_0)} \rightarrow +\infty.$$

*Proof.* Let  $a, \lambda > 0$ ,  $w_0 := (a + \lambda, \dots, a + \lambda, \lambda, \dots, \lambda)$  and let  $\alpha \geq 1$  be the number of coordinate such that  $(w_0)_i = a + \lambda$ . Concrete values for  $a, \lambda$ , and  $\alpha$  will be prescribed later. For such model we obtain that

$$\begin{aligned} \|w_0\|_1 &= d\lambda + \alpha a \text{ and} \\ \|w_0\|_2 &= \sqrt{\alpha(a + \lambda)^2 + (d - \alpha)\lambda^2} \\ &= \sqrt{\alpha a^2 + d\lambda^2 + 2\alpha a\lambda}. \end{aligned}$$

Therefore the shrinkage model obtains an error equals to:

$$E_{shrink}^{\|\cdot\|_\infty}(w_0, r) = \min(\alpha a^2 + d\lambda^2 + 2\alpha a\lambda, r^2(d\lambda + \alpha a)^2).$$

Now let  $w := (\underbrace{a, \dots, a}_{\alpha \text{ times}}, 0, \dots, 0)$ , we obtain that

$$\sqrt{E_{opt}^{\|\cdot\|_\infty}(w_0, r)} \leq \|w - w_0\|_2 + r\|w\|_1 = \sqrt{d}\lambda + r\alpha a$$

Let us now consider the case where

$$r = \lambda = \frac{1}{d^{1/4} \log d}, \quad \alpha = d^{1/2}, \quad a = 1.$$

Then, for large  $d$  we obtain that

$$\begin{aligned} \|w - w_0\|_2 + r\|w\|_1 &\asymp \frac{d^{1/4}}{\log d} \\ r\|w_0\|_1 &\asymp \frac{d^{1/2}}{(\log d)^2}, \quad \|w_0\|_2 \asymp d^{1/4} \end{aligned}$$

from which follows that

$$\frac{E_{shrink}^{\|\cdot\|_\infty}(w_0, r)}{E_{opt}^{\|\cdot\|_\infty}(w_0, r)} \geq \frac{\min(d^{1/2}, \frac{d}{(\log d)^4})}{\frac{d^{1/2}}{(\log d)^2}} \asymp (\log d)^2 \rightarrow +\infty,$$

as claimed. □

## D.5 Proof of Theorem 5.2: Case of Optimality of GD+ to $\ell_p$ -Norm Attacks

**Theorem 5.2.** *Let  $\Sigma$  be any definite positive diagonal matrix and  $p \in [1, +\infty]$ , then we have*

$$\inf_{M \in \mathcal{M}_d(\mathbb{R}), t \geq 0} \frac{E^{\|\cdot\|_p}(M^\top w^M(t, w_0), w_0, r)}{E_{opt}^{\|\cdot\|_p}(r, w_0)} \leq \alpha.$$

For the proof, we will need the following useful lemma.

**Lemma D.2.** *Given an extended-value convex function (EVCF)  $g : \mathbb{R}^d \rightarrow (-\infty, \infty]$ , every minimizer of the EVCF  $g_a : \mathbb{R}^d \rightarrow (-\infty, \infty]$  defined by  $g_a(x) := \|x - a\|_\Sigma + g(x)$  can be written as  $x = \text{prox}_{\Sigma, tg}(a)$ , for some  $t \in [0, \infty]$ .*

In the above Lemma, we consider the proximal operator

$$\text{prox}_{\Sigma, tg}(a) := \arg \min_{x \in \mathbb{R}^d} (1/2)\|x - a\|_\Sigma^2 + tg(x), \quad (40)$$

defined w.r.t to the Bregman divergence on  $\mathbb{R}^d$  given by  $D(z, z') := (1/2)\|z - z'\|_\Sigma^2$ .

