
Probabilistically Robust Learning: Balancing Average- and Worst-case Performance

Alexander Robey¹ Luiz F. O. Chamon² George J. Pappas¹ Hamed Hassani¹

Abstract

Many of the successes of machine learning are based on minimizing an averaged loss function. However, it is well-known that this paradigm suffers from robustness issues that hinder its applicability in safety-critical domains. These issues are often addressed by training against worst-case perturbations of data, a technique known as adversarial training. Although empirically effective, adversarial training can be overly conservative, leading to unfavorable trade-offs between nominal performance and robustness. To this end, in this paper we propose a framework called *probabilistic robustness* that bridges the gap between the accurate, yet brittle average case and the robust, yet conservative worst case by enforcing robustness to most rather than to all perturbations. From a theoretical point of view, this framework overcomes the trade-offs between the performance and the sample-complexity of worst-case and average-case learning. From a practical point of view, we propose a novel algorithm based on risk-aware optimization that effectively balances average- and worst-case performance at a considerably lower computational cost relative to adversarial training. Our results on MNIST, CIFAR-10, and SVHN illustrate the advantages of this framework on the spectrum from average- to worst-case robustness.

1. Introduction

Underlying many of the modern successes of learning is the statistical paradigm of empirical risk minimization (ERM), in which the goal is to minimize a loss function averaged over data (Vapnik, 1999). Although ubiquitous in practice,

¹Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA, USA ²University of California, Berkeley, Berkeley, CA, USA. Correspondence to: Alexander Robey <arobey1@seas.upenn.edu>.

it is now well-known that prediction rules learned by ERM suffer from a severe lack of robustness, which in turn greatly limits their applicability in safety-critical domains (Biggio et al., 2013; Shen et al., 2021). Indeed, this vulnerability has led to a pronounced interest in improving the robustness of modern learning tools (Goodfellow et al., 2014; Madry et al., 2017; Zhang et al., 2019).

To this end, a growing body of work has motivated a learning paradigm known as *adversarial training*, wherein rather than training on the raw data, predictors are trained against worst-case perturbations of data (Goodfellow et al., 2014; Madry et al., 2017). Yet, despite ample empirical evidence showing that adversarial training improves the robustness of learned predictors (Su et al., 2018; Croce et al., 2020; Tang et al., 2021), this framework is not without drawbacks. Indeed, adversarial training is known to be overly conservative (Tsipras et al., 2018; Raghunathan et al., 2019), a property often exhibited by other worst-case approaches ranging from complexity theory (Spielman & Teng, 2004) to robust control (Zhou & Doyle, 1998). Furthermore, there are broad classes of problems for which the sample complexity of learning a robust predictor is arbitrarily large (Cullina et al., 2018; Montasser et al., 2019). Finally, the problem of computing worst-case perturbations of data is nonconvex and underparameterized for most modern learning models including deep neural networks (DNNs).

The fundamental drawbacks of these learning paradigms motivate the need for a new robust learning framework that (i) avoids the conservatism of adversarial robustness without incurring the brittleness of ERM, (ii) provides an interpretable way to balance nominal performance and robustness, and (iii) admits an efficient and effective algorithm. To this end, in this paper we propose a framework called *probabilistic robustness* that bridges the gap between the accurate, yet brittle average-case approach of ERM and the robust, yet conservative worst-case approach of adversarial training. By enforcing robustness to most rather than to all perturbations, we show theoretically and empirically that probabilistic robustness meets the desiderata in (i)–(iii). Indeed, our approach parallels a litany of past work in a variety of fields, including smoothed analysis (Spielman & Teng, 2004) and control theory (Campi & Garatti, 2008), wherein

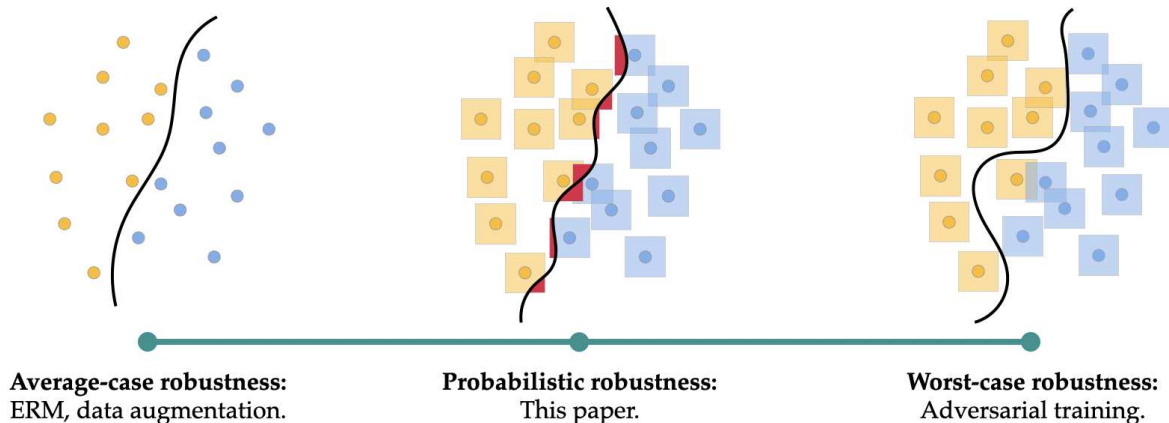


Figure 1. The spectrum of robustness. Illustration of the different decision boundaries engendered by robustness paradigms. *Left:* the two classes (yellow and blue dots) can be separated by a simple decision boundary, though it may not be robust to data perturbations. *Right:* the decision boundary must account for the neighborhood of each data point (yellow and blue boxes), leading to a degraded nominal performance. *Middle:* probabilistic robustness bridges these extremes by allowing a small proportion of perturbations (shown in red) to be misclassified, mitigating the trade-offs between robustness and accuracy both in theory and in practice.

robustness is enforced with high probability rather than in the worst case. In particular, our contributions include:

- **Novel robustness framework.** We introduce *probabilistically robust learning*, a new formulation wherein the goal is to learn predictors that are robust to most rather than to all perturbations (see Fig. 1).
- **(Lack of) Fundamental trade-offs.** We show that in high dimensional settings, the nominal performance of probabilistically robust classifiers is the same as the Bayes optimal classifier, which contrasts with analogous results for adversarially robust classifiers.
- **Sample complexity.** We also show that while the sample complexity of adversarial learning can be arbitrarily high, the sample complexity of our probabilistically robust learning is the same as ERM.
- **Tractable algorithm.** Inspired by risk-aware optimization, we propose a tractable algorithm for probabilistically robust training that spans the full spectrum of robustness (Fig. 1) at a considerably lower computational cost than adversarial training.
- **Thorough experiments.** We provide thorough experiments on MNIST, CIFAR-10, and SVHN. In particular, when we evaluate the ability of algorithms to be robust to 99% of points in ℓ_∞ -balls on CIFAR-10, our algorithm outperforms all baselines by six percentage points.

2. Adversarially Robust Learning

In this paper, we consider the standard supervised learning setting in which data is distributed according to an unknown joint distribution \mathcal{D} over instance-label pairs (x, y) , with instances x drawn from $\mathcal{X} \subseteq \mathbb{R}^d$ and labels y drawn from $\mathcal{Y} \subseteq \mathbb{R}$; in particular, for classification problems we let $\mathcal{Y} =$

$\{1, \dots, K\}$. Our goal is to obtain a hypothesis $h : \mathcal{X} \rightarrow \mathcal{Y}$ belonging to a given hypothesis class \mathcal{H} that correctly predicts the corresponding label y for each instance x . One common approach to this problem is to minimize a suitably-chosen loss function $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}_+$ (e.g., the 0-1, cross-entropy, or squared loss) on average over \mathcal{D} . Explicitly,

$$\min_{h \in \mathcal{H}} \text{SR}(h) \triangleq \mathbb{E}_{(x,y) \sim \mathcal{D}} [\ell(h(x), y)]. \quad (\text{P-NOM})$$

Here $\text{SR}(h)$ denotes the standard risk or *nominal performance* of h .¹ The hypothesis class \mathcal{H} is often comprised of models f_θ parameterized by a vector θ drawn from a compact set $\Theta \subset \mathbb{R}^p$, e.g., linear classifiers or deep neural networks with bounded parameters.

Because the distribution \mathcal{D} is unknown, the objective in (P-NOM) cannot be evaluated in practice. The core idea behind ERM is to use samples (x_j, y_j) drawn i.i.d. from \mathcal{D} to estimate the expectation:

$$\min_{h \in \mathcal{H}} \frac{1}{N} \sum_{j=1}^N \ell(h(x_j), y_j). \quad (\text{P-ERM})$$

One of the fundamental problems in learning theory is to establish the number N of i.i.d. samples needed for (P-ERM) to approximate the value of (P-NOM) with high probability. Problems for which N is finite are called probably approximately correct (PAC) learnable (Vapnik, 1999).

Pitfalls of ERM. While solving (P-ERM) often yields classifiers that are near-optimal for (P-NOM), there is now overwhelming evidence that these hypotheses are sensitive to

¹We assume that ℓ and h satisfy the integrability conditions needed for the expectation in (P-NOM) to be well-defined.

imperceptible perturbations of their input (Biggio et al., 2013; Szegedy et al., 2013). Explicitly, given an instance x and a solution h^* for (P-ERM), one can often find a small perturbations δ such that $h(x + \delta) \neq h(x) = y$.² This issue has been observed in hypotheses ranging from linear models to complex nonlinear models (e.g., DNNs) and has motivated a considerable body of recent work on robust learning (Goodfellow et al., 2014; Madry et al., 2017; Jalal et al., 2017; Zhang et al., 2019; Kamalaruban et al., 2020; Rebuffi et al., 2021).

2.1. Adversarial robustness

Among the approaches that have been proposed to mitigate the sensitivity of hypotheses to input perturbations, there is considerable empirical evidence suggesting that adversarial training is an effective way to obtain adversarially robust classifiers (Su et al., 2018; Athalye et al., 2018; Croce et al., 2020). In this paradigm, hypotheses are trained against worst-case perturbations of data rather than on the raw data itself, giving rise to a robust counterpart of (P-NOM):

$$\min_{h \in \mathcal{H}} \text{AR}(h) \triangleq \mathbb{E}_{(x,y)} \left[\sup_{\delta \in \Delta} \ell(h(x + \delta), y) \right], \quad (\text{P-ROB})$$

where $\Delta \subset \mathbb{R}^d$ is the set of allowable perturbations and we omit the distribution \mathfrak{D} for simplicity. In (P-ROB), $\text{AR}(h)$ denotes the *adversarial risk* of h . Observe that in contrast to (P-NOM), the objective of (P-ROB) explicitly penalizes hypotheses that are sensitive to perturbations in Δ , thus yielding more robust hypotheses. Numerous principled adversarial training algorithms have been proposed for solving (P-ROB) (Goodfellow et al., 2014; Madry et al., 2017; Kannan et al., 2018) and closely-related variants (Moosavi-Dezfooli et al., 2016; Wong & Kolter, 2018; Wang et al., 2019; Zhang et al., 2019).

Pitfalls of adversarial training. Despite the empirical success of adversarial training at defending against worst-case attacks, this paradigm has several limitations. In particular, it is well-known that the improved adversarial robustness offered by (P-ROB) comes at the cost of degraded nominal performance (Tsipras et al., 2018; Dobriban et al., 2020; Javanmard et al., 2020; Yang et al., 2020). Additionally, evaluating the supremum in (P-ROB) can be challenging in practice, since the resulting optimization problem is nonconcave and underparameterized for modern hypothesis classes, e.g., DNNs (Soltanolkotabi et al., 2018). Finally, from a learning theoretic perspective, there exist hypothesis classes for which (P-NOM) is PAC learnable while (P-ROB) is not, i.e., for which (P-NOM) can be approximated using samples whereas (P-ROB) cannot (Cullina et al., 2018; Montasser

²For conciseness, we focus on perturbations of the form $x \mapsto x + \delta$. However, our results also apply to more general models, such as those in (Robey et al., 2020; Wong & Kolter, 2020).

et al., 2019; Diochnos et al., 2019).

