
A Kernel Perspective for Regularizing Deep Neural Networks

Alberto Bietti^{*1} Grégoire Mialon^{*12} Dexiong Chen¹ Julien Mairal¹

Abstract

We propose a new point of view for regularizing deep neural networks by using the norm of a reproducing kernel Hilbert space (RKHS). Even though this norm cannot be computed, it admits upper and lower approximations leading to various practical strategies. Specifically, this perspective (i) provides a common umbrella for many existing regularization principles, including spectral norm and gradient penalties, or adversarial training, (ii) leads to new effective regularization penalties, and (iii) suggests hybrid strategies combining lower and upper bounds to get better approximations of the RKHS norm. We experimentally show this approach to be effective when learning on small datasets, or to obtain adversarially robust models.

1. Introduction

Learning predictive models for complex tasks often requires large amounts of annotated data. For instance, convolutional neural networks are huge-dimensional and typically involve more parameters than training samples, which raises several challenges: achieving good generalization with small datasets is indeed difficult, which limits the deployment of such deep models to many tasks where labeled data is scarce, *e.g.*, in biology (Ching et al., 2018). Besides, imperceptible adversarial perturbations can significantly degrade the prediction quality (Szegedy et al., 2013; Biggio & Roli, 2018). These issues raise the question of regularization as an essential tool to control the complexity of deep models, as well as their stability to small variations of their inputs.

In this paper, we present a new perspective on regularization of deep networks, by viewing convolutional neural networks (CNNs) as elements of a RKHS following the work of Bietti & Mairal (2019) on deep convolutional kernels. For

^{*}Equal contribution ¹Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK, 38000 Grenoble, France ²Département d’informatique de l’ENS, ENS, CNRS, Inria, PSL, 75005 Paris, France. Correspondence to: <firstname.lastname@inria.fr>.

such kernels, the RKHS contains indeed deep convolutional networks similar to generic ones—up to smooth approximations of rectified linear units. Such a point of view provides a natural regularization function, the RKHS norm, which allows us to control the variations of the predictive model and to limit its complexity for better generalization. Besides, the norm also acts as a Lipschitz constant, which provides a direct control on the stability to adversarial perturbations.

In contrast to traditional kernel methods, the RKHS norm cannot be explicitly computed in our setup. Yet, this norm admits numerous approximations—lower bounds and upper bounds—which lead to many strategies for regularization based on penalties, constraints, or combinations thereof. Depending on the chosen approximation, we recover then many existing principles such as spectral norm regularization (Cisse et al., 2017; Yoshida & Miyato, 2017; Miyato et al., 2018a; Sedghi et al., 2019), gradient penalties and double backpropagation (Drucker & Le Cun, 1991; Simon-Gabriel et al., 2019; Gulrajani et al., 2017; Roth et al., 2017; 2018; Arbel et al., 2018), adversarial training (Madry et al., 2018), and we also draw links with tangent propagation (Simard et al., 1998). For all these principles, we provide a unified viewpoint and theoretical insights, and we also introduce new variants, which we show are effective in practice when learning with few labeled data, or in the presence of adversarial perturbations.

Moreover, regularization and robustness are tightly linked in our kernel framework. Specifically, some lower bounds on the RKHS norm lead to robust optimization objectives with worst-case ℓ_2 perturbations; further, we can extend margin-based generalization bounds in the spirit of Bartlett et al. (2017); Boucheron et al. (2005) to the setting of *adversarially robust* generalization (see Schmidt et al., 2018), where an adversary can perturb test data. We also discuss connections between recent regularization strategies for training generative adversarial networks and approaches to generative modeling based on kernel two-sample tests (MMD) (Dzugaite et al., 2015; Li et al., 2017; Bińkowski et al., 2018).

Summary of the contributions.

- We introduce an RKHS perspective for regularizing deep neural networks models which provides a unified view on various practical regularization principles, together with theoretical insight and guarantees;

- By considering lower bounds to the RKHS norm, we obtain new penalties based on adversarial perturbations, adversarial deformations, or gradient norms of prediction functions, which we show to be effective in practice;
- Our RKHS point of view suggests combined strategies based on both upper and lower bounds, which we show often perform empirically best in the context of generalization from small image and biological datasets, by providing a tighter control of the RKHS norm.

Related work. The construction of hierarchical kernels and the study of neural networks in the corresponding RKHS was studied by Mairal (2016); Zhang et al. (2016; 2017); Bietti & Mairal (2019). Some of the regularization strategies we obtain from our kernel perspective are variants of previous approaches to adversarial robustness (Cisse et al., 2017; Madry et al., 2018; Simon-Gabriel et al., 2019; Roth et al., 2018), to improving generalization (Drucker & Le Cun, 1991; Miyato et al., 2018b; Sedghi et al., 2019; Simard et al., 1998; Yoshida & Miyato, 2017), and stable training of generative adversarial networks (Roth et al., 2017; Gulrajani et al., 2017; Arbel et al., 2018; Miyato et al., 2018a). The link between robust optimization and regularization was studied by Xu et al. (2009a;b), focusing mainly on linear models with quadratic or hinge losses. The notion of adversarial generalization was considered by Schmidt et al. (2018), who provide lower bounds on a particular data distribution. Sinha et al. (2018) provide generalization guarantees in the different setting of distributional robustness; compared to our bound, they consider expected loss instead of classification error, and their bounds do not highlight the dependence on the model complexity.

2. Regularization of Deep Neural Networks

In this section, we recall the kernel perspective on deep networks introduced by Bietti & Mairal (2019), and present upper and lower bounds on the RKHS norm of a given model, leading to various regularization strategies. For simplicity, we first consider real-valued networks and binary classification, before discussing multi-class extensions.

2.1. Relation between deep networks and RKHSs

Kernel methods consist of mapping data living in a set \mathcal{X} to a RKHS \mathcal{H} associated to a positive definite kernel K through a mapping function $\Phi : \mathcal{X} \rightarrow \mathcal{H}$, and then learning simple machine learning models in \mathcal{H} . Specifically, when considering a real-valued regression or binary classification problem, classical kernel methods find a prediction function $f : \mathcal{X} \rightarrow \mathbb{R}$ living in the RKHS which can be written in linear form, i.e., such that $f(x) = \langle f, \Phi(x) \rangle_{\mathcal{H}}$ for all x in \mathcal{X} . While explicit mapping to a possibly infinite-dimensional space is of course only an abstract mathematical operation, learning f can be done implicitly by com-

puting kernel evaluations and typically by using convex programming (Schölkopf & Smola, 2001).