*Proof.* Let  $x \in \mathbb{R}^d$  be a minimizer of  $g_a$ . If  $x \in \{0, a\}$ , then we are done. Otherwise, set  $t = t(x) := \|x - a\|_\Sigma$ . Note that by first-order optimality conditions,  $x$  is a minimizer of  $g_a$  iff  $0 \in \partial g_a(x)$  iff  $0 \in \Sigma(x - a)/t + \partial g(x)$  iff  $0 \in (x - a) + t\Sigma^{-1}\partial g(x)$  iff  $x = (I + t\Sigma^{-1}\partial g)^{-1}(a) = \text{prox}_{\Sigma, tg}(a)$ , by definition of the proximal operator. Indeed recall that  $\partial g$  is a maximal monotone operator (a result of Rockafellar from the 70's [Rockafellar \(1970\)](#)), so that  $(I + t\Sigma^{-1}\partial g)^{-1}(a)$  is well-defined and single-value. □

Let us now introduce another important Lemma in order to show the result.

**Lemma D.3.** *For any  $q \in [1, 2]$ ,  $a \in \mathbb{R}^d$ , and  $\Sigma$  positive definite diagonal matrix, we have*

$$\begin{aligned} |\text{prox}_{\Sigma, \|\cdot\|_q}(a)| &\leq |a| \text{ and} \\ \text{sign}(\text{prox}_{\Sigma, \|\cdot\|_q}(a)) &= \text{sign}(a) \end{aligned}$$

where the  $|\cdot|$  and  $\text{sign}$  are the coordinate-wise operators.

*Proof.* A simple computation reveals that

$$\text{prox}_{\Sigma, g} = \Sigma^{-1/2} \circ \text{prox}_{g \circ \Sigma^{-1/2}} \circ \Sigma^{1/2},$$

for any EVCF  $g$ . Using the Moreau decomposition formula, we can further express

$$\begin{aligned} \text{prox}_{\Sigma, g}(a) &= \Sigma^{-1/2} [\Sigma^{1/2} a - \text{prox}_{(g \circ \Sigma^{-1/2})^*}(\Sigma^{1/2} a)] \\ &= a - \Sigma^{-1/2} \text{prox}_{g^* \circ \Sigma^{1/2}}(\Sigma^{1/2} a), \end{aligned}$$

where  $g^*$  is the convex conjugate of  $g$ , and we have used the fact that  $(g \circ \Sigma^{-1/2})^* = g^* \circ \Sigma^{1/2}$ . Now, taking  $g = \|\cdot\|_q$ , one notes that  $g^*$  is the indicator function of the unit-ball  $B_{\|\cdot\|_p}^d$  for the norm  $\|\cdot\|_p$ , where  $p = p(q) \in [2, \infty]$  is the harmonic conjugate of  $q$ . Therefore, we obtain that

$$\text{prox}_{\Sigma, g}(a) = a - \Sigma^{-1/2} \text{proj}_{B_{\|\Sigma^{1/2} \cdot\|_p}^d}(\Sigma^{1/2} a).$$

Let us assume without loss of generality that all the  $a_i \geq 0$ . and let us show now that

$$0 \leq \Sigma^{-1/2} \text{proj}_{B_{\|\Sigma^{1/2} \cdot\|_p}^d}(\Sigma^{1/2} a) \leq a$$

First note that for all the coordinates of the vector  $\Sigma^{-1/2} \text{proj}_{B_{\|\Sigma^{1/2} \cdot\|_p}^d}(\Sigma^{1/2} a)$  are nonnegative. Indeed, if some coordinate of the projection were negative, then by simply replacing it by 0, we could reduce the total cost of projection and therefore contradicts its optimality. Then using the fact that  $\Sigma$  is diagonal, we want to show that for any  $v \in \mathbb{R}^d$

$$0 \leq \text{proj}_{B_{\|\Sigma^{1/2} \cdot\|_p}^d}(\Sigma^{1/2} a) \leq \Sigma^{1/2} a$$