2.2. Between the average and worst case

Aside from the now prevalent framework of adversarial training, many works have proposed alternative methods to mitigate the aforementioned vulnerabilities of learning. A standard technique that dates back to (Holmstrom & Koistinen, 1992) is to use a form of data augmentation (P-NOM):

$$\min_{h \in \mathcal{H}} \mathbb{E}_{(x,y)} \left[\mathbb{E}_{\delta \sim \tau} [\ell(h(x + \delta), y)] \right]. \quad (\text{P-AVG})$$

Here, the inner expectation is taken against a known distribution τ . While many algorithms have been proposed for specific τ (Krizhevsky et al., 2012; Hendrycks et al., 2019; Laidlaw & Feizi, 2019; Chen et al., 2020), they fail to yield classifiers sufficiently robust to small perturbations.

Toward obtaining robust alternatives to (P-AVG), two recent works propose relaxations of (P-ROB) that engender notions of robustness between (P-NOM) and (P-ROB). The first relies on the hierarchy of Lebesgue spaces, i.e.,

$$\min_{h_q \in \mathcal{H}} \mathbb{E}_{(x,y)} \left[\left\| \ell(h_q(x + \delta), y) \right\|_{L^q} \right], \quad (\text{PI})$$

where $\|\cdot\|_{L^q}$ denotes the Lebesgue q -norm taken over Δ with respect to the measure τ (Rice et al., 2021). The second relaxes the supremum using the soft maximum or LogSum-Exp function (Li et al., 2020; 2021):

$$\min_{h_t \in \mathcal{H}} \mathbb{E}_{(x,y)} \left[\frac{1}{t} \log \left(\mathbb{E}_{\delta \sim \tau} \left[e^{t \cdot \ell(h_t(x + \delta), y)} \right] \right) \right]. \quad (\text{PII})$$

While both (PI) and (PII) are strong alternatives to (P-ROB), both suffer from significant practical issues related to optimizing their objectives. More specifically, the objective in (PI) cannot be efficiently computed during training due to the difficulty of evaluating the L^q norm (Rice et al., 2021). And in the case of (PII), the gradient of the objective becomes unstable for large values of t when training DNNs. Furthermore, looking beyond these practical limitations, there is also no clear relationship between the values of q and t and robustness properties of the solutions for (PI) and (PII), making these parameters difficult to choose or interpret. These limitations motivate the need for an alternative formulation of robust learning.

Remark 2.1. *Formally, the limiting cases of (PI) and (PII) are not (P-NOM) and (P-ROB). Indeed, for $q = 1$ and $t \rightarrow 0$, the objectives of both problems approach the objective of (P-AVG). For $q = \infty$ and $t \rightarrow \infty$, the objectives of (PI) and (PII) can be written in terms of the essential supremum*

$$\min_{h_r \in \mathcal{H}} \mathbb{E}_{(x,y)} \left[\text{ess sup}_{\delta \sim \tau} \ell(h_r(x + \delta), y) \right], \quad (\text{PIII})$$

where $\text{ess sup}_{\delta \sim \tau} f(\delta)$ denotes an almost everywhere upper bound of f , i.e., an upper bound except perhaps on a set

of τ -measure zero. Note that the essential supremum is a weaker adversary ($\text{ess sup} \leq \text{sup}$), although for rich enough hypothesis classes, the value of (P-ROB) and (PIII) can be the same (Bungert et al., 2021, Lemma 3.8).

3. Probabilistically robust learning

The discussion in the previous section identifies three desiderata for a new robust learning framework:

- (i) **Interpolation.** The framework should strike a balance between nominal and adversarial performance.
- (ii) **Interpretability.** This interpolation should be precisely controlled by an interpretable parameter.
- (iii) **Tractability.** The framework should admit a computationally tractable training algorithm.

While (PI) and (PII) do achieve (i), neither meets the criteria in (ii) or (iii). On the other hand, the probabilistic robustness framework introduced in this section satisfies all of these desiderata. Moreover, as we show in Section 4, it benefits from numerous theoretical properties.

3.1. A probabilistic perspective on robustness

The core idea behind probabilistic robustness is to replace the worst-case view of robustness with a probabilistic perspective. This idea has a long history in numerous fields, including chance-constrained optimization in operations research (Charnes et al., 1958; Miller & Wagner, 1965) and control theory (Campi & Garatti, 2008; Ben-Tal et al., 2009; Ramponi, 2018; Schildbach et al., 2014) and smoothed analysis in algorithmic complexity theory (Spielman & Teng, 2004). In each of these domains, probabilistic approaches are founded on the premise that a few rare events are disproportionately responsible for the performance degradation and increased complexity of adversarial solutions. In the context of robust learning, this argument is supported by recent theoretical and empirical observations suggesting that low-dimensional regions of small volume in the data space are responsible for the prevalence of adversarial examples (Gilmer et al., 2018; Khoury & Hadfield-Menell, 2018; Shamir et al., 2021). This suggests that because the adversarial training formulation (P-ROB) does not differentiate between perturbations, it is prone to yielding conservative solutions that overcompensate for rare events.

To begin our exposition of probabilistic robustness, first consider the case of the 0-1 loss. Here, adversarial training has shifted focus from the nominal 0-1 loss $\mathbb{I}[h(x) \neq y]$ to the adversarially robust 0-1 loss $\mathbb{I}[\exists \delta \in \Delta \text{ s.t. } h(x + \delta) \neq y]$. Concretely, this robust loss takes value one if there exists any perturbation δ in a neighborhood Δ of a fixed instance x that causes misclassification, and value zero otherwise. Motivated by our discussion in the previous paragraph, we seek a relaxed variant of the robust loss which will allow us




Figure 2. **The ρ -ess sup operator.** In this cartoon, we fix $(x, y) \in \Omega$ to show the perturbation set Δ on the x -axis and the value of $\ell(h(x + \delta), y)$ on the y -axis. The solid line shows the value of $\sup_{\delta \in \Delta} \ell(h(x + \delta), y)$, the least upper bound for $\ell(h(x + \delta), y)$. The dashed line shows, for a fixed $\rho > 0$, the value of $\rho\text{-ess sup}_{\delta \sim \tau} \ell(h(x + \delta), y)$, the smallest number u such that $\ell(h(x + \delta), y)$ takes on values larger than u (shown in red) on a subset (shown in yellow) with volume not exceeding ρ .

to quantify the robust performance of a candidate hypothesis while ignoring regions of insignificant volume in Δ . To do so, we first introduce a probability distribution τ (e.g., the uniform distribution) over Δ to assess the local probability of error $\mathbb{P}_{\delta \sim \tau}[h(x + \delta) \neq y]$ around each instance x . Then, for a fixed tolerance level $\rho \in [0, 1)$, the goal of probabilistic robustness is to minimize the probability that the event $\mathbb{P}_{\delta \sim \tau}[h(x + \delta) = y] < 1 - \rho$ will occur; that is, the goal is to ensure that most perturbations do not cause an instance x to be misclassified. As such, the smaller the value of ρ , the more stringent the requirement on robustness. In this way, under the 0-1 loss, our probabilistically robust learning task can then be formulated as follows:

$$\min_{h \in \mathcal{H}} \mathbb{E}_{(x, y)} \left[\mathbb{I}[\mathbb{P}_{\delta \sim \tau}[h(x + \delta) \neq y] > \rho] \right]. \quad (\text{PIV})$$

It is then straightforward to see that under the 0-1 loss, probabilistically robust learning is an instance of (P-NOM) in the sense that we are minimizing the expectation of a particular loss $\mathbb{I}[\mathbb{P}_{\delta \sim \tau}[h(x + \delta) \neq y] > \rho]$.

3.2. Generalizing to general loss functions

With this intuition in mind, we now generalize (PIV) to arbitrary loss functions. To do so, let (Ω, \mathcal{B}) define a measurable space, where $\Omega = \mathcal{X} \times \mathcal{Y}$ and \mathcal{B} denotes the Borel σ -algebra of Ω . Observe that for fixed $(x, y) \in \Omega$, the supremum $t^* := \sup_{\delta \in \Delta} \ell(h(x + \delta), y)$ from (P-ROB) can be written in epigraph form as

$$t^* = \min_{t \in \mathbb{R}} t \text{ s.t. } \ell(h(x + \delta), y) \leq t \quad \forall \delta \in \Delta. \quad (3.1)$$

This formulation makes explicit the fact that the supremum is the least upper bound of $\ell(h(x + \delta), y)$ (see Fig. 2).

As in the development of (PIV), however, we do not need t to upper bound $\ell(h(x + \delta), y)$ for all $\delta \in \Delta$, but only for a proportion $1 - \rho$ of the volume of Δ . We therefore consider

the following relaxation of (3.1):

$$\begin{aligned} u^*(\rho) &= \min_{u \in \mathbb{R}} u & (3.2) \\ \text{s.t.} \quad & \mathbb{P}_{\delta \sim \tau}[\ell(h(x + \delta), y) \leq u] > 1 - \rho \end{aligned}$$

In contrast to (3.1), the upper bound in (3.2) can ignore perturbations for which $\ell(h(x + \delta), y)$ is large (red regions in Fig. 2) as long as these perturbations occupy a subset of Δ that has probability less than ρ (yellow regions in Fig. 2). Thus, note that for $\rho \geq \rho'$, it holds that

$$u^*(\rho) \leq u^*(\rho') \leq u^*(0) \leq t^*$$

and that $u^*(0)$ is the essential supremum from measure theory. In view of this connection, we call $u^*(\rho)$ in (3.2) the ρ -essential supremum (ρ -ess sup) (dashed line in Fig. 2) and formalize its definition below.

Definition 3.1. Let $(\Omega, \mathcal{B}, \mathfrak{p})$ be a measure space and let $f : \Omega \rightarrow \mathbb{R}$ be a measurable function. Define the set

$$U_\rho = \{u \in \mathbb{R} \mid \mathfrak{p}(f^{-1}(u, \infty)) \leq \rho\}.$$

Then, the ρ -ess sup is defined as³

$$\rho\text{-ess sup}_{x \sim \mathfrak{p}} f(x) = \begin{cases} \inf U_\rho & \rho \in [0, 1) \\ \inf \{f(x) : x \in \text{supp}(\mathfrak{p})\} & \rho = 1 \end{cases}$$

where $\text{supp}(\mathfrak{p})$ denotes the support of \mathfrak{p} .

For a given tolerance level $\rho \in [0, 1)$, probabilistically robust learning can now be formalized in full generality as

$$\min_{h_p \in \mathcal{H}} \text{PR}(h_p; \rho) \triangleq \mathbb{E}_{(x,y)} \left[\rho\text{-ess sup}_{\delta \sim \tau} \ell(h_p(x + \delta), y) \right] \quad (\text{P-PRL})$$

In this problem, τ is defined on the measurable space $(\Delta, \mathcal{B}_\Delta)$, where \mathcal{B}_Δ is the restriction of the σ -algebra \mathcal{B} to Δ . For consistency, we define the probabilistic robustness problem (P-PRL) for $\rho = 1$ as (P-AVG).