Moreover, the RKHS norm $\|f\|_{\mathcal{H}}$ acts as a natural regularization function, which controls the variations of model predictions according to the geometry induced by Φ :

$$|f(x) - f(x')| \leq \|f\|_{\mathcal{H}} \cdot \|\Phi(x) - \Phi(x')\|_{\mathcal{H}}. \quad (1)$$

Unfortunately, our setup does not allow us to use the RKHS norm in a traditional way since evaluating the kernel is intractable. Instead, we propose a different approach that considers explicit parameterized representations of functions contained in the RKHS, given by generic CNNs, and leverage properties of the RKHS and the kernel mapping in order to regularize when learning the network parameters.

Consider indeed a real-valued deep convolutional network $f : \mathcal{X} \rightarrow \mathbb{R}$, where \mathcal{X} is simply \mathbb{R}^d , with rectified linear unit (ReLU) activations and no bias units. By constructing an appropriate multi-layer hierarchical kernel, Bietti & Mairal (2019) show that the corresponding RKHS \mathcal{H} contains a CNN with the same architecture and parameters as f , but with activations that are smooth approximations of ReLU. Although the model predictions might not be strictly equal, we will abuse notation and denote this approximation with smooth ReLU by f as well, with the hope that the regularization procedures derived from the RKHS model will be effective in practice on the original CNN f .

Besides, the mapping $\Phi(\cdot)$ is shown to be non-expansive:

$$\|\Phi(x) - \Phi(x')\|_{\mathcal{H}} \leq \|x - x'\|_2, \quad (2)$$

so that controlling $\|f\|_{\mathcal{H}}$ provides some robustness to additive ℓ_2 -perturbations, by (1). Additionally, with appropriate pooling operations, Bietti & Mairal (2019) show that the kernel mapping is also stable to deformations, meaning that the RKHS norm also controls robustness to translations and other transformations including scaling and rotations, which can be seen as deformations when they are small.

In contrast to standard kernel methods, where the RKHS norm is typically available in closed form, this norm is difficult to compute in our setup, and requires approximations. The following sections present upper and lower bounds on $\|f\|_{\mathcal{H}}$, with linear convolutional operations denoted by W_k for $k = 1, \dots, L$, where L is the number of layers. Defining $\theta := \{W_k : k = 1, \dots, L\}$, we then leverage these bounds to approximately solve the following penalized or constrained optimization problems on a training set $(x_i, y_i), i = 1, \dots, n$:

$$\min_{\theta} \frac{1}{n} \sum_{i=1}^n \ell(y_i, f_{\theta}(x_i)) + \lambda \|f_{\theta}\|_{\mathcal{H}}^2 \quad \text{or} \quad (3)$$

$$\min_{\theta: \|f_{\theta}\|_{\mathcal{H}} \leq C} \frac{1}{n} \sum_{i=1}^n \ell(y_i, f_{\theta}(x_i)). \quad (4)$$

We also note that while the construction of [Bietti & Mairal \(2019\)](#) considers VGG-like networks ([Simonyan & Zisserman, 2014](#)), the regularization algorithms we obtain in practice can be easily adapted to different architectures such as residual networks ([He et al., 2016](#)).

2.2. Exploiting lower bounds of the RKHS norm

In this section, we devise regularization algorithms by leveraging lower bounds on $\|f\|_{\mathcal{H}}$, obtained by relying on the following variational characterization of Hilbert norms:

$$\|f\|_{\mathcal{H}} = \sup_{\|u\|_{\mathcal{H}} \leq 1} \langle f, u \rangle_{\mathcal{H}}.$$

At first sight, this definition is not useful since the set $U = \{u \in \mathcal{H} : \|u\|_{\mathcal{H}} \leq 1\}$ may be infinite-dimensional and the inner products $\langle f, u \rangle_{\mathcal{H}}$ cannot be computed in general. Thus, we devise tractable lower bound approximations by considering smaller sets $\bar{U} \subset U$.

Adversarial perturbation penalty. Thanks to the non-expansiveness of Φ , we can consider the subset $\bar{U} \subset U$ defined as $\bar{U} = \{\Phi(x + \delta) - \Phi(x) : x \in \mathcal{X}, \|\delta\|_2 \leq 1\}$, leading to the bound

$$\|f\|_{\mathcal{H}} \geq \|f\|_{\delta}^2 := \sup_{x \in \mathcal{X}, \|\delta\|_2 \leq 1} f(x + \delta) - f(x), \quad (5)$$

which is reminiscent of adversarial perturbations. Adding a regularization parameter $\epsilon > 0$ in front of the norm then corresponds to different sizes of perturbations:

$$\epsilon \|f\|_{\mathcal{H}} = \sup_{\|u\|_{\mathcal{H}} \leq \epsilon} \langle f, u \rangle_{\mathcal{H}} \geq \sup_{x \in \mathcal{X}, \|\delta\|_2 \leq \epsilon} f(x + \delta) - f(x). \quad (6)$$

Using this lower bound or its square as a penalty in the objective (3) when training a CNN provides a way to regularize. Optimizing over adversarial perturbations has been useful to obtain robust models (e.g., the PGD method of [Madry et al., 2018](#)); yet our approach differs in two important ways:

(i) it involves a penalty that is decoupled from the loss term such that in principle, our penalty could be used beyond the supervised empirical risk paradigm. In contrast, PGD optimizes the robust formulation (7) below, which fits training data while considering perturbations on the loss.

(ii) our penalty involves a global maximization problem on the input space \mathcal{X} , as opposed to only maximizing on perturbations near training data. In practice, optimizing over \mathcal{X} is however difficult and instead, we replace \mathcal{X} by random mini-batches of examples, yielding further lower bounds on the RKHS norm. These examples may be labeled or not, in contrast to PGD that perturb labeled examples only. When using such a mini-batch, a gradient of the penalty can be obtained by first finding maximizers $\hat{x}, \hat{\delta}$ (where \hat{x} is an element of the mini-batch and $\hat{\delta}$ is a perturbation), and then

computing gradients of $f_{\theta}(\hat{x} + \hat{\delta}) - f_{\theta}(\hat{x})$ with respect to θ by using back-propagation. In practice, we compute the perturbations δ for each example x by using a few steps of projected gradient ascent with constant step-lengths.