where the inequality is coordinate-wise. Now remarks that we have

$$\text{proj}_{B_{\|\Sigma^{1/2} \cdot\|_p}^d}(\Sigma^{1/2} a) = \Sigma^{-1/2} \text{proj}_{B_{\|\cdot\|_p}^d}(\Sigma a)$$

where

$$\text{proj}_{B_{\|\cdot\|_p}^d}^{\Sigma^{-1}}(v) := \text{argmin}_{\|x\|_p \leq 1} \|x - v\|_{\Sigma^{-1}}$$

Therefore for any  $p \geq 2$ , the first order condition gives on  $x := \text{proj}_{B_{\|\cdot\|_p}^d}^{\Sigma^{-1}}(\Sigma a)$  that for all coordinate  $i \in [d]$ ,

$$(1/\lambda_i)(x_i - \lambda_i a_i) + \nu x_i^{p-1} = 0, \quad x_i \geq 0,$$

for some  $\nu \geq 0$ . The above further gives

$$(1/\lambda_i)x_i + \nu x_i^{p-1} = a_i, \quad \forall i \in [d].$$

It follows that  $\tilde{x} := \text{proj}_{B_{\|\Sigma^{1/2} \cdot\|_p}^d}(\Sigma^{1/2} a)$  must satisfy:

$$1/(\sqrt{\lambda_i})\tilde{x}_i + \nu(\sqrt{\lambda_i}\tilde{x}_i)^{p-1} = a_i$$

which is equivalent to

$$\tilde{x}_i + \sqrt{\lambda_i}\nu(\sqrt{\lambda_i}\tilde{x}_i)^{p-1} = \sqrt{\lambda_i}a_i$$

However if  $\tilde{x}_i > \sqrt{\lambda_i}a_i$ , then by nonnegativity of  $\sqrt{\lambda_i}\nu(\sqrt{\lambda_i}\tilde{x}_i)^{p-1}$  we obtain a contraction and the result follows.  $\square$

The following is a generalization to the previous lemma.

**Lemma D.4.** *Let  $C$  be a closed convex subset of  $\mathbb{R}^d$  which is symmetric about the coordinate axes. Let  $a \in \mathbb{R}^d$  and set  $b = \text{proj}_C(a)$ . Then, it holds coordinate-wise that  $\text{sign}(b) = \text{sign}(a)$  and  $|b| \leq |a|$ .*

For example, if  $\Sigma$  is diagonal and the attacker's norm  $\|\cdot\|$  is coordinate-wise even (meaning that  $\|x\| = \|z\|$  whenever  $z$  is obtained from  $x$  by flipping the sign of some or no coordinate), then the ball  $B_{\|\Sigma^{1/2}\cdot\|}^d$  is symmetric about the coordinate axes, and we recover the previous lemma.

*Proof of Lemma D.4.* Recall the Kolmogorov characterization of the equation  $b = \text{proj}_C(a)$ , namely

$$(a - b)^\top (c - b) \leq 0, \forall c \in C. \quad (41)$$

Now fix an index  $i \in [d]$ . If  $b_i = a_i$ , there is nothing to show. Otherwise, suppose  $a_i \geq 0$ . Then  $b_i \geq 0$ , since  $C$  is symmetric about the  $i$ th axis in  $\mathbb{R}^d$  (see details further below). Also,  $C$  must contain  $b - te_i$ , for sufficient small  $t > 0$ ; otherwise  $b_i = a_i$ . Here,  $e_i$  is the  $i$ th standard basis vector in  $\mathbb{R}^d$ . Applying (41) with  $c = b - te_i$  then gives  $(a - b)^\top (-te_i) \leq 0$ , i.e.  $-ta_i + tb_i \leq 0$ , i.e.  $0 \leq b_i \leq a_i$ . Similarly, if  $a_i \leq 0$ , consider  $c = b + te_i$  instead, to deduce  $a_i \geq b_i \geq 0$ . We conclude that  $|b_i| \leq |a_i|$  as claimed.