By construction, it is clear that (P-PRL) satisfies desideratum (i): For $\rho = 0$, we recover (PIII) and for all $\rho \in (0, 1)$, we obtain a strict relaxation of the robustness criteria in (P-ROB). Furthermore, if τ is symmetric—meaning that the mean and median coincide, as is the case for the uniform and Gaussian distributions—then we recover (P-AVG) for $\rho = 1/2$. However, we note that for general distribution τ the objective of (P-PRL) does not approach that of (P-AVG). In practice, this is inconsequential because we are primarily interested in values of ρ close to zero in order to guarantee

³While we define ρ -ess sup as the infimum for $\rho = 1$, this is done only for consistency as this value will play no significant role in subsequent derivations.

robustness in large neighborhoods of the data. Additionally, as we show in Section 5, the algorithm we put forward to solve (P-PRL) yields solutions that exactly recover the average case. Moreover, we show in Sections 5 and 6 that this algorithm fulfills desideratum (iii).

As for the interpretability of ρ in desideratum (ii), notice that the relaxation in (P-PRL) explicitly minimizes the loss over a neighborhood of τ -measure at least $1 - \rho$ of each data point. Thus, in contrast to (PI) or (PII), this relaxation has a practical interpretation. This interpretability is clearest in the 0-1 loss case (PIV), which effectively minimizes $\mathbb{P}_{(x,y)}[\mathbb{P}_{\delta \sim \tau}[h_p(x + \delta) \neq y] > \rho]$. In this way, probabilistic robustness measures the probability of making an error in a neighborhood of each point and only declares failure if that probability is too large, i.e., larger than ρ . This is in contrast to directly measuring the probability of error as in (P-NOM) or requiring that the probability of failure vanishes as in (PIII).

4. Statistical properties of probabilistic robustness

In this section, we characterize the behavior of probabilistic robustness in different settings to show that, in addition to meeting the practical desiderata enumerated in Section 3, this framework also enjoys significant statistical advantages over its worst-case counterpart. In particular, in line with a myriad of past work (Su et al., 2018; Bhagoji et al., 2019; Dobriban et al., 2020; Javanmard et al., 2020; Cullina et al., 2018; Montasser et al., 2020), we first observe that the security guarantee of adversarial robustness comes at the cost of degraded nominal performance as well as an arbitrarily large sample complexity. However, in stark contrast to these results, we show that even for arbitrarily small ρ , there exists classes of problems for which probabilistic robustness can be achieved with the same sample complexity as classical learning and at a vanishingly small cost in nominal performance relative to the Bayes optimal classifier. In the sequel, we first analyze probabilistically robust learning in the two fundamental settings of binary classification and linear regression (Section 4.1), followed by a learning theoretic characterization of its sample complexity (Section 4.2).

4.1. Nominal performance vs. robustness trade-offs

In this section, we consider perturbation sets of the form

$$\Delta = \{\delta \in \mathbb{R}^d : \|\delta\|_2 \leq \epsilon\} \quad (4.1)$$

for a fixed $\epsilon > 0$ and we let τ be the uniform distribution over Δ . We consider binary classification problems with data distributed as

$$x \mid y \sim \mathcal{N}(y\mu, \sigma^2 I_d), \quad y = \begin{cases} +1 & \text{w.p. } \pi \\ -1 & \text{w.p. } 1 - \pi \end{cases}, \quad (4.2)$$

where $\pi \in [0, 1]$ is the proportion of the $y = +1$ class, I_d is the d -dimensional identity matrix, and $\sigma > 0$ is the within-class standard deviation. We assume without loss of generality that the class means $\pm\mu$ are centered about the origin and, by scaling, that $\sigma = 1$. In this setting, it is well-known that the Bayes optimal classifier is

$$h_{\text{Bayes}}^*(x) = \text{sign}(x^\top \mu - q/2) \quad (4.3)$$

where $q = \ln[(1 - \pi)/\pi]$ (Anderson, 2003). Moreover, (Dobriban et al., 2020) recently showed that the optimal adversarially robust classifier is

$$h_r^*(x) = \text{sign}(x^\top \mu [1 - \epsilon / \|\mu\|_2]_+ - q/2), \quad (4.4)$$

where $[z]_+ = \max(0, z)$. In the following proposition, we obtain a closed-form expression for the optimal probabilistically robust linear classifier.

Proposition 4.1. *Suppose the data is distributed according to (4.2) and let $\epsilon < \|\mu\|_2$. Then, for $\rho \in [0, 1/2]$,*

$$h_p^*(x) = \text{sign}\left(x^\top \mu \left(1 - \frac{v_\rho}{\|\mu\|_2}\right)_+ - \frac{q}{2}\right) \quad (4.5)$$

is the optimal linear solution for (PIV), where v_ρ is the Euclidean distance from the origin to a spherical cap of Δ with measure ρ . Moreover, it holds that

$$\text{PR}(h_p^*; \rho) - \text{SR}(h_{\text{Bayes}}^*) = \begin{cases} O\left(\frac{1}{\sqrt{d}}\right), & \rho \in (0, \frac{1}{2}] \\ O(1), & \rho = 0. \end{cases} \quad (4.6)$$

Concretely, Prop. 4.1 conveys three messages.

Firstly, (4.5) shows that the optimal probabilistically robust linear classifier corresponds to the Bayes classifier with an effective mean $\mu \mapsto \mu(1 - v_\rho/\|\mu\|_2)_+$. Secondly, h_p^* depends on the tolerance level ρ through the measure of a spherical cap of Δ . Indeed, it is straightforward to check that $v_{1/2} = 0$ and $v_0 = \epsilon$, and thus (4.5) recovers h_{Bayes}^* and h_r^* respectively. Thus, in this setting, not only does (P-PRL) interpolate between (P-ROB) and (P-NOM) as ρ varies from 0 to $1/2$, but so too do its optimal solutions.

Finally, (4.6) shows that the best achievable probabilistic robustness is essentially the same as the best achievable nominal performance in high dimensions, regardless of the value of ρ provided that it remains strictly positive. However, in the adversarially robust setting of $\rho = 0$, the gap between robustness and accuracy does not vanish, which lays bare the conservatism engendered by forcing classifiers to account for a small set of rare events. In this way, a phase transition occurs at $\rho = 0$ in the sense that for any $\rho > 0$, the gap between nominal performance and probabilistic robustness vanishes in high dimensions, despite the fact that we protect against an arbitrary proportion $1 - \rho$ of perturbations.

In Appendix A.2, we study the distinct yet related problem of linear regression with Gaussian features. In this setting, we observe exactly the same phenomenon, wherein the trade-off between nominal performance and probabilistic robustness vanishes in high dimensions.

4.2. Sample complexity of probabilistic robustness

From a learning theoretic perspective, the behavior of adversarial learning is considerably different from that of classical learning. Indeed, the sample complexity of adversarial learning, i.e., the number of samples needed for the empirical counterpart of (P-ROB) to approximate its solution with high probability, can often be arbitrarily large relative to (P-NOM) (Cullina et al., 2018; Yin et al., 2019; Montasser et al., 2019). The following proposition shows that unlike in the case of adversarial robustness, the sample complexity of probabilistically robust learning can match that of classical learning even for arbitrarily small $\rho > 0$.

Proposition 4.2. *Let ℓ be the 0-1 loss and let ν be fully supported on Δ and absolutely continuous with respect to the Lebesgue measure. For any constant $\rho_o \in (0, 1/2)$, there exists a hypothesis class \mathcal{H}_o such that the sample complexity of probabilistically robust learning at level ρ is*

$$N = \begin{cases} \Theta(\log_2(1/\rho_o)), & \rho = 0 \\ \Theta(1), & \rho_o \leq \rho \leq 1 - \rho_o \end{cases}$$

In particular, $\Theta(1)$ is the sample complexity of (P-NOM).

A formal statement of this result and the requisite preliminaries are provided in Appendix B. Concretely, Proposition 4.2 shows that there exist learning problems for which the sample complexity of ERM and PRL are the same, and for which the sample complexity of adversarial training is much larger⁴ than PRL and ERM. However, the result also highlights the fact that when protecting against an overwhelmingly large proportion $1 - \rho_o$ of perturbations, PRL can transition from having the sample complexity of classical learning to that of adversarial learning depending on the value of ρ . Note that the hypothesis class \mathcal{H}_o of the problem depends on ρ , meaning that although the sample complexity can depend on ρ , there are still problems for which the sample complexity of PRL is exponentially smaller than adversarial training. This implies that the conservatism of adversarial learning can manifest itself not only in the form of nominal performance degradation (Tsipras et al., 2018), but also in terms of learning complexity.

⁴In fact, in the case of $\rho = 0$, i.e., adversarial robustness, (Montasser et al., 2019) shows that the problem can be unlearnable, i.e., have infinite sample complexity.

Algorithm 1 Probabilistically Robust Learning (PRL)

```

1: Hyperparameters: sample size  $M$ , step sizes  $\eta_\alpha, \eta >$ 
   0, robustness parameter  $\rho > 0$ , neighborhood distribu-
   tion  $\tau$ , num. of inner optimization steps  $T$ , batch size  $B$ 
2: repeat
3:   for minibatch  $\{(x_j, y_j)\}_{j=1}^B$  do
4:     for  $T$  steps do
5:       Draw  $\delta_k \sim \tau$ ,  $k = 1, \dots, M$ 
6:        $g_{\alpha_n} \leftarrow 1 - \frac{1}{\rho M} \sum_{k=1}^M \mathbb{I}[\ell(f_\theta(x_j + \delta_k), y_j) \geq \alpha_j]$ 
7:        $\alpha_j \leftarrow \alpha_j - \eta_\alpha g_{\alpha_j}$ , for  $n = 1, \dots, B$ 
8:     end for
9:      $g \leftarrow \frac{1}{\rho MB} \sum_{j,k} \nabla_\theta [\ell(f_\theta(x_j + \delta_k), y_j) - \alpha_j]_+$ 
10:     $\theta \leftarrow \theta - \eta g$ 
11:  end for
12: until convergence
    
```

5. A tractable, risk-aware algorithm

So far, we have established that probabilistically robust learning has numerous desirable practical and theoretical properties. However, the stochastic, non-convex, non-smooth nature of the ρ -ess sup means that in practice solving (P-PRL) presents a significant challenge. Nevertheless, in this section we show that the ρ -ess sup admits a tight convex upper bound that can be efficiently optimized using stochastic gradient methods. Given this insight, we propose a novel algorithm for probabilistically robust learning which is guaranteed to interpolate between (P-AVG) and (P-ROB).

5.1. A convex upper bound for the ρ -ess sup

Toward obtaining a practical algorithm for training probabilistically robust predictors, we first consider the relationship between probabilistic robustness and risk mitigation in portfolio optimization (Krokhmal et al., 2002). To this end, notice that the ρ -ess sup is closely related to the inverse cumulative distribution function (CDF): If F_z is the CDF of a random variable z with distribution \mathbf{p} , then ρ -ess sup $_{z \sim \mathbf{p}} z = F_z^{-1}(\rho)$. For an appropriately-chosen distribution \mathbf{p} , $F_z^{-1}(\rho)$ is known as the value-at-risk (VaR) in the portfolio optimization literature. However, VaR is seldom used in practice due to its computational and theoretical limitations. Indeed, VaR is often replaced with a tractable, convex upper bound known as the conditional value-at-risk (CVaR) (Rockafellar et al., 2000; Rockafellar & Uryasev, 2002). Concretely, given a function f and a continuous distribution \mathbf{p} , CVaR can be interpreted as the expected value of f on the tail of the distribution, i.e.,

$$\text{CVaR}_\rho(f; \mathbf{p}) = \mathbb{E}_{z \sim \mathbf{p}} [f(z) \mid f(z) \geq F_z^{-1}(\rho)]. \quad (5.1)$$

It is straightforward to show that $\text{CVaR}_0(f; \mathbf{p}) = \mathbb{E}_{z \sim \mathbf{p}} [f(z)]$ and $\text{CVaR}_1(f; \mathbf{p}) = \text{ess sup}_{z \sim \mathbf{p}} f(z)$. In view

of this property, it is not surprising that CVaR is an upper bound on ρ -ess sup, a result we summarize below:

Proposition 5.1 ((Nemirovski & Shapiro, 2007)). *CVaR is the tightest convex upper bound of ρ -ess sup, i.e.,*

$$\rho\text{-ess sup}_{z \sim \mathbf{p}} f(z) \leq \text{CVaR}_{1-\rho}(f; \mathbf{p}) \quad (5.2)$$

with equality when $\rho = 0$ or $\rho = 1$.