Robust optimization yields another lower bound. In some contexts, our penalized approach is related to solving the robust optimization problem

$$\min_{\theta} \frac{1}{n} \sum_{i=1}^n \sup_{\|\delta\|_2 \leq \epsilon} \ell(y_i, f_{\theta}(x_i + \delta)), \quad (7)$$

which is commonly considered for training adversarially robust classifiers ([Wong & Kolter, 2018](#); [Madry et al., 2018](#); [Raghunathan et al., 2018](#)). In particular, [Xu et al. \(2009b\)](#) show that the penalized and robust objectives are equivalent in the case of the hinge loss with linear predictors, when the data is non-separable. They also show the equivalence for kernel methods when considering the (intractable) full perturbation set U around each point in the RKHS $\Phi(x_i)$, that is, predictions $\langle f, \Phi(x_i) + u \rangle_{\mathcal{H}}$ with $u \in U$. Intuitively, when a training example (x_i, y_i) is misclassified, we are in the “linear” part of the hinge loss, such that

$$\sup_{\|u\|_{\mathcal{H}} \leq \epsilon} \ell(y_i, \langle f, \Phi(x_i) + u \rangle_{\mathcal{H}}) = \ell(y_i, f(x_i)) + \epsilon \|f\|_{\mathcal{H}}.$$

For other losses such as the logistic loss, a regularization effect is still present even for correctly classified examples, though it may be smaller since the loss has a reduced slope for such points. This leads to an *adaptive* regularization mechanism that may automatically reduce the amount of regularization when the data is easily separable. However, the robust optimization approach might only encourage local stability around training examples, while the global quantity $\|f\|_{\mathcal{H}}$ may become large in order to better fit the data. We note that a perfect fit of the data with large complexity does not prevent generalization (see, e.g., [Belkin et al., 2018a;b](#)); yet, such mechanisms are still poorly understood. Nevertheless, it is easy to show that the robust objective (7) lower bounds the penalized objective with penalty $\epsilon \|f\|_{\mathcal{H}}$.

Gradient penalties. Taking $\bar{U} = \{\frac{\Phi(x) - \Phi(y)}{\|x - y\|_2} : x, y \in \mathcal{X}\}$, which is a subset of U by Eq. (2)—it turns out that this is the same set as for adversarial perturbation penalties, since Φ is homogeneous ([Bietti & Mairal, 2019](#)) and $\mathcal{X} = \mathbb{R}^d$ —we obtain a lower bound based on the Lipschitz constant of f :

$$\|f\|_{\mathcal{H}} \geq \sup_{x, y \in \mathcal{X}} \frac{f(x) - f(y)}{\|x - y\|_2} \geq \|\nabla f\| := \sup_{x \in \mathcal{X}} \|\nabla f(x)\|_2, \quad (8)$$

where the second inequality becomes an equality when \mathcal{X} is convex, and the supremum is taken over points where f is differentiable. Although we are unaware of previous work using this exact lower bound for a generic regularization

penalty, we note that variants replacing the supremum over x by an expectation over data have been recently used to stabilize the training of generative adversarial networks (Gulrajani et al., 2017; Roth et al., 2017), and we provide insights in Section 3.2 on the benefits of RKHS regularization in such a setting. Related penalties have been considered in the context of robust optimization, for regularization or robustness, noting that a penalty based on the gradient of the loss function $x \mapsto \ell(y, f(x))$ can give a good approximation of (7) when ϵ is small (Drucker & Le Cun, 1991; Lyu et al., 2015; Roth et al., 2018; Simon-Gabriel et al., 2019).

Penalties based on deformation stability. We may also obtain new penalties by considering more exotic sets $\bar{U} = \{\Phi(\tilde{x}) - \Phi(x) : x \in \mathcal{X}, \tilde{x} \text{ is a small deformation of } x\}$, where the amount of deformation is dictated by the stability bounds of Bietti & Mairal (2019) in order to ensure that $\bar{U} \subset U$. More precisely, such bounds depend on the maximum displacement and Jacobian norm of the diffeomorphisms considered. These can be easily computed for various parameterized families of transformations, such as translations, scaling or rotations, leading to simple ways to control the regularization strength through the parameters of these transformations. One can also consider infinitesimal deformations from such parameterized transformations, which approximately yields the *tangent propagation* regularization strategy of Simard et al. (1998). These approaches are detailed in Appendix B. If instead we consider the robust optimization formulation (7), we obtain a form of *data augmentation* where transformations are optimized instead of sampled, as done by (Engstrom et al., 2017).

Extensions to multiple classes and beyond We now extend the regularization strategies based on lower bounds to multi-valued networks, in order to deal with multiple classes. For that purpose, we consider a multi-class penalty $\|f_1\|_{\mathcal{H}}^2 + \dots + \|f_K\|_{\mathcal{H}}^2$ for an \mathbb{R}^K -valued function $f = (f_1, f_2, \dots, f_K)$, and we define

$$\|f\|_{\delta}^2 := \sum_{k=1}^K \|f_k\|_{\delta}^2 \quad \text{and} \quad \|\nabla f\|^2 := \sum_{k=1}^K \|\nabla f_k\|^2,$$

where $\|f_k\|_{\delta}$ is the adversarial penalty (5), and $\|\nabla f_k\|$ is defined in (8). For deformation stability penalties, we proceed in a similar manner, and for robust optimization formulations (7), the extension is straightforward, given that multi-class losses such as cross-entropy can be directly optimized in an adversarial training or gradient penalty setup.

Finally, we note that while the kernel approach we introduce considers the Euclidian geometry in the input space, it is possible to consider heuristic alternatives for other geometries, such as ℓ_{∞} perturbations, as discussed in Appendix D.

2.3. Exploiting upper bounds with spectral norms

Instead of lower bounds, one may use instead the following upper bound from Bietti & Mairal (2019, Proposition 14):

$$\|f\|_{\mathcal{H}} \leq \omega(\|W_1\|, \dots, \|W_L\|), \quad (9)$$

where ω is increasing in all of its arguments, and $\|W_k\|$ is the spectral norm of the linear operator W_k . Here, we simply consider the spectral norm on the filters, given by $\|W\| := \sup_{\|x\|_2 \leq 1} \|Wx\|_2$. Other generalization bounds relying on similar quantities have been proposed for controlling complexity (Bartlett et al., 2017; Neyshabur et al., 2018), suggesting that using them for regularization is relevant even beyond our kernel perspective, as observed by Cisse et al. (2017); Sedghi et al. (2019); Yoshida & Miyato (2017). Extensions to multiple classes are simple to obtain by simply considering spectral norms up to the last layer.

