
Certified Robustness to Label-Flipping Attacks via Randomized Smoothing

Elan Rosenfeld¹ Ezra Winston¹ Pradeep Ravikumar¹ J. Zico Kolter^{1,2}

Abstract

Machine learning algorithms are known to be susceptible to data poisoning attacks, where an adversary manipulates the training data to degrade performance of the resulting classifier. In this work, we present a unifying view of randomized smoothing over arbitrary functions, and we leverage this novel characterization to propose a new strategy for building classifiers that are *pointwise-certifiably robust* to general data poisoning attacks. As a specific instantiation, we utilize our framework to build linear classifiers that are robust to a strong variant of label flipping, where each test example is targeted independently. In other words, for each test point, our classifier includes a certification that its prediction would be the same had some number of training labels been changed adversarially. Randomized smoothing has previously been used to guarantee—with high probability—test-time robustness to adversarial manipulation of the input to a classifier; we derive a variant which provides a deterministic, analytical bound, sidestepping the probabilistic certificates that traditionally result from the sampling subprocedure. Further, we obtain these certified bounds with minimal additional runtime complexity over standard classification and no assumptions on the train or test distributions. We generalize our results to the multi-class case, providing the first multi-class classification algorithm that is certifiably robust to label-flipping attacks.

1. Introduction

Modern classifiers, despite their widespread empirical success, are known to be susceptible to adversarial attacks. In this paper, we are specifically concerned with so-called “data poisoning” attacks (formally, *causative* attacks [Bar-

reno et al. 2006; Papernot et al. 2018]), where the attacker manipulates some aspects of the training data in order to cause the learning algorithm to output a faulty classifier. Automated machine-learning systems which rely on large, user-generated datasets—e.g. email spam filters, product recommendation engines, and fake review detectors—are particularly susceptible to such attacks. For example, by maliciously flagging legitimate emails as spam and mislabeling spam as innocuous, an adversary can trick a spam filter into mistakenly letting through a particular email.

Data poisoning attacks in the literature include label-flipping attacks (Xiao et al., 2012), where the labels of a training set can be adversarially manipulated to decrease performance of the trained classifier; general data poisoning, where both the training inputs and labels can be manipulated (Steinhardt et al., 2017); and backdoor attacks (Chen et al., 2017; Tran et al., 2018), where the training set is corrupted so as to cause the classifier to deviate from its expected behavior only when triggered by a specific pattern. However, unlike the alternative test-time adversarial setting, where reasonably effective provable defenses exist, comparatively little work has been done on building classifiers that are certifiably robust to targeted data poisoning attacks.

In this work we propose a framework for building classifiers that are certifiably robust to a given class of data poisoning attacks, such as label-flipping or backdoor attacks. In particular, we propose what we refer to as a *pointwise certified defense*—this means that with each prediction, the classifier includes a certificate guaranteeing that its prediction would not be different had it been trained on adversarially manipulated data up to some “radius of perturbation” (a formal definition is presented in Section 3). We then demonstrate a specific instantiation of this protocol, constructing linear classifiers that are pointwise-certifiably robust to label-flipping attacks; i.e., each prediction is certified robust against a certain number of training label flips.

Prior works on certified defenses make statistical guarantees over the entire test distribution, but they make no guarantees as to the robustness of a prediction on any particular test point; thus, a determined adversary could still cause a specific test point to be misclassified. We therefore consider the threat of a worst-case adversary that can make a training set perturbation to target *each test point individually*. This

¹Carnegie Mellon University ²Bosch Center for AI. Correspondence to: Elan Rosenfeld <elan@cmu.edu>.

motivates a defense that can certify each of its individual predictions, as we present here. Compared to traditional robust classification, this framework is superior for a task such as determining who receives a coveted resource (a loan, parole, etc.), as it provides a guarantee for each individual, rather than at the population level. This work represents the first such pointwise certified defense to *any type of data poisoning attack*; we expect significant advances can be made on both attacks and defenses within this threat model.

Our approach leverages randomized smoothing (Cohen et al., 2019), a technique that has previously been used to guarantee test-time robustness to adversarial manipulation of the input to a deep network. However, where prior uses of randomized smoothing randomize over the input to the classifier for test-time guarantees, we instead randomize over *the entire training procedure of the classifier*. Specifically, by randomizing over the labels during this training process, we obtain an overall classification pipeline that is certified to not change its prediction when some number of labels are adversarially manipulated in the training set. Previous applications of randomized smoothing perform sampling to provide probabilistic bounds, due to the intractability of integrating the decision regions of a deep network. We instead derive an analytical bound, providing truly guaranteed robustness. Although a naive implementation of this approach would not be tractable, we show how to obtain these certified bounds with minimal additional runtime complexity over standard classification, suffering only a linear cost in the number of training points.