5.2. Minimizing the conditional value at risk

The main computational advantage of CVaR is that it admits the following convex, variational characterization:

$$\text{CVaR}_\rho(f; \mathbf{p}) = \inf_{\alpha \in \mathbb{R}} \alpha + \frac{\mathbb{E}_{z \sim \mathbf{p}} [[f(z) - \alpha]_+]}{1 - \rho}. \quad (5.3)$$

Given this form, CVaR can be computed efficiently by using stochastic gradient-based techniques on (5.3). This is the basis of the probabilistically robust training method detailed in Algorithm 1, which tackles the statistical problem

$$\min_{h_p \in \mathcal{H}} \mathbb{E}_{(x,y)} \left[\text{CVaR}_{1-\rho}(\ell(h_p(x + \delta), y); \tau) \right] \quad (\text{P-CVaR})$$

for parameterized, differentiable hypothesis classes $\mathcal{H} = \{f_\theta : \theta \in \Theta\}$. Notice that like (P-ROB), (P-CVaR) is a *composite* optimization problem involving an inner minimization over α to compute CVaR and an outer minimization over θ to train the predictor. However, unlike (P-ROB), the inner problem in (P-CVaR) is *convex* regardless of \mathcal{H} , and moreover the gradient of the objective in (5.3) can be computed in closed form. To this end, in lines 5–6 of Algorithm 1, we compute CVaR via stochastic gradient descent (SGD) by sampling perturbations $\delta_k \sim \tau$ (Thomas & Learned-Miller, 2019). Then, in lines 9–10, we run SGD on the outer problem using an empirical approximation of the expectation based on a finite set of i.i.d. samples $\{(x_j, y_j)\} \sim \mathcal{D}$ as in (P-ERM).

6. Experiments

We conclude our work by thoroughly evaluating the performance of the algorithm proposed in the previous section on three standard benchmarks: MNIST, CIFAR-10, and SVHN. Throughout, we consider the perturbation set $\Delta = \{\delta \in \mathbb{R}^d : \|\delta\|_\infty \leq \epsilon\}$ under the uniform distribution τ ; for MNIST, we use $\epsilon = 0.3$ and for CIFAR-10 and SVHN, we use $\epsilon = 8/255$. Further details concerning hyperparameter selection are deferred to the appendix.

Baseline algorithms. We consider a range of baselines, including three variants of ERM: standard ERM (Vapnik, 1999), tilted ERM (denoted TERM) (Li et al., 2020; 2021), and ERM with data augmentation (denoted ERM+DA) wherein we run ERM on randomly perturbed instances.




Figure 3. **CVaR as a metric for test-time robustness.** We plot the conditional value at risk $\text{CVaR}_{0.95}(\ell(h(x + \delta), y); \tau)$ averaged over the test data points in CIFAR-10, SVHN, and MNIST respectively. Observe that PRL is more effectively able to minimize the objective in (P-CVaR) than any of the baselines.

Furthermore, we compare to the L^q norm-based Hamiltonian Monte Carlo (N-HMC) method of (Rice et al., 2021). We also run various state-of-the-art adversarial training algorithms, including FGSM (Goodfellow et al., 2014), PGD (Madry et al., 2017), TRADES (Zhang et al., 2019), MART (Wang et al., 2019), and DALE (Robey et al., 2021a).

Evaluation metrics. To evaluate the algorithms we consider, for each dataset we record the clean accuracy and the adversarial accuracy against a PGD adversary. We also record the accuracy of each algorithm on perturbed samples in two ways. Firstly, for each data point we randomly draw 100 samples from τ and then record the average accuracy across perturbed samples $x + \delta$; we denote these accuracies by “Aug.” in the relevant tables. And secondly, to explicitly measure probabilistic robustness, we propose the following *quantile accuracy* metric, the form of which follows directly from the probabilistically robust 0-1 loss defined in (PIV):

$$\text{ProbAcc}(\rho) = \mathbb{I}[\mathbb{P}_{\delta \sim \tau}[h(x + \delta) = y] \geq 1 - \rho]. \quad (6.1)$$

In words, this metric describes the proportion of instances which are probabilistically robust with tolerance level ρ , and therefore this will be our primary metric for evaluating probabilistic robustness for a given tolerance level ρ .

Clean, robust, and quantile accuracies. In Tables 1–3, we record the clean, robust, and probabilistic error metrics described above for PRL and a range of baselines. Throughout, the value of ρ was chosen by cross-validation; see Appendix C for details. Given these results, several remarks are in order. Firstly, across each of these tables, it is clear that the PRL algorithm does not incur the same degradation in nominal performance as does adversarial training; indeed,

Table 1. **Classification results for CIFAR-10.**

| Algorithm | Test Accuracy | | | ProbAcc(ρ) | | |
|-----------|---------------|--------------|--------------|-------------------|--------------|--------------|
| | Clean | Aug. | Adv. | 0.1 | 0.05 | 0.01 |
| ERM | 94.38 | 91.31 | 1.25 | 86.35 | 84.20 | 79.17 |
| ERM+DA | 94.21 | 91.15 | 1.08 | 86.35 | 84.15 | 79.19 |
| TERM | 93.19 | 89.95 | 8.93 | 84.42 | 82.11 | 76.46 |
| N-HMC | 85.07 | 84.41 | 3.24 | 79.50 | 77.96 | 74.76 |
| FGSM | 84.96 | 84.65 | 43.50 | 83.76 | 83.50 | 82.85 |
| PGD | 84.38 | 84.15 | 47.07 | 83.18 | 82.90 | 82.32 |
| TRADES | 80.42 | 80.25 | 48.54 | 79.38 | 79.12 | 78.65 |
| MART | 81.54 | 81.32 | 48.90 | 80.44 | 80.21 | 79.62 |
| DALE | 84.83 | 84.69 | 50.02 | 83.77 | 83.53 | 82.90 |
| PRL | 93.82 | 93.77 | 0.71 | 91.45 | 90.63 | 88.55 |

Table 2. **Classification results for SVHN.**

| Algorithm | Test Accuracy | | | ProbAcc(ρ) | | |
|-----------|---------------|--------------|--------------|-------------------|--------------|--------------|
| | Clean | Aug. | Adv. | 0.1 | 0.05 | 0.01 |
| ERM | 94.44 | 94.28 | 2.72 | 92.16 | 91.40 | 89.42 |
| ERM+DA | 94.69 | 94.43 | 2.08 | 92.65 | 92.01 | 89.92 |
| TERM | 91.85 | 91.58 | 18.33 | 89.01 | 88.04 | 85.85 |
| N-HMC | 90.32 | 90.55 | 18.30 | 88.79 | 87.61 | 86.12 |
| FGSM | 80.69 | 85.55 | 32.82 | 80.18 | 78.02 | 74.87 |
| PGD | 91.19 | 91.29 | 44.89 | 90.15 | 89.68 | 83.82 |
| TRADES | 86.16 | 86.47 | 54.89 | 85.09 | 84.76 | 83.82 |
| MART | 90.20 | 90.44 | 45.23 | 89.81 | 88.82 | 84.32 |
| DALE | 93.85 | 93.72 | 51.98 | 92.52 | 91.08 | 89.19 |
| PRL | 95.00 | 94.81 | 3.11 | 93.28 | 92.97 | 91.74 |

Table 3. **Classification results for MNIST.**

| Algorithm | Test Accuracy | | | ProbAcc(ρ) | | |
|-----------|---------------|--------------|--------------|-------------------|--------------|--------------|
| | Clean | Aug. | Adv. | 0.1 | 0.05 | 0.01 |
| ERM | 99.37 | 98.82 | 0.01 | 97.96 | 97.96 | 96.66 |
| ERM+DA | 99.42 | 99.13 | 5.23 | 98.46 | 98.12 | 97.30 |
| TERM | 99.20 | 98.55 | 11.27 | 97.15 | 96.42 | 94.15 |
| N-HMC | 99.33 | 99.25 | 3.91 | 98.85 | 98.71 | 98.23 |
| FGSM | 98.86 | 98.72 | 19.34 | 98.00 | 97.83 | 97.25 |
| PGD | 99.16 | 99.10 | 94.45 | 99.05 | 98.63 | 98.34 |
| TRADES | 99.10 | 99.04 | 94.76 | 98.71 | 98.61 | 98.33 |
| MART | 98.94 | 98.98 | 94.13 | 98.59 | 98.39 | 97.98 |
| PRL | 99.32 | 99.25 | 26.03 | 99.27 | 99.01 | 98.54 |

on CIFAR-10 and MNIST, the clean accuracy of PRL is within one percentage point of ERM, and for SVHN, the clean accuracy of PRL surpasses that of ERM. A second observation is that across these datasets, PRL offers significant improvements in the ProbAcc(ρ) metric. This improvement manifests most clearly on CIFAR-10, wherein PRL improves by more than six percentage points over all baseline algorithms for $\rho = 0.01$. Moreover, the gap between the ProbAcc of PRL and that of the baselines increases as ρ decreases, indicating that PRL is particularly effective for more stringent robustness requirements.

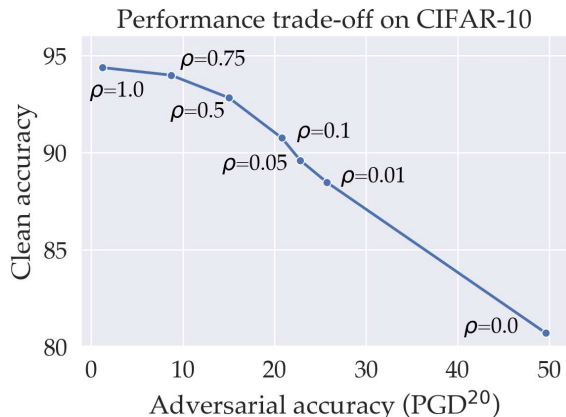


Figure 4. Trade-offs between adversarial and clean accuracy. By sweeping over ρ , we show that our approach bridges the average and worst case by trading-off clean and adversarial accuracy. Thus, as ρ decreases, trained classifiers improve robustness to adversarial perturbations at the cost of decreasing clean performance.

We also highlight the fact that PRL consistently outperforms both TERM and (Rice et al., 2021) on the clean, augmented, and quantile accuracy metrics.⁵ This demonstrates that PRL facilitates a strong empirical trade-off between robustness and accuracy relative to other methods that seek to interpolate between the average and worst case.

CVaR as a metric for test-time robustness. As we showed in Section 5, $\text{CVaR}_{1-\rho}$ is an upper bound for the ρ -ess sup. In this way, CVaR can be used as a surrogate for assessing the test-time robustness of trained classifiers. To this end, in Figure 3 we plot $\text{CVaR}_{0.95}(\ell(h(x + \delta), y); \tau)$ averaged over the test data on CIFAR-10, SVHN, and MNIST. This plot shows that PRL displays significantly lower values of $\text{CVaR}_{0.95}$ among all of the algorithms we considered, which reinforces the message from Tables 1-3 that PRL is most successful at imposing probabilistic robustness.