Penalizing the spectral norms. One way to control the upper bound (9) when learning a neural network f_{θ} is to consider a regularization penalty based on spectral norms

$$\min_{\theta} \frac{1}{n} \sum_{i=1}^n \ell(y_i, f_{\theta}(x_i)) + \lambda \sum_{l=1}^L \|W_l\|^2, \quad (10)$$

where λ is a regularization parameter. To optimize this cost, one can obtain (sub)gradients of the penalty by computing singular vectors associated to the largest singular value of each W_l . We consider the method of Yoshida & Miyato (2017), which computes such singular vectors approximately using one or two iterations of the power method, as well as a more costly approach using the full SVD.

Constraining the spectral norms with a continuation approach. In the constrained setting, we want to optimize:

$$\min_{\theta} \frac{1}{n} \sum_{i=1}^n \ell(y_i, f_{\theta}(x_i)) \quad \text{s.t.} \quad \|W_l\| \leq \tau; l \in 1, \dots, L,$$

where τ is a user-defined constraint. This objective may be optimized by projecting each W_l in the spectral norm ball of radius τ after each gradient step. Such a projection is achieved by truncating the singular values to be smaller than τ (see Appendix C). We found that the loss was hardly optimized with this approach, and therefore introduce a continuation approach with an exponentially decaying schedule for τ reaching a constant τ_0 after a few epochs, which we found to be important for good empirical performance.

2.4. Combining upper and lower bounds.

One advantage of lower bound penalties is that they are independent of the model parameterization, making them flexible enough to use with more complex architectures. In

addition, the connection with robust optimization can provide a useful mechanism for adaptive regularization. However, they do not provide a guaranteed control on the RKHS norm, unlike the upper bound strategies. This is particularly true for robust optimization approaches, which may favor small training loss and local stability over global stability through $\|f\|_{\mathcal{H}}$. Nevertheless, we observed that our new approaches based on separate penalties sometimes do help in controlling upper bounds as well (see Section 4).

While these upper bound strategies are useful for limiting model complexity, we found them empirically less effective for robustness (see Section 4.2). However, we observed that combining with lower bound approaches can overcome this weakness, perhaps due to a better control of local stability. In particular, such combined approaches often provide the best generalization performance in small data scenarios, as well as better guarantees on adversarially robust generalization thanks to a tighter control of the RKHS norm.

3. Theoretical Guarantees and Insights

In this section, we study how the kernel perspective allows us to extend standard margin-based generalization bounds to an adversarial setting in order to provide theoretical guarantees on adversarially robust generalization. We then discuss how our kernel approach provides novel interpretations for training generative adversarial networks.

3.1. Guarantees on adversarial generalization

While various methods have been introduced to empirically gain robustness to adversarial perturbations, the ability to generalize with such perturbations, also known as *adversarial generalization* (Schmidt et al., 2018), still lacks theoretical understanding. Margin-based bounds have been useful to explain the generalization behavior of learning algorithms that can fit the training data well, such as kernel methods, boosting and neural networks (Koltchinskii & Panchenko, 2002; Boucheron et al., 2005; Bartlett et al., 2017). Here, we show how such arguments can be adapted to obtain guarantees on adversarial generalization, *i.e.*, on the expected classification error in the presence of an ℓ_2 -bounded adversary, based on the RKHS norm of a learned model. For a binary classification task with labels in $\mathcal{Y} = \{-1, 1\}$ and data distribution \mathcal{D} , we would like to bound the expected adversarial error of a classifier f , given for some $\epsilon > 0$ by

$$\text{err}_{\mathcal{D}}(f, \epsilon) := P_{(x,y) \sim \mathcal{D}}(\exists \|\delta\|_2 \leq \epsilon : yf(x + \delta) < 0). \quad (11)$$

Leveraging the fact that f is $\|f\|_{\mathcal{H}}$ -Lipschitz, we now show how to further bound this quantity using empirical margins, following the usual approach to obtaining margin bounds for kernel methods (*e.g.*, Boucheron et al., 2005). Consider a training dataset $(x_1, y_1), \dots, (x_n, y_n) \in \mathcal{X} \times \mathcal{Y}$. Defining

$L_n^\gamma(f) := \frac{1}{n} \sum_{i=1}^n \mathbf{1}\{y_i f(x_i) < \gamma\}$, we have the following bound, proved in Appendix E:

Proposition 1 (Adversarially robust margin bound). *With probability $1 - \delta$ over a dataset $\{(x_i, y_i)\}_{i=1, \dots, n}$, we have, for all choices of $\gamma > 0$ and $f \in \mathcal{H}$,*

$$\text{err}_{\mathcal{D}}(f, \epsilon) \leq L_n^{\gamma+2\epsilon\|f\|_{\mathcal{H}}}(f) + \tilde{O}\left(\frac{\|f\|_{\mathcal{H}} \bar{B}}{\gamma \sqrt{n}}\right), \quad (12)$$

where $\bar{B} = \sqrt{\frac{1}{n} \sum_{i=1}^n K(x_i, x_i)}$ and \tilde{O} hides a term depending logarithmically on $\|f\|_{\mathcal{H}}$, γ , and δ .

When $\epsilon = 0$, we obtain the usual margin bound, while $\epsilon > 0$ yields a bound on adversarial error $\text{err}_{\mathcal{D}}(f, \epsilon)$, for some neural network f learned from data. Note that other complexity measures based on products of spectral norms may be used instead of $\|f\|_{\mathcal{H}}$, as well as multi-class extensions, following Bartlett et al. (2017); Neyshabur et al. (2018). In concurrent work, Khim & Loh (2018); Yin et al. (2019) derive similar bounds in the context of fully-connected networks. In contrast to these works, which bound complexity of a modified function class, our bound uses the complexity of the original class and leverages smoothness properties of functions to derive the margin bound.