To conclude the proof we need to provide one omitted detail, namely, that  $b_i$  has the same sign as  $a_i$ . We proceed by *reductio ad absurdum*. So, suppose  $a_i b_i < 0$ . Let  $b'$  be obtained from  $b$  by flipping the sign  $i$ th coordinate. Because  $C$  is symmetric around the  $i$ th axis, it must contain  $b'$ . One then computes

$$\|b' - a\|_2^2 - \|b - a\|_2^2 = (b_i - a_i)^2 - (b_i + a_i)^2 = (b_i + a_i)^2 - (b_i - a_i)^2 = 4a_i b_i < 0,$$

which contradicts the optimality of  $b$  as the point of  $C$  which is closest to  $a$ .  $\square$

We are now ready to prove Theorem 5.2.

*Proof of Theorem 5.2.* Because  $\Sigma := \text{diag}(\lambda_1, \dots, \lambda_d)$  is diagonal let  $M = \text{diag}(\sigma_1, \dots, \sigma_d)$  with  $\sigma_i > 0$  and let us consider  $B = M\Sigma^{-1}$ . In that case we obtain that:

$$\begin{aligned} B^{1/2} w_c^{B^{1/2}}(t, w_0) &:= B^{1/2} (I_d - \exp(-tB^{1/2}\Sigma B^{1/2}t)) B^{-1/2} w_0 \\ &= (I_d - \exp(-tM)) w_0 \end{aligned}$$

Then using Lemma D.3, the minimizer,  $w_{opt}$  of  $w \mapsto \tilde{E}^{\|\cdot\|}(w, w_0, r)$  satisfies in a coordinate-wise manner:

$$|w_{opt}| \leq |w_0|, \text{sign}(w_{opt}) = \text{sign}(w_0).$$

If  $(w_0)_i = 0$  then by the above conditions we have necessarily  $(w_{opt})_i = 0$  and therefore  $\sigma_i > 0$  can be chosen arbitrarily.

Now, let us consider the case where  $(w_0)_i \neq 0$ . Fix  $t > 0$ . If  $|(w_0)_i| > |(w_{opt})_i| > 0$ , we can just take  $\sigma_i$  such that:

$$1 - \exp(-t\sigma_i) = \frac{|(w_{opt})_i|}{|(w_0)_i|}, \text{ i.e. } \sigma_i = \sigma_i(t) = -\frac{1}{t} \log \left( 1 - \frac{|(w_{opt})_i|}{|(w_0)_i|} \right).$$

On the other hand, if  $|(w_0)_i| = |(w_{opt})_i|$ , we can simply consider a positive sequence of  $(\sigma_{i,m})_{m \geq 0}$  such that  $\sigma_{i,m} \xrightarrow{m \rightarrow \infty} \infty$ . Finally, if  $|(w_0)_i| > |(w_{opt})_i| = 0$ , then we can consider a positive sequence of  $(\sigma_{i,m})_{m \geq 0}$  such that  $\sigma_{i,m} \xrightarrow{m \rightarrow \infty} 0$ .

Finally by considering such sequence we obtain that

$$\lim_{m \rightarrow \infty} B_m^{1/2} w^{B_m}(t, w_0) \rightarrow w_{opt}$$

Then, by using the fact that

$$\begin{aligned} \lim_{m \rightarrow \infty} E^{\|\cdot\|_p}(B_m^{1/2} w^{B_m}(t, w_0), w_0, r) &\leq \lim_{m \rightarrow +\infty} \tilde{E}^{\|\cdot\|_p}(B_m^{1/2} w^{B_m}(t, w_0), w_0, r) \\ &= \inf_w \tilde{E}^{\|\cdot\|_p}(w, w_0), w_0, r) \\ &\leq \alpha \inf_w E^{\|\cdot\|_p}(w, w_0), w_0, r), \end{aligned}$$