Ablation study: the role of ρ in Algorithm 1. In Section 5, we claimed that our algorithm interpolates between the average- and worst-case problems in (P-AVG) and (P-ROB) respectively. To verify this claim, we study the trade-off between nominal accuracy and adversarial accuracy for varying values of ρ . In Figure 4, we show that as ρ decreases, our algorithm improves adversarial accuracy at the cost of degrading nominal performance.

7. Conclusion

In this paper, motivated by the brittleness of ERM and the conservatism of adversarial training, we proposed a new framework called probabilistically robust learning in which robustness is enforced with high probability over pertur-

bations rather than in the worst case. Our analysis of this framework showed that PRL provably mitigates the well-known trade-off between robustness and accuracy and can have a sample complexity that is exponentially lower than that of adversarial training. We also proposed an algorithm motivated by risk-aware optimization which shows strong performance on a variety of metrics designed to evaluate intermediate robustness between the average and worst case.

There are numerous directions for future work. One fruitful direction is to further explore the use of risk-aware optimization in other areas of learning. Although these techniques have been applied to problems in reinforcement learning (Chow et al., 2015) and active learning (Curi et al., 2019), there is ample opportunity to use these methods to improve performance in fields like domain generalization, domain adaptation, and fair learning. Another exciting direction is to devise tighter and/or more efficient schemes to optimize the ρ -ess sup in (P-PRL). Furthermore, although we focused on additive, norm-bounded perturbations in this paper, there is ample opportunity to extend this framework to deal with more general distribution shifts, such as those studied in (Robey et al., 2020; 2021b; Wong & Kolter, 2020; Goyal et al., 2020; Zhou et al., 2022).

A final direction of interest is to study the connections between alternative definitions of robustness, including both (Li et al., 2020; 2021) and (Rice et al., 2021) as well as existing notions of astuteness (Wang et al., 2018) and of neighborhood-preserving Bayes optimal classifiers (Bhattacharjee & Chaudhuri, 2021) for non-parametric methods. Indeed, a unification of these robustness frameworks may represent a significant advance in our understanding of the robustness of machine learning models.

Acknowledgements

The research of Hamed Hassani and Alexander Robey is supported by NSF grants 1837253, 1943064, AFOSR grant FA9550-20-1-0111, DCIST-CRA, and the AI Institute for Learning-Enabled Optimization at Scale (TILOS). George J. Pappas and Alexander Robey are also supported by NSF grant 2038873 as well as the NSF-Simons Foundation’s Mathematical and Scientific Foundations of Deep Learning (MoDL) program on Transferable, Hierarchical, Expressive, Optimal, Robust, Interpretable NETWORKS (THEORINET).

⁵We selected t and q for TERM and N-HMC by cross-validation; see Appendix C for details.

References

- Anderson, T. W. *An Introduction to Multivariate Statistical Analysis*. Wiley New York, 2003. 6, 13
- Athalye, A., Carlini, N., and Wagner, D. Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. In *International conference on machine learning*, pp. 274–283. PMLR, 2018. 3
- Ben-Tal, A., El Ghaoui, L., and Nemirovski, A. *Robust optimization*. Princeton university press, 2009. 4
- Bhagoji, A. N., Cullina, D., and Mittal, P. Lower bounds on adversarial robustness from optimal transport. *arXiv preprint arXiv:1909.12272*, 2019. 5
- Bhattacharjee, R. and Chaudhuri, K. Consistent non-parametric methods for maximizing robustness. *Advances in Neural Information Processing Systems*, 34, 2021. 9
- Biggio, B., Corona, I., Maiorca, D., Nelson, B., Šrđić, N., Laskov, P., Giacinto, G., and Roli, F. Evasion attacks against machine learning at test time. In *Joint European conference on machine learning and knowledge discovery in databases*, pp. 387–402. Springer, 2013. 1, 3
- Bungert, L., Trillos, N. G., and Murray, R. The geometry of adversarial training in binary classification. *arXiv preprint arXiv:2111.13613*, 2021. 4
- Campi, M. C. and Garatti, S. The exact feasibility of randomized solutions of uncertain convex programs. *SIAM Journal on Optimization*, 19(3):1211–1230, 2008. 1, 4
- Charnes, A., Cooper, W. W., and Symonds, G. H. Cost horizons and certainty equivalents: an approach to stochastic programming of heating oil. *Management science*, 4(3): 235–263, 1958. 4
- Chen, S., Dobriban, E., and Lee, J. H. A group-theoretic framework for data augmentation. *Journal of Machine Learning Research*, 21(245):1–71, 2020. 3
- Chow, Y., Tamar, A., Mannor, S., and Pavone, M. Risk-sensitive and robust decision-making: a cvar optimization approach. *arXiv preprint arXiv:1506.02188*, 2015. 9
- Croce, F., Andriushchenko, M., Sehwag, V., DeBenedetti, E., Flammarion, N., Chiang, M., Mittal, P., and Hein, M. Robustbench: a standardized adversarial robustness benchmark. *arXiv preprint arXiv:2010.09670*, 2020. 1, 3
- Cullina, D., Bhagoji, A. N., and Mittal, P. Pac-learning in the presence of evasion adversaries. *arXiv preprint arXiv:1806.01471*, 2018. 1, 3, 5, 6
- Curi, S., Levy, K., Jegelka, S., Krause, A., et al. Adaptive sampling for stochastic risk-averse learning. *arXiv preprint arXiv:1910.12511*, 2019. 9
- Diochnos, D. I., Mahloujifar, S., and Mahmoody, M. Lower bounds for adversarially robust pac learning. *arXiv preprint arXiv:1906.05815*, 2019. 3
- Dobriban, E., Hassani, H., Hong, D., and Robey, A. Provable tradeoffs in adversarially robust classification. *arXiv preprint arXiv:2006.05161*, 2020. 3, 5, 6, 13
- Gilmer, J., Metz, L., Faghri, F., Schoenholz, S. S., Raghu, M., Wattenberg, M., and Goodfellow, I. Adversarial spheres. *arXiv preprint arXiv:1801.02774*, 2018. 4
- Goodfellow, I. J., Shlens, J., and Szegedy, C. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014. 1, 3, 8
- Gowal, S., Qin, C., Huang, P.-S., Cemgil, T., Dvijotham, K., Mann, T., and Kohli, P. Achieving robustness in the wild via adversarial mixing with disentangled representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1211–1220, 2020. 9
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016. 20
- Hendrycks, D., Mu, N., Cubuk, E. D., Zoph, B., Gilmer, J., and Lakshminarayanan, B. Augmix: A simple data processing method to improve robustness and uncertainty. *arXiv preprint arXiv:1912.02781*, 2019. 3
- Holmstrom, L. and Koistinen, P. Using additive noise in back-propagation training. *IEEE transactions on neural networks*, 3(1):24–38, 1992. 3
- Jalal, A., Ilyas, A., Daskalakis, C., and Dimakis, A. G. The robust manifold defense: Adversarial training using generative models. *arXiv preprint arXiv:1712.09196*, 2017. 3
- Javanmard, A., Soltanolkotabi, M., and Hassani, H. Precise tradeoffs in adversarial training for linear regression. In *Conference on Learning Theory*, pp. 2034–2078. PMLR, 2020. 3, 5, 15, 16
- Kamalaruban, P., Huang, Y.-T., Hsieh, Y.-P., Rolland, P., Shi, C., and Cevher, V. Robust reinforcement learning via adversarial training with langevin dynamics. *arXiv preprint arXiv:2002.06063*, 2020. 3
- Kannan, H., Kurakin, A., and Goodfellow, I. Adversarial logit pairing. *arXiv preprint arXiv:1803.06373*, 2018. 3

- Khoury, M. and Hadfield-Menell, D. On the geometry of adversarial examples. *arXiv preprint arXiv:1811.00525*, 2018. 4
- Krizhevsky, A., Hinton, G., et al. Learning multiple layers of features from tiny images. 2009. 20
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25: 1097–1105, 2012. 3
- Krokhmal, P., Palmquist, J., and Uryasev, S. Portfolio optimization with conditional value-at-risk objective and constraints. *Journal of risk*, 4:43–68, 2002. 7
- Laidlaw, C. and Feizi, S. Functional adversarial attacks. *arXiv preprint arXiv:1906.00001*, 2019. 3
- Li, T., Beirami, A., Sanjabi, M., and Smith, V. Tilted empirical risk minimization. *arXiv preprint arXiv:2007.01162*, 2020. 3, 7, 9
- Li, T., Beirami, A., Sanjabi, M., and Smith, V. On tilted losses in machine learning: Theory and applications. *arXiv preprint arXiv:2109.06141*, 2021. 3, 7, 9
- Madry, A., Makelov, A., Schmidt, L., Tsipras, D., and Vladu, A. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017. 1, 3, 8
- Miller, B. L. and Wagner, H. M. Chance constrained programming with joint constraints. *Operations Research*, 13(6):930–945, 1965. 4
- MNISTWebPage. The MNIST database of handwritten digits Home Page. <http://yann.lecun.com/exdb/mnist/>. 20
- Montasser, O., Hanneke, S., and Srebro, N. Vc classes are adversarially robustly learnable, but only improperly. In *Conference on Learning Theory*, pp. 2512–2530. PMLR, 2019. 1, 3, 6, 17
- Montasser, O., Goel, S., Diakonikolas, I., and Srebro, N. Efficiently learning adversarially robust halfspaces with noise. In *International Conference on Machine Learning*, pp. 7010–7021. PMLR, 2020. 5
- Moosavi-Dezfooli, S.-M., Fawzi, A., and Frossard, P. Deepfool: a simple and accurate method to fool deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2574–2582, 2016. 3
- Nemirovski, A. and Shapiro, A. Convex approximations of chance constrained programs. *SIAM Journal on Optimization*, 17(4):969–996, 2007. 7
- Netzer, Y., Wang, T., Coates, A., Bissacco, A., Wu, B., and Ng, A. Y. Reading digits in natural images with unsupervised feature learning. 2011. 20
- Raghunathan, A., Xie, S. M., Yang, F., Duchi, J. C., and Liang, P. Adversarial training can hurt generalization. *arXiv preprint arXiv:1906.06032*, 2019. 1
- Ramponi, F. A. Consistency of the scenario approach. *SIAM Journal on Optimization*, 28(1):135–162, 2018. 4
- Rebuffi, S.-A., Gowal, S., Calian, D. A., Stimberg, F., Wiles, O., and Mann, T. Fixing data augmentation to improve adversarial robustness. *arXiv preprint arXiv:2103.01946*, 2021. 3
- Rice, L., Bair, A., Zhang, H., and Kolter, J. Z. Robustness between the worst and average case. *Advances in Neural Information Processing Systems*, 34, 2021. 3, 8, 9, 20
- Robey, A., Hassani, H., and Pappas, G. J. Model-based robust deep learning: Generalizing to natural, out-of-distribution data. *arXiv preprint arXiv:2005.10247*, 2020. 3, 9
- Robey, A., Chamon, L., Pappas, G., Hassani, H., and Ribeiro, A. Adversarial robustness with semi-infinite constrained learning. *Advances in Neural Information Processing Systems*, 34, 2021a. 8
- Robey, A., Pappas, G. J., and Hassani, H. Model-based domain generalization. *Advances in Neural Information Processing Systems*, 34:20210–20229, 2021b. 9
- Rockafellar, R. T. and Uryasev, S. Conditional value-at-risk for general loss distributions. *Journal of banking & finance*, 26(7):1443–1471, 2002. 7
- Rockafellar, R. T., Uryasev, S., et al. Optimization of conditional value-at-risk. *Journal of risk*, 2:21–42, 2000. 7
- Schildbach, G., Fagiano, L., Frei, C., and Morari, M. The scenario approach for stochastic model predictive control with bounds on closed-loop constraint violations. *Automatica*, 50(12):3009–3018, 2014. 4
- Shalev-Shwartz, S. and Ben-David, S. *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014. 17
- Shamir, A., Melamed, O., and BenShmuel, O. The dimpled manifold model of adversarial examples in machine learning. *arXiv preprint arXiv:2106.10151*, 2021. 4
- Shen, Y., Zheng, L., Shu, M., Li, W., Goldstein, T., and Lin, M. C. Improving robustness of learning-based autonomous steering using adversarial images. *arXiv preprint arXiv:2102.13262*, 2021. 1