One can then study the effectiveness of a regularization algorithm by inspecting cumulative distribution (CDF) plots of the *normalized margins* $\bar{\gamma}_i = y_i f(x_i) / \|f\|_{\mathcal{H}}$, for different strengths of regularization (an example is given in Figure 2, Section 4.2). According to the bound (12), one can assess expected adversarial error with ϵ -bounded perturbations by looking at the part of the plot to the right of $\bar{\gamma} = 2\epsilon$. In particular, the value of the CDF at such a value of $\bar{\gamma}$ is representative of the bound for large n (since the second term is negligible), while for smaller n , the best bound is obtained for a larger value of $\bar{\gamma}$, which also suggests that the right side of the plots is indicative of performance on small datasets.

When the RKHS norm can be well approximated, our bound provides a certificate on test error in the presence of adversaries. While such an approximation is difficult to obtain in general, the guarantee is most useful when lower and upper bounds of the RKHS norm are controlled together.

3.2. New insights on generative adversarial networks

Generative adversarial networks (GANs) attempt to learn a *generator* neural network $G_\phi : \mathcal{Z} \rightarrow \mathcal{X}$, so that the distribution of $G_\phi(z)$ with $z \sim D_z$ a noise vector resembles a data distribution D_x . In this section, we discuss connections between recent regularization techniques for training GANs, and approaches to learning generative models based on a MMD criterion (Gretton et al., 2012), in view of our RKHS framework. Our goal is to provide a new insight on these methods, but not necessarily to provide a new one.

Various recent approaches have relied on regularization strategies on a *discriminator* network in order to improve the stability of GAN training and the quality of the produced samples. Some of these resemble the approaches presented in Section 2 such as gradient penalties (Gulrajani et al., 2017; Roth et al., 2017) and spectral norm regularization (Miyato et al., 2018a). We provide an RKHS interpretation of these methods as optimizing an MMD distance with the convolutional kernel introduced in Section 2:

$$\min_{\phi} \sup_{\|f\|_{\mathcal{H}} \leq 1} \mathbb{E}_{x \sim D_x} [f(x)] - \mathbb{E}_{z \sim D_z} [f(G_{\phi}(z))]. \quad (13)$$

When learning from an empirical distribution over n samples, the MMD criterion is known to have much better sample complexity than the Wasserstein-1 distance considered by Arjovsky et al. (2017) for high-dimensional data such as images (Sriperumbudur et al., 2012). While the MMD approach has been used for training generative models, it generally relies on a generic kernel function, such as a Gaussian kernel, that appears explicitly in the objective (Dziugaite et al., 2015; Li et al., 2017; Bińkowski et al., 2018). Although using a learned feature extractor can improve this, the Gaussian kernel might be a poor choice when dealing with natural signals such as images, while the hierarchical kernel we consider in our paper is better suited for this type of data, by providing useful invariance and stability properties. Leveraging the variational form of the MMD (13) with this kernel suggests for instance using convolutional networks as the discriminator f , with constraints on the spectral norms in order to ensure $\|f\|_{\mathcal{H}} \leq C$ for some C , as done by Miyato et al. (2018a) through normalization.

4. Experiments

We tested the regularization strategies presented in Section 2 in the context of improving generalization on small datasets and training robust models. Our goal is to use common architectures used for large datasets and improve their performance in different settings through regularization. Our Pytorch implementation of the various strategies is available at https://github.com/albietz/kernel_reg.

For the adversarial training strategies, the inner maximization problems are solved using 5 steps of projected gradient ascent with constant step-lengths. In the case of the lower bound penalties $\|f\|_{\delta}^2$ and $\|\nabla f\|^2$, we also maximize over examples in the mini-batch, only considering the maximal element when computing gradients with respect to parameters. For the robust optimization problem (7), we use PGD with ℓ_2 perturbations, as well as the corresponding ℓ_2 (squared) gradient norm penalty on the loss. For the upper bound approaches with spectral norms (SNs), we consider the SN projection strategy with decaying τ , as well as the SN penalty (10), either using power iteration (PI) or a full SVD for computing gradients.

Table 1. Regularization on CIFAR10 with 1 000 examples for VGG-11 and ResNet-18. Each entry shows the test accuracy with/without data augmentation when all hyper-parameters are optimized on a validation set. See also Section A.1 in the appendix for additional results and statistical testing.

Method	1k VGG-11	1k ResNet-18
No weight decay	50.70 / 43.75	45.23 / 37.12
Weight decay	51.32 / 43.95	44.85 / 37.09
SN penalty (PI)	54.64 / 45.06	47.01 / 39.63
SN projection	54.14 / 46.70	47.12 / 37.28
VAT	50.88 / 43.36	47.47 / 42.82
PGD- ℓ_2	51.25 / 44.40	45.80 / 41.87
grad- ℓ_2	55.19 / 43.88	49.30 / 44.65
$\ f\ _{\delta}^2$ penalty	51.41 / 45.07	48.73 / 43.72
$\ \nabla f\ ^2$ penalty	54.80 / 46.37	48.99 / 44.97
PGD- ℓ_2 + SN proj	54.19 / 46.66	47.47 / 41.25
grad- ℓ_2 + SN proj	55.32 / 46.88	48.73 / 42.78
$\ f\ _{\delta}^2$ + SN proj	54.02 / 46.72	48.12 / 43.56
$\ \nabla f\ ^2$ + SN proj	55.24 / 46.80	49.06 / 44.92

4.1. Improving generalization on small datasets

We consider the datasets CIFAR10 and MNIST when using a small number of training examples, as well as 102 datasets of biological sequences that suffer from small sample size.

CIFAR10. In this setting, we use 1 000 and 5 000 examples of the CIFAR10 dataset, with or without data augmentation. We consider a VGG network (Simonyan & Zisserman, 2014) with 11 layers, as well as a residual network (He et al., 2016) with 18 layers, which achieve 91% and 93% test accuracy respectively when trained on the full training set with standard data augmentation (horizontal flips + random crops). We do not use batch normalization layers in order to prevent any interaction with spectral norms. Each strategy derived in Section 2 is trained for 500 epochs using SGD with momentum and batch size 128, halving the step-size every 40 epochs. In order to study the potential effectiveness of each method, we assume that a reasonably large validation set is available to select hyper-parameters; thus, we keep 10 000 annotated examples for this purpose. We also show results using a smaller validation set in Appendix A.1.