A further distinction of our approach is that the applicability of our robustness guarantees do not rely upon stringent model assumptions or the quality of the features. Existing work on robust linear classification or regression provides certificates that only hold under specific model assumptions, e.g., recovering the best-fit linear coefficients, which is most useful when the data exhibit a linear relationship in the feature space. In contrast, our classifier makes *no assumptions* about the separability of the data or quality of the features; this means our certificates remain valid when applying our classifier to arbitrary features, which in practice allows us to leverage advances in unsupervised feature learning (Le, 2013; Chen et al., 2020) and transfer learning (Donahue et al., 2014). We apply our classifier to pre-trained and unsupervised deep features to demonstrate its feasibility for classification of highly non-linear data such as ImageNet.

We evaluate our proposed classifier on several benchmark datasets common to the data poisoning literature. On the Dogfish binary classification challenge from ImageNet, our classifier maintains 81.3% certified accuracy in the face of an adversary who could reduce an undefended classifier to less than 1%. Additional experiments on MNIST and CIFAR10 demonstrate our algorithm’s effectiveness for

multi-class classification. Moreover, our classifier maintains a reasonably competitive non-robust accuracy (e.g., 94.5% on MNIST 1/7 versus 99.1% for the undefended classifier).

2. Related Work

Data poisoning attacks A *data poisoning attack* (Muñoz González et al., 2017; Yang et al., 2017) is an attack where an adversary corrupts some portion of the training set or adds new inputs, with the goal of degrading the performance of the learned model. The adversary is assumed to have perfect knowledge of the learning algorithm, so security by *design*—as opposed to *obscurity*—is the only viable defense against such attacks. The adversary is also typically assumed to have access to the training set and, in some cases, the test set.

Previous work has investigated attacks and defenses for data poisoning attacks applied to feature selection (Xiao et al., 2015), SVMs (Biggio et al., 2011; Xiao et al., 2012), linear regression (Liu et al., 2017), and PCA (Rubinstein et al., 2009), to name a few. Some attacks can even achieve success with “clean-label” attacks, inserting adversarially perturbed, seemingly correctly labeled training examples that cause the classifier to perform poorly (Shafahi et al., 2018; Zhu et al., 2019). Interestingly, our defense can also be viewed as (the first) certified defense to such attacks: perturbing an image such that the resulting features no longer match the label is theoretically equivalent to changing the label such that it no longer matches the image’s features. For an overview of data poisoning attacks and defenses in machine learning, see Biggio et al. (2014).

Label-flipping attacks A *label-flipping attack* is a specific type of data poisoning attack where the adversary is restricted to changing the training labels. The classifier is then trained on the corrupted training set, with no knowledge of which labels have been tampered with. For example, an adversary could mislabel spam emails as innocuous, or flag real product reviews as fake.

Unlike random label noise, for which many robust learning algorithms have been successfully developed (Natarajan et al., 2013; Liu & Tao, 2016; Patrini et al., 2017), adversarial label-flipping attacks can be specifically targeted to exploit the structure of the learning algorithm, significantly degrading performance. Robustness to such attacks is therefore harder to achieve, both theoretically and empirically (Xiao et al., 2012; Biggio et al., 2011). A common defense technique is *sanitization*, whereby a defender attempts to identify and remove or relabel training points that may have had their labels corrupted (Paudice et al., 2019; Taheri et al., 2019). Unfortunately, recent work has demonstrated that this is often not enough against a sufficiently powerful adversary (Koh et al., 2018). Further, *no existing defenses*

provide pointwise guarantees regarding their robustness.

Certified defenses Existing works on certified defenses to adversarial data poisoning attacks typically focus on the regression case and provide broad statistical guarantees over the entire test distribution. A common approach to such certifications is to show that a particular algorithm recovers some close approximation to the best linear fit coefficients (Diakonikolas et al., 2019; Prasad et al., 2018; Shen & Sanghavi, 2019), or that the expected loss on the test distribution is bounded (Klivans et al., 2018; Chen & Paschalidis, 2018). These results generally rely on assumptions on the data distribution: some assume sparsity in the coefficients (Karmalkar & Price, 2018; Chen et al., 2013) or corruption vector (Bhatia et al., 2015); others require limited effects of outliers (Steinhardt et al., 2017). As mentioned above, all of these methods fail to provide guarantees for individual test points. Additionally, these statistical guarantees are not as meaningful when their model assumptions do not hold.

Randomized smoothing Since the discovery of adversarial examples (Szegedy et al., 2013; Goodfellow et al., 2015), the research community has been investigating techniques for increasing the adversarial robustness of complex models such as deep networks. After a series of heuristic defenses, followed by attacks breaking them (Athalye et al., 2018; Carlini & Wagner, 2017), focus began to shift towards the development of *provable* robustness.

One approach which has gained popularity in recent work is randomized smoothing. Rather than certifying the original classifier f , randomized smoothing defines a new classifier g whose prediction at an input x is the class assigned the most probability when x is perturbed with noise from some distribution μ and passed through f . That is, $g(x) = \arg \max_c \mathbb{P}_{\epsilon \sim \mu}(f(x + \epsilon) = c)$. This new classifier g is then certified as robust, ideally without sacrificing too much accuracy compared to f . The original formulation was presented by Lecuyer et al. (2018) and borrowed ideas from differential privacy. The above definition is due to Li et al. (2018) and was popularized by Cohen et al. (2019), who derived a tight robustness guarantee.