- Soltanolkotabi, M., Javanmard, A., and Lee, J. D. Theoretical insights into the optimization landscape of over-parameterized shallow neural networks. *IEEE Transactions on Information Theory*, 65(2):742–769, 2018. 3
- Spielman, D. A. and Teng, S.-H. Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time. *Journal of the ACM (JACM)*, 51(3): 385–463, 2004. 1, 4
- Su, D., Zhang, H., Chen, H., Yi, J., Chen, P.-Y., and Gao, Y. Is robustness the cost of accuracy?—a comprehensive study on the robustness of 18 deep image classification models. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 631–648, 2018. 1, 3, 5
- Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., and Fergus, R. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013. 3
- Tang, S., Gong, R., Wang, Y., Liu, A., Wang, J., Chen, X., Yu, F., Liu, X., Song, D., Yuille, A., et al. Robust-tart: Benchmarking robustness on architecture design and training techniques. *arXiv preprint arXiv:2109.05211*, 2021. 1
- Thomas, P. and Learned-Miller, E. Concentration inequalities for conditional value at risk. In *International Conference on Machine Learning*, pp. 6225–6233. PMLR, 2019. 7
- Tsipras, D., Santurkar, S., Engstrom, L., Turner, A., and Madry, A. Robustness may be at odds with accuracy. *arXiv preprint arXiv:1805.12152*, 2018. 1, 3, 6
- Vapnik, V. *The nature of statistical learning theory*. Springer science & business media, 1999. 1, 2, 7
- Wang, Y., Jha, S., and Chaudhuri, K. Analyzing the robustness of nearest neighbors to adversarial examples. In *International Conference on Machine Learning*, pp. 5133–5142. PMLR, 2018. 9
- Wang, Y., Zou, D., Yi, J., Bailey, J., Ma, X., and Gu, Q. Improving adversarial robustness requires revisiting misclassified examples. In *International Conference on Learning Representations*, 2019. 3, 8, 20
- Wong, E. and Kolter, J. Z. Learning perturbation sets for robust machine learning. *arXiv preprint arXiv:2007.08450*, 2020. 3, 9
- Wong, E. and Kolter, Z. Provable defenses against adversarial examples via the convex outer adversarial polytope. In *International Conference on Machine Learning*, pp. 5286–5295. PMLR, 2018. 3
- Yang, Y.-Y., Rashtchian, C., Zhang, H., Salakhutdinov, R., and Chaudhuri, K. A closer look at accuracy vs. robustness. *arXiv preprint arXiv:2003.02460*, 2020. 3
- Yin, D., Kannan, R., and Bartlett, P. Rademacher complexity for adversarially robust generalization. In *International conference on machine learning*, pp. 7085–7094. PMLR, 2019. 6
- Zeiler, M. D. Adadelata: an adaptive learning rate method. *arXiv preprint arXiv:1212.5701*, 2012. 20
- Zhang, H., Yu, Y., Jiao, J., Xing, E., El Ghaoui, L., and Jordan, M. Theoretically principled trade-off between robustness and accuracy. In *International Conference on Machine Learning*, pp. 7472–7482. PMLR, 2019. 1, 3, 8, 20
- Zhou, A., Tajwar, F., Robey, A., Knowles, T., Pappas, G. J., Hassani, H., and Finn, C. Do deep networks transfer invariances across classes? *arXiv preprint arXiv:2203.09739*, 2022. 9
- Zhou, K. and Doyle, J. C. *Essentials of robust control*, volume 104. Prentice hall Upper Saddle River, NJ, 1998. 1

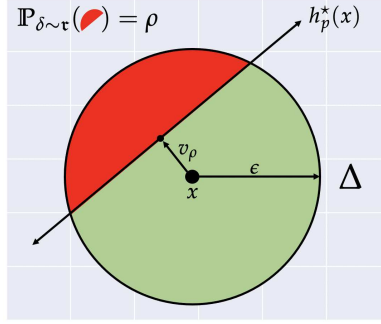


Figure 5. Spherical cap of an ℓ_2 -ball with radius $\epsilon > 0$ in two-dimensions.

A. Proofs concerning trade-offs in binary classification and linear regression

A.1. Binary classification under a Gaussian mixture model

In this subsection, we provide proofs for the results for binary classification in Section 4.1. In general, our proof of the closed form expression for h_p^* in Proposition 4.1 follows along the same lines as the proof of Theorem 4.2 in (Dobriban et al., 2020). In particular, our contribution is to generalize the proof techniques to the setting of probabilistic robustness, thereby subsuming the results in (Dobriban et al., 2020) as a special case when $\rho = 0$.

Lemma A.1. For any $\rho \in [0, 1/2]$, it holds that among all linear classifiers,

$$h^*(x) = \text{sign} \left(x^\top \mu \left(1 - \frac{v_\rho}{\|\mu\|_2} \right)_+ - \frac{q}{2} \right) \quad (\text{A.1})$$

is optimal for (P-PRL), where v_ρ is the distance from the center of Δ to a spherical cap of volume ρ .

Proof. To begin, observe that the the probabilistically robust risk $\text{PR}(h_p; \rho)$ can be written in the following way:

$$\text{PR}(h_p; \rho) = \mathbb{E}_{(x,y) \sim \mathcal{D}} \left[\mathbb{I} \left[\mathbb{P}_{\delta \sim \tau} [h(x + \delta) \neq y] > \rho \right] \right] \quad (\text{A.2})$$

$$= \mathbb{P}_{(x,y) \sim \mathcal{D}} \left[\mathbb{P}_{\delta \sim \tau} [h(x + \delta) \neq y] > \rho \right] \quad (\text{A.3})$$

$$= \mathbb{P}[y = +1] \cdot \mathbb{P}_{x|y=+1} \left[\mathbb{P}_{\delta \sim \tau} [h(x + \delta) = -1] > \rho \right] \quad (\text{A.4})$$

$$+ \mathbb{P}[y = -1] \cdot \mathbb{P}_{x|y=-1} \left[\mathbb{P}_{\delta \sim \tau} [h(x + \delta) = +1] > \rho \right]. \quad (\text{A.5})$$

Note that it is enough to solve this problem in one dimension, as the problem in d -dimensions can be easily reduced to a one-dimensional problem. Thus, our goal is to find the value of a threshold c that minimizes the probabilistically robust risk. In this one-dimensional case, the probabilistically robust risk can be written as

$$\text{PR}(h_p; \rho) = \pi \cdot \mathbb{P}_{x|y=+1} [x \leq c + \rho] + (1 - \pi) \cdot \mathbb{P}_{x|y=-1} [x \geq c - \rho] \quad (\text{A.6})$$

$$= \pi \cdot \mathbb{P}_{x|y=+1} [x - \rho \leq c] + (1 - \pi) \cdot \mathbb{P}_{x|y=-1} [x + \rho \geq c]. \quad (\text{A.7})$$

Recall that as $x|y \sim \mathcal{N}(y\mu, \sigma^2 I)$. Therefore,

$$\text{PR}(h_p; \rho) = \pi \cdot \mathbb{P}_{x \sim \mathcal{N}(\mu - \rho, \sigma^2 I)} [x \leq c] + (1 - \pi) \cdot \mathbb{P}_{x \sim \mathcal{N}(-\mu + \rho, \sigma^2 I)} [x \geq c] \quad (\text{A.8})$$

This is exactly the same as the problem of non-robust classification between tow Gaussians with means $\mu' = \mu - \rho$ and $-\mu'$ (by assumption, we have that $\mu \geq 0$). As is well known (see, e.g., (Anderson, 2003)), the optimal classifier for this setting is

$$h_p^*(x) = \text{sign}[x \cdot (\mu - \rho) - q/2] \quad (\text{A.9})$$

where $q = \ln[(1 - \pi)/\pi]$. And indeed, when moving from the one-dimensional case, one need only recognize that a linear classifier which ignores a set of volume ρ in Δ will create a spherical cap of volume ρ in Δ . This is illustrated by the red region in Figure 5. The form of $h_p^*(x)$ given in (4.5) follows from (A.9) as a direct analog for the d -dimensional case. \square

To prove the second part of Proposition 4.1, we seek to characterize the distance v_ρ from the center of Δ to a spherical cap of volume ρ . To this end, we have the following result.

Lemma A.2. *Let $B(0, \gamma) = \{\delta \in \Delta : \|\delta\|_2 \leq \gamma\}$ for any number $\gamma > 0$. Define v_ρ to be the distance from the origin such that the fraction of the volume of the corresponding spherical cap is ρ (see Figure 5). Then we have that*

$$v_\rho = \begin{cases} \epsilon & \rho = 0 \\ \frac{\epsilon}{\sqrt{d}} \Phi^{-1}(1 - \rho)(1 - o_d(1)) & \rho \in (0, 1/2] \end{cases} \quad (\text{A.10})$$

Proof. By inspection, the result is clear for $\rho = 0$. Thus, we consider the case when $\rho \in [0, 1/2]$. Note that for any number $\gamma < \epsilon$, we have that

$$\frac{\mathbb{P}_{\delta \sim \tau}(B(0, \epsilon)) - \mathbb{P}_{\delta \sim \tau}(B(0, \epsilon - \gamma))}{\mathbb{P}_{\delta \sim \tau}(B(0, \epsilon))} = 1 - \left(1 - \frac{\gamma}{\epsilon}\right)^d. \quad (\text{A.11})$$

As a result, by taking $\gamma = \epsilon \cdot (\ln(d)/d)$ we obtain

$$1 - \left(1 - \frac{\gamma}{\epsilon}\right)^d = 1 - \exp\left(d \ln\left(1 - \frac{\gamma}{\epsilon}\right)\right) \quad (\text{A.12})$$

$$= 1 - \exp\left(-\frac{d\gamma}{\epsilon} + O\left(d\left(\frac{\gamma}{\epsilon}\right)^2\right)\right) \quad (\text{A.13})$$

$$= 1 - O(1/d). \quad (\text{A.14})$$

In this way, we have shown that the uniform distribution over any ball centered at the origin can be approximated up to $o_d(1)$ by the uniform distribution on the sphere.