Table 1 shows the test accuracies on 1 000 examples for upper and lower bound approaches, as well as combined ones. We also include virtual adversarial training (VAT, Miyato et al., 2018b). We provide extended tables in Appendix A.1 with additional methods, other geometries, results for 5 000 examples, as well as hypothesis tests for comparing pairs of methods and assessing the significance of our findings. Overall, we find that the combined lower bound + SN constraints approaches often yield better results than either method separately. For lower bound approaches alone, we found our $\|f\|_{\delta}^2$ and $\|\nabla f\|^2$ penalties to often work best, particularly without data augmentation, while robust optimization strategies can be preferable with data

Table 2. Regularization on 300 or 1 000 examples from MNIST, using deformations from Infinite MNIST. (*) indicates that random deformations were included as training examples, while $\|f\|_\tau^2$ and $\|D_\tau f\|^2$ use them as part of the regularization penalty. See Section A.2 in the appendix for more results and statistical testing.

Method	300 VGG	1k VGG
Weight decay	89.32	94.08
SN projection	90.69	95.01
grad- ℓ_2	93.63	96.67
$\ f\ _\delta^2$ penalty	94.17	96.99
$\ \nabla f\ ^2$ penalty	94.08	96.82
Weight decay (*)	92.41	95.64
grad- ℓ_2 (*)	95.05	97.48
$\ D_\tau f\ ^2$ penalty	94.18	96.98
$\ f\ _\tau^2$ penalty	94.42	97.13
$\ f\ _\tau^2 + \ \nabla f\ ^2$	94.75	97.40
$\ f\ _\tau^2 + \ f\ _\delta^2$	95.23	97.66
$\ f\ _\tau^2 + \ f\ _\delta^2$ (*)	95.53	97.56
$\ f\ _\tau^2 + \ f\ _\delta^2 + \text{SN proj}$	95.20	97.60
$\ f\ _\tau^2 + \ f\ _\delta^2 + \text{SN proj} (*)$	95.40	97.77

augmentation, perhaps thanks to the adaptive regularization effect discussed earlier, which may be helpful in this easier setting. Gradient penalties often outperform adversarial perturbation strategies, possibly because of the closed form gradients which may improve optimization. We also found that adversarial training strategies tend to poorly control SNs compared to gradient penalties, particularly PGD (see also Section 4.2). SN constraints alone can also work well in some cases, particularly for VGG architectures, and often outperform SN penalties. SN penalties can work well nevertheless and provide computational benefits when using the power iteration variant.

Infinite MNIST. In order to assess the effectiveness of lower bound penalties based on deformation stability, we consider the Infinite MNIST dataset (Loosli et al., 2007), which provides an “infinite” number of transformed generated examples for each of the 60 000 MNIST training digits. Here, we use a 5-layer VGG-like network with average pooling after each 3x3 convolution layer, in order to more closely match the architecture assumptions of Bietti & Mairal (2019) for deformation stability. We consider two lower bound penalties that leverage the digit transformations in Infinite MNIST: one based on “adversarial” deformations around each digit, denoted $\|f\|_\tau^2$; and a tangent propagation (Simard et al., 1998) variant, denoted $\|D_\tau f\|^2$, which provides an approximation to $\|f\|_\tau^2$ for small deformations based on gradients along a few tangent vector directions given by deformations (see Appendix B for details). Table 2 shows the obtained test accuracy for subsets of MNIST of size 300 and 1 000. Overall, we find that combining both adversarial penalties $\|f\|_\tau^2$ and $\|f\|_\delta^2$ performs best, which suggests that it is helpful to obtain tighter lower approximations of the RKHS norm by considering perturbations of

Table 3. Regularization on protein homology detection tasks, with or without data augmentation (DA). Fixed hyperparameters are selected using the first half of the datasets, and we report the average auROC50 score on the second half. See Section A.3 in the appendix for more details and statistical testing.

Method	No DA	DA
No weight decay	0.421	0.541
Weight decay	0.432	0.544
SN proj	0.583	0.615
PGD- ℓ_2	0.488	0.554
grad- ℓ_2	0.551	0.570
$\ f\ _\delta^2$	0.577	0.611
$\ \nabla f\ ^2$	0.566	0.598
PGD- ℓ_2 + SN proj	0.615	0.622
grad- ℓ_2 + SN proj	0.581	0.634
$\ f\ _\delta^2$ + SN proj	0.631	0.639
$\ \nabla f\ ^2$ + SN proj	0.576	0.617

different kinds. Explicitly controlling the spectral norms can further improve performance, as does training on deformed digits, which may yield better margins by exploiting the additional knowledge that small deformations preserve labels. Note that data augmentation alone (with some weight decay) does quite poorly in this case, even compared to our lower bound penalties which do not use deformations.

Protein homology detection. Remote homology detection between protein sequences is an important problem to understand protein structure. Given a protein sequence, the goal is to predict whether it belongs to a superfamily of interest. We consider the Structural Classification Of Proteins (SCOP) version 1.67 dataset (Murzin et al., 1995), which we process as described in Appendix A.3 in order to obtain 102 balanced binary classification tasks with 100 protein sequences each, thus resulting in a low-sample regime. Protein sequences were also cut to 400 amino acids.

Sequences are represented with a one-hot encoding strategy—that is, a sequence of length l is represented as a binary matrix in $\{0, 1\}^{20 \times l}$, where 20 is the number of different amino acids (alphabet size of the sequences). Such a structure can then be processed by convolutional neural networks (Alipanahi et al., 2015). In this paper, we do not try to optimize the structure of the network for the task, since our goal is only to evaluate the effect of regularization strategies. Therefore, we use a simple convolutional network with 3 convolutional layers followed by global max-pooling and a final fully-connected layer (we use filters of size 5, and a max-pooling layer after the second convolutional layer).

Training was done using Adam with a learning rate fixed to 0.01, and a weight decay parameter tuned for each method. Since hyper-parameter selection per dataset is difficult due to the low sample size, we use the same parameters across datasets. This allows us to use the first 51 datasets as a validation set for hyper-parameter tuning, and we report




Figure 1. Robustness trade-off curves of different regularization methods for VGG11 on CIFAR10. Each plot shows test accuracy vs adversarial test accuracy for ℓ_2 -bounded, 40-step PGD adversaries with a fixed ϵ_{test} . Different points on a curve correspond to training with different regularization strengths. The regularization increases monotonically along a given curve, and the leftmost points correspond to the strongest regularization. For PGD- ℓ_2 + SN projection, we vary ϵ with a fixed $\tau = 0.8$.

average performance with these fixed choices on the remaining 51 datasets. The standard performance measure for this task is the auROC50 score (area under the ROC curve up to 50% false positives). We note that the selection of hyperparameters has a transductive component, since some of the sequences in the test datasets may also appear in the datasets used for validation (possibly with a different label).