and the desired result follows.  $\square$



## E Proof of Theorem 6.1: Consistency of Proposed Two-Stage Estimator (Algorithm 1)

**Theorem 6.1.** For all  $r \geq 0$ , it holds that

$$\begin{aligned} \tilde{E}_{opt}^{\|\cdot\|}(w_0, r) &\leq \tilde{E}^{\|\cdot\|}(\hat{w}, w_0, r) \leq \tilde{E}_{opt}^{\|\cdot\|}(w_0, r) + \Delta, \\ E_{opt}^{\|\cdot\|}(w_0, r) &\leq E^{\|\cdot\|}(\hat{w}, w_0, r) \leq \alpha E_{opt}^{\|\cdot\|}(w_0, r) + \Delta, \end{aligned}$$

where  $\alpha := 2/(1+c_0) \approx 1.1124$ ,  $\Delta = O(e_1^2 + e_2^2)$ , where the hidden constant in the big- $O$  is of order  $\max(\|\Sigma\|_{op}^2, \|w_0\|_{\Sigma}^2)$ .

*Proof.* For  $r \geq 0$ , vectors  $w, a \in \mathbb{R}^d$  and a psd matrix  $C \in \mathbb{R}^{d \times d}$ , define  $F(w, a, C, r) := \|w - a\|_C + r\|w\|_*$ . Note that the adversarial risk proxy  $\tilde{E}$  defined in (9) can be written as  $\tilde{E}^{\|\cdot\|}(w, w_0, r) = F(w, w_0, \Sigma, r)^2$ , for all  $w \in \mathbb{R}^d$ . Further more, by Lemma 3.1, we know that

$$E^{\|\cdot\|}(w, w_0, r) \leq F(w, w_0, \Sigma, r)^2 \leq \alpha \cdot E^{\|\cdot\|}(w, w_0, r). \quad (42)$$

Thus, the second line of inequalities in the theorem follows from the first. Thus, we only prove the first line.

So, let  $w_{opt}$  be any minimizer of  $F(w, w_0, \Sigma, r)$  over  $w \in \mathbb{R}^d$ . Observe that

$$\begin{aligned} F(\hat{w}, w_0, \Sigma, r) - F(w_{opt}, w, w_0, \Sigma, r) &= F(\hat{w}, w_0, \Sigma, r) - F(\hat{w}, \hat{w}_0, \hat{\Sigma}, r) + F(\hat{w}, \hat{w}_0, \hat{\Sigma}, r) - F(w_{opt}, \hat{w}_0, \hat{\Sigma}, r) \\ &\quad + F(w_{opt}, \hat{w}_0, \hat{\Sigma}, r) - F(w_{opt}, w_0, \Sigma, r) \\ &\leq F(\hat{w}, w_0, \Sigma, r) - F(\hat{w}, \hat{w}_0, \hat{\Sigma}, r) + F(w_{opt}, \hat{w}_0, \hat{\Sigma}, r) - F(w_{opt}, w_0, \Sigma, r), \end{aligned} \quad (43)$$

where we have used the fact that  $F(\hat{w}, \hat{w}_0, \hat{\Sigma}, r) - F(w_{opt}, \hat{w}_0, \hat{\Sigma}, r) \leq 0$ , since  $\hat{w}$  minimizes  $w \mapsto F(w, \hat{w}_0, \Sigma, r)$  by construction. We deduce that

$$F(\hat{w}, w_0, \Sigma, r) \leq F(w_{opt}, w, w_0, \Sigma, r) + e_4(\hat{w}) + e_4(w_{opt}), \quad (44)$$

where  $e_4(w) := |F(w, w_0, \Sigma, r) - F(w, \hat{w}_0, \hat{\Sigma}, r)|$ . The rest of the proof is divided into two steps.