Now let (X_1, \dots, X_n) be the random vector generated by uniformly sampling a point on the sphere of radius ϵ . We note that, up to $o_d(1)$ terms, the distribution of each of the coordinates, e.g. X_1 , is $\epsilon Z/\sqrt{d}$, where Z is the normal random variable. Again, up to $o_d(1)$ terms, i.e. when the dimension grows large, the volume of the spherical cap at distance v_ρ can be approximated by

$$\mathbb{P}(X_1 \geq v_\rho) = \mathbb{P}\left(Z \frac{\epsilon}{\sqrt{d}} \geq v_\rho\right) = 1 - \Phi\left(v_\rho \frac{\sqrt{d}}{\epsilon}\right). \quad (\text{A.15})$$

where Φ denotes the Gaussian CDF. As a result, for the RHS of the above relation to be equal to ρ , we must have

$$v_\rho = \frac{\epsilon \Phi^{-1}(1 - \rho)}{\sqrt{d}}. \quad (\text{A.16})$$

This concludes the proof. \square

From the above lemma, we can conclude the following phase-transition behavior. When $\rho = 0$, there is a constant gap between the adversarially robust risk $\text{AR}(h_r^*)$ and the best attainable clean risk $\text{SR}(h_{\text{Bayes}}^*)$. Indeed, this gap does not vanish as the dimension d grows, resulting in a non-trivial trade-off between adversarial robustness and accuracy. However, for $\rho > 0$, the gap between the probabilistically robust accuracy $\text{PR}(h_p^*; \rho)$ and the clean risk is of the form $\text{PR}(h_p^*; \rho) - \text{SR}(h_{\text{Bayes}}^*) = O(1/\sqrt{d})$, and as a result, as the dimension d grows, the trade-off between robustness and accuracy vanishes (a blessing of high dimensions).

A.2. Linear regression with Gaussian features

We next consider the setting of linear regression, wherein it is assumed that there exists an underlying parameter vector $\theta_0 \in \Theta \subset \mathbb{R}^d$, and that the data is subsequently generated according to the following model:

$$x \sim \mathcal{N}(0, I_d), \quad y = \theta_0^\top x + z, \quad z \sim \mathcal{N}(0, \sigma^2) \quad (\text{A.17})$$

where $\sigma > 0$ is a fixed noise level. Furthermore, we consider hypotheses of the form $f_\theta(x) = \theta^\top x$ for $\theta \in \Theta$, and we use the squared loss $\ell(f_\theta(x), y) = (f_\theta(x) - y)^2 = (\theta^\top x - y)^2$. In this setting, it is straightforward to calculate that at optimality $\text{SR}(f_\theta) = \sigma^2$, which is achieved for $\theta = \theta_0$. Moreover, in the more general probabilistic robustness setting, we have the following complementary result:

Proposition A.3. Suppose that the data is distributed according to (A.17). Let $\theta^* \in \Theta$ denote the optimal solution obtained by solving (P-PRL) over Θ . Then for any $\rho > 0$,

$$\text{PR}(f_{\theta^*}; \rho) - \text{SR}(f_{\theta_0}) = \begin{cases} O(1/\sqrt{d}) & \rho > 0 \\ O(1) & \rho = 0 \end{cases} \quad (\text{A.18})$$

In this way, as in the previous subsection, it holds that for any $\rho > 0$, the gap between probabilistic robustness and nominal performance vanishes in high dimensions. On the other hand, as was recently shown in (Javanmard et al., 2020), there exists a non-trivial gap between adversarial robustness and clean accuracy that does not vanish to zero by increasing the dimension in this setting.

To prove this result, we consider the following variational form of the problem in (P-PRL):

$$\begin{aligned} \min_{\theta \in \mathbb{R}^d, t \in L^1} \quad & \mathbb{E}_{(x,y) \sim \mathfrak{D}} [t(x, y)] \\ \text{subject to} \quad & \mathbb{P}_{\delta \sim \tau} \{ (\theta^\top (x + \delta) - y)^2 \leq t(x, y) \} \geq 1 - \rho \quad \forall (x, y) \in \Omega. \end{aligned} \quad (\text{PV})$$

where L^1 denotes the space of Lebesgue integral functions. We can then characterize $t(x, y)$ as follows:

Lemma A.4. We have the following characterization for $t(x, y)$:

$$t(x, y) = \begin{cases} (|\theta^\top x - y| + \epsilon \|\theta\|_2)^2 & \text{if } \rho = 0, \\ (\theta^\top x - y)^2 + \frac{\epsilon^2 \|\theta\|_2^2 (\theta^\top x - y)}{\sqrt{d}} (Q^{-1}(1 - \rho) + o_d(1)) & \text{if } \rho \in (0, 1], \end{cases}$$

Almost surely for any (x, y) .

Proof. Let's first consider the case in which $\rho > 0$. Since $\delta \sim \mathbb{P}_{\delta \sim \tau}$ is the uniform distribution over the Euclidean ball of radius ϵ , we know that for any $\theta \in \mathbb{R}^d$ we have

$$\theta^\top \delta \xrightarrow{d} \mathcal{N}\left(0, \frac{\epsilon^2 \|\theta\|_2^2}{d}\right),$$

where the convergence is in distribution. This is because the uniform distribution over the Euclidean ball of radius ϵ converges to the Gaussian distribution $N(0, \frac{\epsilon^2}{d} I_d)$. As a result, up to $o_d(1)$ terms, we have $\theta^\top \delta \sim \frac{\epsilon \|\theta\|_2}{\sqrt{d}} Z$, where Z is the normal random variable.

$$\begin{aligned} & \mathbb{P}_{\delta \sim \tau} \{ (\theta^\top (x + \delta) - y)^2 \leq t(x, y) \} \\ &= \mathbb{P}_Z \left\{ \frac{\epsilon^2 \|\theta\|_2^2}{d} Z^2 + 2(\theta^\top x - y) \frac{\epsilon \|\theta\|_2}{\sqrt{d}} Z + (\theta^\top x - y)^2 \leq t(x, y) \right\} + o_d(1) \\ &= \mathbb{P}_Z \left\{ 2(\theta^\top x - y) \frac{\epsilon \|\theta\|_2}{\sqrt{d}} Z + (\theta^\top x - y)^2 \leq t(x, y) \right\} + o_d(1) \\ &= \mathbb{P}_Z \left\{ \frac{\epsilon \|\theta\|_2}{\sqrt{d}} Z + (\theta^\top x - y)^2 \leq t(x, y) \right\} + o_d(1) \\ &= \mathbb{P}_Z \left\{ Z \leq \sqrt{d} \frac{t(x, y) - (\theta^\top x - y)^2}{2\epsilon \|\theta\|_2 (\theta^\top x - y)} \right\} + o_d(1) \\ &= Q \left(\sqrt{d} \frac{t(x, y) - (\theta^\top x - y)^2}{2\epsilon \|\theta\|_2 (\theta^\top x - y)} \right) + o_d(1) \end{aligned}$$

where Q is the quantile function for Z , i.e. the inverse of the normal CDF. As a result, equating the above with $1 - \rho$ we obtain

$$t(x, y) = (\theta^\top x - y)^2 + \frac{2\epsilon \|\theta\|_2 (\theta^\top x - y)}{\sqrt{d}} (Q^{-1}(1 - \rho) + o_d(1)).$$

For the case $\beta = 0$, this is a simple optimization problem where the closed form solution provided in the lemma is its solution. \square

From the above lemma, it is easy to see that for $\theta = \theta_0$, the value of the objective in (??) becomes $\sigma^2 + O(1/\sqrt{d})$ for any value of $\rho > 0$. This means that the gap between the probabilistically robust risk and the clean risk is of the form $\text{PR}(h_p^*; \rho) - \text{SR}(h_r) = O(1/\sqrt{d})$, and as a result, as the dimension d grows, the trade-off between probabilistic robustness and accuracy vanishes. We further remark that for $\rho = 0$, i.e. the adversarial setting, there exists a non-trivial gap between the robust and clean accuracies that does not vanish to zero by increasing dimension. This is indeed clear from the above lemma, and it has also been shown in (Javanmard et al., 2020).

B. Learning theory proofs

We restate Proposition 4.2 formally to detail what we mean by *sample complexity*. Note that this is exactly the number of sample required for PAC learning.

Proposition B.1. *Consider the probabilistically robust learning problem (P-PRL) with the 0-1 loss function and a robust measure τ fully supported on Δ and absolutely continuous with respect to the Lebesgue measure (i.e., non-atomic). Let P_r^* be its optimal value and consider its empirical version*

$$\hat{P}_r^* = \min_{h_p \in \mathcal{H}} \frac{1}{N} \sum_{n=1}^N \left[\rho\text{-ess sup}_{\delta \sim \tau} \ell(h_p(x_n + \delta), y_n) \right],$$

based on i.i.d. samples $(x_n, y_n) \sim \mathcal{D}$. For any threshold $\rho_o \in (0, 0.5)$, there exists a hypothesis class \mathcal{H}_o such that the sample complexity of probabilistically robust learning, i.e., the number of samples N needed for $|P_r^* - \hat{P}_r^*| \leq \epsilon$ with high probability, is

$$N = \begin{cases} \Theta(\log_2(1/\rho_o)/\epsilon^2), & \rho = 0 \\ \Theta(1/\epsilon^2), & \rho_o \leq \rho \leq 1 - \rho_o \end{cases}$$

In particular, $\Theta(1/\epsilon^2)$ is the sample complexity of (P-NOM).

Proof. Let us begin by reducing the task of determining the sample complexity of these problems to that of determining the VC dimension of their objectives:

Lemma B.2 ((Shalev-Shwartz & Ben-David, 2014, Thm. 6.8)). *Consider*

$$\begin{aligned} P^* &= \min_{h \in \mathcal{H}} \mathbb{E}_{(x,y) \sim \mathcal{D}} [g(h, x, y)] \\ \hat{P}^* &= \min_{h \in \mathcal{H}} \frac{1}{N} \sum_{n=1}^N g(h, x_n, y_n) \end{aligned} \tag{B.1}$$

where $g : \mathcal{H} \times \mathcal{X} \times \mathcal{Y} \rightarrow \{0, 1\}$ and the samples $\{(x_n, y_n)\} \sim \mathcal{D}$ are i.i.d. The number of samples N needed for $|P^* - \hat{P}^*| \leq \epsilon$ with probability $1 - \delta$ over the sample set $\{(x_n, y_n)\}$ is

$$C_1 \frac{d_{VC} + \log(1/\delta)}{\epsilon^2} \leq N \leq C_2 \frac{d_{VC} + \log(1/\delta)}{\epsilon^2},$$

for universal constants C_1, C_2 . The VC dimension d_{VC} is defined the largest d such that $\Pi(d) = 2^d$ for the growth function

$$\Pi(d) = \max_{\{(x_n, y_n)\} \subset (\mathcal{X} \times \mathcal{Y})^d} |\mathcal{S}(\{(x_n, y_n)\})|,$$

where $\mathcal{S}(\{(x_n, y_n)\}) = \{u \in \{0, 1\}^m \mid \exists h \in \mathcal{H} \text{ such that } u_n = \ell(h(x_n), y_n)\}$.

We now proceed by defining \mathcal{H}_o using a modified version of the construction in (Montasser et al., 2019, Lemma 2). Let $m = \lceil \log_2(1/\rho_o) \rceil + 1$ and pick $\{c_1, \dots, c_m\} \in \mathcal{X}^m$ such that, for $\Delta_i = c_i + \Delta$, it holds that $\Delta_i \cap \Delta_j = \emptyset$ for $i \neq j$. Within each Δ_i , define 2^{m-1} disjoint sets \mathcal{A}_i^b of measure $\tau(\mathcal{A}_i^b) \leq \rho_o/m$ labeled by the binary digits $b \in \{0, 1\}^m$ whose i -th digit is one. In other words, Δ_1 contains sets with signature $1b_2 \dots b_m$ and Δ_3 contains sets with signature $b_1 b_2 1 \dots b_m$. Observe that there are indeed 2^{m-1} sets \mathcal{A}_i^b within each Δ_i , that their signatures span all possible m -digits binary numbers, and there are at most m sets with the same signature (explicitly, for $b = 11 \dots 1$). Additionally, note that it is indeed possible to fit the \mathcal{A}_i^b inside each Δ_i given that

$$2^{m-1} \cdot \frac{\rho_o}{m} < \frac{2^{\log_2(1/\rho_o)+1} \rho_o}{\log_2(1/\rho_o) + 1} \leq \frac{2}{\log_2(1/\rho_o) + 1} \leq 1$$

for $\rho_o < 0.5$, where we used the fact that $\log_2(1/\rho_o) \leq \lceil \log_2(1/\rho_o) \rceil < \log_2(1/\rho_o) + 1$.