The results are shown in Table 3. The procedure used for data augmentation (right column) is described in Appendix A.3. We found that the most effective approach is the adversarial perturbation penalty, together with SN constraints. In particular, we found it to outperform the gradient penalty $\|\nabla f\|^2$, perhaps because in this case gradient penalties are only computed on a discrete set of possible points given by one-hot encodings, while adversarial perturbations may increase stability to wider regions, potentially covering different possible encoded sequences.

4.2. Training adversarially robust models

We consider the same VGG architecture as in Section 4.1, trained on CIFAR10 with data augmentation, with different regularization strategies. Each method is trained for 300 epochs using SGD with momentum and batch size 128, dividing the step-size in half every 30 epochs. This strategy was successful in reaching convergence for all methods.

Figure 1 shows the test accuracy of the different methods in the presence of ℓ_2 -bounded adversaries, plotted against standard accuracy. We can see that the robust optimization approaches tend to work better in high-accuracy regimes, perhaps because the local stability that they encourage is sufficient on this dataset, while the $\|f\|_\delta^2$ penalty can be useful in large-perturbation regimes. We find that upper bound approaches alone do not provide robust models, but combining the SN constraint approach with a lower bound




Figure 2. (left) Comparison of lower and upper bound quantities ($\|f\|_\delta$ vs the product of spectral norms). (right) CDF plot of normalized empirical margins for the $\|\nabla f\|^2$ penalty with different regularization strengths, normalized by $\|f\|_\delta$. We consider 1000 fixed training examples when computing $\|f\|_\delta$.

strategy (in this case PGD- ℓ_2) helps improve robustness perhaps thanks to a more explicit control of stability. The plots also confirm that gradient penalties on the loss may be preferable for small regularization strengths (they achieve higher accuracy while improving robustness for small ϵ_{test}), while for stronger regularization, the gradient approximation no longer holds and the adversarial training approaches such as PGD (and its combination with SN constraints) are preferred. More experiments confirming these findings are available in Section A.4 of the appendix.

Norm comparison and adversarial generalization. Figure 2 (left) compares lower and upper bound quantities for different regularization strengths. Note that for PGD, in contrast to other methods, we can see that the product of spectral norms (representative of an upper bound on $\|f\|_{\mathcal{H}}$) increases when the lower bound $\|f\|_\delta$ decreases. This suggests that a network learned with PGD with large ϵ may have large RKHS norm, possibly because the approach tries to separate ϵ -balls around the training examples, which may require a more complex model than simply separating the training examples (see also Madry et al., 2018). This large discrepancy between upper and lower bounds highlights the fact that such models may only be stable locally near training data, though this happens to be enough for robustness on many test examples on CIFAR10.

In contrast, for other methods, and in particular the lower bound penalties $\|f\|_\delta^2$ and $\|\nabla f\|^2$, the upper and lower bounds appear more tightly controlled, suggesting a more appropriate control of the RKHS norm. This makes our guarantees on adversarial generalization more meaningful, and thus we may look at the empirical distributions of normalized margins $\bar{\gamma}$ obtained using $\|f\|_\delta$ for normalization (as an approximation of $\|f\|_{\mathcal{H}}$), shown in Figure 2 (right). The curves suggest that for small $\bar{\gamma}$, and hence small ϵ_{test} , smaller values of λ are preferred, while stronger regularization helps for larger $\bar{\gamma}$, yielding lower test error guarantees in the presence of stronger adversaries according to our bounds in Section 3.1. This qualitative behavior is indeed observed in the results of Figure 1 on test data for the $\|\nabla f\|^2$ penalty.

Acknowledgements

This work was supported by the ERC grant number 714381 (SOLARIS project) and by the MSR-Inria joint centre.

References

- Alipanahi, B., DeLong, A., Weirauch, M. T., and Frey, B. J. Predicting the sequence specificities of dna-and rna-binding proteins by deep learning. *Nature biotechnology*, 33(8):831, 2015.
- Altschul, S. F., Madden, T. L., Schäffer, A. A., Zhang, J., Zhang, Z., Miller, W., and Lipman, D. J. Gapped blast and psi-blast: a new generation of protein database search programs. *Nucleic acids research*, 25(17):3389–3402, 1997.
- Arbel, M., Sutherland, D. J., Bińkowski, M., and Gretton, A. On gradient regularizers for MMD GANs. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- Arjovsky, M., Chintala, S., and Bottou, L. Wasserstein GAN. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2017.
- Bartlett, P. L., Foster, D. J., and Telgarsky, M. J. Spectrally-normalized margin bounds for neural networks. In *Advances in Neural Information Processing Systems (NIPS)*, 2017.
- Belkin, M., Hsu, D., and Mitra, P. Overfitting or perfect fitting? risk bounds for classification and regression rules that interpolate. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018a.
- Belkin, M., Ma, S., and Mandal, S. To understand deep learning we need to understand kernel learning. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2018b.
- Bietti, A. and Mairal, J. Group invariance, stability to deformations, and complexity of deep convolutional representations. *Journal of Machine Learning Research (JMLR)*, 20(25):1–49, 2019.
- Biggio, B. and Roli, F. Wild patterns: Ten years after the rise of adversarial machine learning. *Pattern Recognition*, 84:317–331, 2018.
- Bińkowski, M., Sutherland, D. J., Arbel, M., and Gretton, A. Demystifying MMD GANs. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2018.
- Boucheron, S., Bousquet, O., and Lugosi, G. Theory of classification: A survey of some recent advances. *ESAIM: probability and statistics*, 9:323–375, 2005.
- Ching, T. et al. Opportunities and obstacles for deep learning in biology and medicine. *Journal of The Royal Society Interface*, 15(141), 2018.
- Cisse, M., Bojanowski, P., Grave, E., Dauphin, Y., and Usunier, N. Parseval networks: Improving robustness to adversarial examples. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2017.
- Drucker, H. and Le Cun, Y. Double backpropagation increasing generalization performance. In *International Joint Conference on Neural Networks (IJCNN)*, 1991.
- Dziugaite, G. K., Roy, D. M., and Ghahramani, Z. Training generative neural networks via maximum mean discrepancy optimization. In *Conference on Uncertainty in Artificial Intelligence (UAI)*, 2015.
- Engstrom, L., Tsipras, D., Schmidt, L., and Madry, A. A rotation and a translation suffice: Fooling cnns with simple transformations. *arXiv preprint arXiv:1712.02779*, 2017.
- Gretton, A., Borgwardt, K. M., Rasch, M. J., Schölkopf, B., and Smola, A. A kernel two-sample test. *Journal of Machine Learning Research (JMLR)*, 13(Mar):723–773, 2012.
- Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., and Courville, A. C. Improved training of Wasserstein GANs. In *Advances in Neural Information Processing Systems (NIPS)*, 2017.
- Håndstad, T., Hestnes, A. J., and Sætrum, P. Motif kernel generated by genetic programming improves remote homology and fold detection. *BMC bioinformatics*, 8(1):23, 2007.
- 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 (CVPR)*, 2016.
- Kakade, S. M., Sridharan, K., and Tewari, A. On the complexity of linear prediction: Risk bounds, margin bounds, and regularization. In *Advances in Neural Information Processing Systems (NIPS)*, 2009.
- Khim, J. and Loh, P.-L. Adversarial risk bounds via function transformation. *arXiv preprint arXiv:1810.09519*, 2018.
- Koltchinskii, V. and Panchenko, D. Empirical margin distributions and bounding the generalization error of combined classifiers. *The Annals of Statistics*, 30(1):1–50, 2002.
- Li, C.-L., Chang, W.-C., Cheng, Y., Yang, Y., and Póczos, B. Mmd gan: Towards deeper understanding of moment matching network. In *Advances in Neural Information Processing Systems (NIPS)*, 2017.