**Step 1: Controlling  $e_4(w_{opt})$ .** Now, for any  $w \in \mathbb{R}^d$ , repeated application of the triangle-inequality gives

$$\begin{aligned} e_4(w) &:= |F(w, w_0, \Sigma, r) - F(w, \hat{w}_0, \hat{\Sigma}, r)| = \left| \|w - w_0\|_{\Sigma} - \|w - \hat{w}_0\|_{\hat{\Sigma}} \right| \\ &= \left| \|w - w_0\|_{\Sigma} - \|w - \hat{w}_0\|_{\Sigma} + \|w - \hat{w}_0\|_{\Sigma} - \|w - \hat{w}_0\|_{\hat{\Sigma}} \right| \\ &\leq \|\Sigma\|_{op} \cdot e_1 + \|w - \hat{w}_0\|_2 \cdot e_2 \lesssim e_1 + \|w - \hat{w}_0\|_2 \cdot e_2. \end{aligned} \quad (45)$$

Let  $\gamma_r(w_0) := \inf_{w \in \mathbb{R}^d} F(w, w_0, \Sigma, r)$ . We deduce that

$$\begin{aligned} e_4(w_{opt}) &\leq e_1 + \|w_{opt} - \hat{w}_0\|_2 \cdot e_2 \leq e_1 + (\|w_{opt} - w_0\|_2 + \|\hat{w}_0 - w_0\|_2) \cdot e_2 \\ &\leq e_1 + (\gamma_r(w_0) + e_1) \cdot e_2 \lesssim e_1 + \gamma_r(w_0) \cdot e_2 \lesssim e_1 + \|w_0\|_{\Sigma} \cdot e_2, \end{aligned} \quad (46)$$

where the last step is thanks to the fact that  $\gamma_r(w_0) \leq \|w_0\|_{\Sigma}$ .

**Step 2: Controlling  $e_4(\hat{w})$ .** If we take  $w = \hat{w}$  to be a minimizer of the function  $F(w, \hat{w}_0, \hat{\Sigma}, r)$  over  $\mathbb{R}^d$ , then

$$\begin{aligned} \lambda_{\min}(\hat{\Sigma}) \|w - \hat{w}_0\|_{\hat{\Sigma}}^2 &\leq \|w - \hat{w}_0\|_{\hat{\Sigma}}^2 = \|w - \hat{w}_0\|_{\Sigma}^2 + (w - \hat{w}_0)^{\top} (\hat{\Sigma} - \Sigma) (w - \hat{w}_0) \\ &\leq \|w - \hat{w}_0\|_{\Sigma}^2 + \|w - \hat{w}_0\|_{\Sigma}^2 \cdot e_2 \\ &\leq \gamma_r(\hat{w}_0)^2 + \|w - \hat{w}_0\|_{\Sigma}^2 \cdot e_2 \\ &\leq \|\hat{w}_0\|_{\Sigma}^2 + \|w - \hat{w}_0\|_{\Sigma}^2 \cdot e_2. \end{aligned}$$

We deduce that

$$(1 - e_2) \|w - \hat{w}_0\|_{\hat{\Sigma}}^2 \leq \|\hat{w}_0\|_{\Sigma}^2 \leq (\|w_0\|_{\Sigma} + e_1)^2, \quad (47)$$

and so  $\|w - \hat{w}_0\|_2 \lesssim \|w_0\|_{\Sigma} + e_1$ . Plugging this into (45) gives

$$e_4(\hat{w}) \lesssim e_1 + \|w_0\|_{\Sigma} \cdot e_2. \quad (48)$$

Putting things together establishes that

$$F(\hat{w}, w_0, \Sigma, r) \leq F(w_{opt}, w_0, \Sigma, r) + O(e_1) + O(\|w_0\|_{\Sigma} \cdot e_2), \quad (49)$$

from which the result follows.  $\square$