We can now construct the hypothesis class $\mathcal{H}_o = \{h_b \mid b \in \{0, 1\}^m\}$ by taking

$$h_b(x) = \begin{cases} 1, & x \notin \bigcup_{i=1}^m \mathcal{A}_i^b \\ 0, & x \in \bigcup_{i=1}^m \mathcal{A}_i^b \end{cases} \tag{B.2}$$

Let us proceed first for the probabilistically robust loss $g(h, x, y) = \rho\text{-ess sup}_{\delta \sim \tau} \mathbb{I}[h(x + \delta) \neq y]$.

We begin by showing that if $\rho = 0$, then $d_{\text{VC}} > m$. Indeed, consider the set of points $\{(c_i, 1)\}_{i=1, \dots, m} \subset (\mathcal{X} \times \mathcal{Y})^m$. In this case, the cardinality of $\mathcal{S}(\{(c_i, 1)\})$ is 2^m , i.e., this set can be shattered by \mathcal{H} . Indeed, for any signature $b \in \{0, 1\}^m$, we have

$$\rho\text{-ess sup}_{\delta \sim \tau} \mathbb{I}[h_b(c_i + \delta) \neq 1] = \text{ess sup}_{\delta \sim \tau} \mathbb{I}[h_b(c_i + \delta) \neq 1] = b_i, \quad \text{for } i = 1, \dots, m,$$

since $h_b(c_i + \delta) = 0$ for all $c_i + \delta \in \mathcal{A}_i^b$, a set of positive measure. Using Lemma B.2, we therefore conclude that $N = \Theta(m/\epsilon^2)$.

Let us now show that $d_{\text{VC}} = 1$ for $\rho \geq \rho_o$ by showing that $\Pi(2) < 4$. To do so, we take two arbitrary points $(x_1, y_1), (x_2, y_2) \in \mathcal{X} \times \mathcal{Y}$ and proceed case-by-case. To simplify the exposition, let $\mathcal{A} = \bigcup_{i=1}^m \bigcup_{b \in \{0,1\}^m} \mathcal{A}_i^b$.

- Suppose that $(x_1 + \Delta) \cap \mathcal{A} = \emptyset$. Then, observe from (B.2) that $h(x_1 + \delta) = 1$ for all $h \in \mathcal{H}$ and $\delta \in \Delta$. Hence, for all $h \in \mathcal{H}$, we obtain

$$\rho\text{-ess sup}_{\delta \sim \tau} \mathbb{I}[h(x_1 + \delta) \neq y_1] = \mathbb{I}[1 \neq y_1].$$

Hence, depending on the value of y_1 , $\mathcal{S}(\{(x_1, y_1), (x_2, y_2)\})$ can either contain sets of the form $(0, q)$ or $(1, q)$, for $q = \{0, 1\}$, but not both. For such points, we therefore have $|\mathcal{S}(\{(x_1, y_1), (x_2, y_2)\})| \leq 2 < 4$. The same argument holds for $(x_2 + \Delta) \cap \mathcal{A} = \emptyset$.

- Suppose then that both $(x_1 + \Delta) \cap \mathcal{A} \neq \emptyset$ and $(x_2 + \Delta) \cap \mathcal{A} \neq \emptyset$. Then, $(x_j + \Delta)$ can intersect at most m sets \mathcal{A}_i^b with the same signature b (explicitly, $b = 11 \dots 1$). But, by construction, $\tau(\mathcal{A}_i^b) \leq \rho_o/m$, which implies that $\tau(\bigcup_i \mathcal{A}_i^{11 \dots 1}) \leq \rho_o$. We then consider the possible labels separately:

- for $y_j = 1$, we know from (B.2) that $\mathbb{I}[h(x_j + \delta) \neq 1] = 1$ only when $x_j + \delta \in \mathcal{A}_i^b$. But since $\rho \geq \rho_o$, these sets can be ignored when computing the ρ -ess sup and we get that

$$\rho\text{-ess sup}_{\delta \sim \tau} \mathbb{I}[h(x_j + \delta) \neq 1] = 0, \quad \text{for all } h \in \mathcal{H};$$

- for $y_j = 0$, we obtain from (B.2) that $\mathbb{I}[h_b(x_j + \delta) \neq 0] = 1$ everywhere except when $x_j + \delta \in \mathcal{A}_i^b$. Recall that for any signature b , it holds that $\tau(\{\delta \mid x_j + \delta \in \bigcup_i \mathcal{A}_i^b\}) \leq \rho_o$ (possibly with equality if $b = 11 \dots 1$). Hence, since $\rho < 1 - \rho_o$, it holds that

$$\rho\text{-ess sup}_{\delta \sim \tau} \mathbb{I}[h(x_j + \delta) \neq 0] = 1, \quad \text{for all } h \in \mathcal{H}.$$

Since the ρ -ess sup does not vary over \mathcal{H} for either x_1 or x_2 , we conclude that $|\mathcal{S}(\{(x_1, y_1), (x_2, y_2)\})| \leq 2 < 4$, i.e., \mathcal{H} cannot shatter these points.

This implies $d_{\text{VC}} \leq 1$ and since $|\mathcal{H}| > 1$, we obtain $d_{\text{VC}} = 1$. Using Lemma B.2, we therefore conclude that $N = \Theta(1/\epsilon^2)$.

Finally, we consider the case of (P-NOM) for the nominal loss $g(h, x, y) = \mathbb{I}[h(x) \neq y]$. Once again, we take two points $(x_1, y_1), (x_2, y_2) \in \mathcal{X} \times \mathcal{Y}$ and proceed case-by-case.

- Suppose, without loss of generality, that $x_1 \notin \mathcal{A}$. Then, (B.2) yields $h(x_1) = 1$ for all $h \in \mathcal{H}$ and $\ell(h(x_1), y_1) = \mathbb{I}[1 \neq y_1]$. Depending on the value of y_1 , $\mathcal{S}(\{(x_1, y_1), (x_2, y_2)\})$ only has sets of the form $(0, q)$ or $(1, q)$, $q = \{0, 1\}$, but not both. For such points, $|\mathcal{S}(\{(x_1, y_1), (x_2, y_2)\})| \leq 2 < 4$. The same holds for x_2 .
- Suppose now that $x_1 \in \mathcal{A}_i^{b_1}$ and $x_2 \in \mathcal{A}_i^{b_2}$. Then, $h_b(x_j) = 0$ for $b = b_j$ or $h_b(x_j) = 1$ for $b \neq b_j$. Hence,
 - if $b_1 = b_2$, then $h_b(x_1) = h_b(x_2)$. Depending on the value of the labels, $\mathcal{S}(\{(x_1, y_1), (x_2, y_2)\})$ only has sets of the form (q, q) or $(1 - q, q)$, $q = \{0, 1\}$, but not both. For such points, we once again have $|\mathcal{S}(\{(x_1, y_1), (x_2, y_2)\})| \leq 2 < 4$;
 - if $b_1 \neq b_2$, then $h_b(x_1) = 1 - h_b(x_2)$. Depending on the value of the labels, $\mathcal{S}(\{(x_1, y_1), (x_2, y_2)\})$ has, once again, only sets of the form (q, q) or $(1 - q, q)$, $q = \{0, 1\}$, but not both. Hence, $|\mathcal{S}(\{(x_1, y_1), (x_2, y_2)\})| \leq 2 < 4$ for such points;

In none of these cases \mathcal{H} is able to shatter two points, meaning that $d_{VC} \leq 1$. Since $|\mathcal{H}| > 1$, it holds that $d_{VC} = 1$ which is indeed the same value as when $\rho \geq \rho_o$. \square

C. Hyperparameter selection and implementation details

In this appendix, we discuss hyperparameter selection and computational details. All experiments were run across two four-GPU workstations, comprising a total of eight Quadro RTX 5000 GPUs. Our code is available at: <https://github.com/arobey1/advbench>.

C.1. MNIST

For the MNIST dataset ([MNISTWebPage](#)), we used a four-layer CNN architecture with two convolutional layers and two feed-forward layers. To train these models, we use the Adadelta optimizer ([Zeiler, 2012](#)) to minimize the cross-entropy loss for 150 epochs with no learning rate decay and an initial learning rate of 1.0. All classifiers were evaluated with a 10-step PGD adversary. To compute the augmented accuracy, we sampled ten samples from τ per data point, and to compute the ProbAcc metric, we sample 100 perturbations per data point.

C.2. CIFAR-10 and SVHN

For CIFAR-10 ([Krizhevsky et al., 2009](#)) and SVHN ([Netzer et al., 2011](#)), we used the ResNet-18 architecture ([He et al., 2016](#)). We trained using SGD and an initial learning rate of 0.01 and a momentum of 0.9. We also used weight decay with a penalty weight of 3.5×10^{-3} . All classifiers were trained for 115 epochs, and we decayed the learning rate by a factor of 10 at epochs 55, 75, and 90.

C.3. Baseline algorithms

In the experiments section, we trained a number of baseline algorithms. In what follows, we list the hyperparameters we used for each of these algorithms:

- **PGD.** For MNIST, we ran seven steps of gradient ascent at training time with a step size of $\alpha = 0.1$. On CIFAR-10 and SVHN, we ran ten steps of gradient ascent at training time with a step size of $\alpha = 2/255$.
- **TRADES.** We used the same step sizes and number of steps as stated about for PGD. Following the literature ([Zhang et al., 2019](#)), we used a weight of $\beta = 6.0$ for all datasets.
- **MART.** We used the same step sizes and number of steps as stated about for PGD. Following the literature ([Wang et al., 2019](#)), we used a weight of $\lambda = 5.0$ for all datasets.
- **DALE.** We used the same step sizes and number of steps as stated about for PGD. For all datasets, we used a margin of $\rho = 0.1$ (note that this ρ is different from the ρ used in the definition of probabilistically robust learning). For MNIST, we used a dual step size of $\eta_p = 1.0$; for CIFAR-10 and SVHN, we used $\eta_p = 0.01$. For MNIST, we used a temperature of $\sqrt{2\eta T}$ of 10^{-3} ; for CIFAR-10 and SVHN, we used 10^{-5} .
- **TERM.** We chose the value of t in (PII) by cross-validation on the set $\{0.1, 0.5, 1.0, 5.0, 10.0, 50.0\}$.
- **N-HMC** We chose the value of q in (PI) by cross-validation on the set $\{10^1, 10^2, 10^3\}$ (which is the same range considered in the experiments in ([Rice et al., 2021](#))).

C.4. Hyperparameters for probabilistically robust learning

We ran sweeps over a range of hyperparameters for Algorithm 1. By selecting M from $\{1, 2, 5, 10, 20\}$, we found more samples from τ engendered higher levels of robustness. Thus, we use $M = 20$ throughout. We use a step size of $\eta_\alpha = 1.0$ throughout. T was also selected by cross validation from $\{1, 2, 5, 10, 20\}$. In general, it seemed to be the case that more than 10 steps did not result in significant improvements in robustness. Finally, in Tables 1–3, we selected ρ by cross-validation on $\{0.01, 0.05, 0.1, 0.5, 1.0, 2.0, 3.0, 4.0\}$ (note that in practice, ρ can be chosen to be larger than one). We found that perhaps surprisingly, larger values of ρ tended to engender higher levels of probabilistic robustness through the metric ProbAcc. However, this may be due to the instability of training for small values of ρ . In Figure 4, we show the robustness trade-offs for a sweep over different values of ρ .