- Loosli, G., Canu, S., and Bottou, L. Training invariant support vector machines using selective sampling. In Bottou, L., Chapelle, O., DeCoste, D., and Weston, J. (eds.), *Large Scale Kernel Machines*, pp. 301–320. MIT Press, Cambridge, MA., 2007.
- Lyu, C., Huang, K., and Liang, H.-N. A unified gradient regularization family for adversarial examples. In *IEEE International Conference on Data Mining (ICDM)*, 2015.
- Madry, A., Makelov, A., Schmidt, L., Tsipras, D., and Vladu, A. Towards deep learning models resistant to adversarial attacks. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2018.
- Mairal, J. End-to-end kernel learning with supervised convolutional kernel networks. In *Advances in Neural Information Processing Systems (NIPS)*, 2016.
- Miyato, T., Kataoka, T., Koyama, M., and Yoshida, Y. Spectral normalization for generative adversarial networks. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2018a.
- Miyato, T., Maeda, S.-i., Ishii, S., and Koyama, M. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 2018b.
- Murzin, A. G., Brenner, S. E., Hubbard, T., and Chothia, C. Scop: a structural classification of proteins database for the investigation of sequences and structures. *Journal of molecular biology*, 247(4):536–540, 1995.
- Neyshabur, B., Bhojanapalli, S., McAllester, D., and Srebro, N. A PAC-Bayesian approach to spectrally-normalized margin bounds for neural networks. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2018.
- Raghunathan, A., Steinhardt, J., and Liang, P. Certified defenses against adversarial examples. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2018.
- Roth, K., Lucchi, A., Nowozin, S., and Hofmann, T. Stabilizing training of generative adversarial networks through regularization. In *Advances in Neural Information Processing Systems (NIPS)*, 2017.
- Roth, K., Lucchi, A., Nowozin, S., and Hofmann, T. Adversarially robust training through structured gradient regularization. *arXiv preprint arXiv:1805.08736*, 2018.
- Schmidt, L., Santurkar, S., Tsipras, D., Talwar, K., and Madry, A. Adversarially robust generalization requires more data. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- Schölkopf, B. and Smola, A. J. *Learning with kernels: support vector machines, regularization, optimization, and beyond*. 2001.
- Sedghi, H., Gupta, V., and Long, P. M. The singular values of convolutional layers. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2019.
- Simard, P. Y., LeCun, Y. A., Denker, J. S., and Victorri, B. Transformation invariance in pattern recognition—tangent distance and tangent propagation. In *Neural networks: tricks of the trade*, pp. 239–274. Springer, 1998.
- Simon-Gabriel, C.-J., Ollivier, Y., Bottou, L., Schölkopf, B., and Lopez-Paz, D. First-order adversarial vulnerability of neural networks and input dimension. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2019.
- Simonyan, K. and Zisserman, A. Very deep convolutional networks for large-scale image recognition. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2014.
- Sinha, A., Namkoong, H., and Duchi, J. Certifying some distributional robustness with principled adversarial training. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2018.
- Sriperumbudur, B. K., Fukumizu, K., Gretton, A., Schölkopf, B., Lanckriet, G. R., et al. On the empirical estimation of integral probability metrics. *Electronic Journal of Statistics*, 6:1550–1599, 2012.
- 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.
- Tsipras, D., Santurkar, S., Engstrom, L., Turner, A., and Madry, A. Robustness may be at odds with accuracy. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2019.
- Wong, E. and Kolter, J. Z. Provable defenses against adversarial examples via the convex outer adversarial polytope. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2018.
- Xu, H., Caramanis, C., and Mannor, S. Robust regression and lasso. In *Advances in Neural Information Processing Systems (NIPS)*, 2009a.
- Xu, H., Caramanis, C., and Mannor, S. Robustness and regularization of support vector machines. *Journal of Machine Learning Research (JMLR)*, 10(Jul):1485–1510, 2009b.

- Yin, D., Ramchandran, K., and Bartlett, P. Rademacher complexity for adversarially robust generalization. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2019.
- Yoshida, Y. and Miyato, T. Spectral norm regularization for improving the generalizability of deep learning. *arXiv preprint arXiv:1705.10941*, 2017.
- Zhang, Y., Lee, J. D., and Jordan, M. I. ℓ_1 -regularized neural networks are improperly learnable in polynomial time. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2016.
- Zhang, Y., Liang, P., and Wainwright, M. J. Convexified convolutional neural networks. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2017.