3. A General View of Randomized Smoothing

Our first contribution is a general viewpoint of randomized smoothing, unifying all existing applications of the framework. Under our notation, randomized smoothing constructs an operator $G(\mu, \phi)$ that maps a binary-valued¹ function $\phi : \mathcal{X} \rightarrow \{0, 1\}$ and a *smoothing measure* $\mu : \mathcal{X} \rightarrow \mathbb{R}_+$,

¹For simplicity, we present the methodology here with binary-valued functions, which will correspond eventually to binary classification problems. The extension to the multi-class setting requires additional notation, and thus is deferred to the appendix.

with $\int_{\mathcal{X}} \mu(x) dx = 1$, to the expected value of ϕ under μ (that is, $G(\mu, \phi)$ represents the “vote” of ϕ weighted by μ). For example, ϕ could be a binary image classifier and μ could be some small, random pixel noise applied to the to-be-classified image. We also define a “hard threshold” version $g(\mu, \phi)$ that returns the most probable output (the majority vote winner). Formally,

$$G(\mu, \phi) = \mathbb{E}_{x \sim \mu}[\phi(x)] = \int_{\mathcal{X}} \mu(x) \phi(x) dx,$$

$$g(\mu, \phi) = \mathbf{1}\{G(\mu, \phi) \geq 1/2\},$$

where $\mathbf{1}\{\cdot\}$ is the indicator function. Where it is clear from context, we will omit the arguments, writing simply G or g . Intuitively, for two similar measures μ, ρ , we would expect that for most ϕ , even though $G(\mu, \phi)$ and $G(\rho, \phi)$ may not be equal, the threshold function g should satisfy $g(\mu, \phi) = g(\rho, \phi)$. Further, the degree to which μ and ρ can differ while still preserving this property should increase as $G(\mu, \phi)$ approaches either 0 or 1, because this increases the “margin” with which the function ϕ is 0 or 1 respectively over the measure μ . More formally, we define a general randomized smoothing guarantee as follows:

Definition 1. Let $\mu : \mathcal{X} \rightarrow \mathbb{R}_+$ be a smoothing measure over \mathcal{X} , with $\int_{\mathcal{X}} \mu(x) dx = 1$.² Then a randomized smoothing robustness guarantee is a specification of a distance function over probability measures $d(\mu, \rho)$ and a function $f : [0, 1] \rightarrow \mathbb{R}_+$ such that for all $\phi : \mathcal{X} \rightarrow \{0, 1\}$,

$$g(\mu, \phi) = g(\rho, \phi) \text{ whenever } d(\mu, \rho) \leq f(G(\mu, \phi)). \quad (1)$$

Informally, (1) says that the majority vote winner of ϕ weighted by μ and ρ will be the same, so long as μ and ρ are “close enough” as a function of the margin with which the majority wins. We will sometimes use p in place of $G(\mu, \phi)$, representing the fraction of the vote that the majority class receives (analogous to p_A in Cohen et al. (2019)).

Instantiations of randomized smoothing This definition is rather abstract, so we highlight concrete examples of how it can be applied to achieve certified guarantees against adversarial attacks.

Example 1. The randomized smoothing guarantee of Cohen et al. (2019) uses the smoothing measures $\mu = \mathcal{N}(x_0, \sigma^2 I)$, a Gaussian around the point x_0 to be classified, and $\rho = \mathcal{N}(x_0 + \delta, \sigma^2 I)$, the same measure perturbed by δ . They prove that (1) holds for all classifiers ϕ if we define

$$d(\mu, \rho) = \frac{1}{\sigma} \|\delta\|_2 \equiv \sqrt{2\text{KL}(\mu \parallel \rho)}, \quad f(p) = |\Phi^{-1}(p)|,$$

²There is no theoretical reason to restrict μ to be a probability measure. While this and all previous works only consider probability measures, the framework we present here could easily be extended to allow for more general measures μ, ρ and functions ϕ .

where $\text{KL}(\cdot)$ denotes KL divergence and Φ^{-1} denotes the inverse CDF of the Gaussian distribution.

Although this work focused on the case of randomized smoothing of continuous data via Gaussian noise, this is by no means a requirement. Lee et al. (2019) consider an alternative approach for dealing with discrete variables.

Example 2. *The randomized smoothing guarantee of Lee et al. (2019) uses the factorized smoothing measure in d dimensions $\mu_{\alpha,K}(\mathbf{x}) = \prod_{i=1}^d \mu_{\alpha,K,i}(\mathbf{x}_i)$, defined with respect to parameters $\alpha \in [0, 1]$, $K \in \mathbb{N}$, and a base input $\mathbf{z} \in \{0, \dots, K\}^d$, where*

$$\mu_{\alpha,K,i}(\mathbf{x}_i) = \begin{cases} \alpha, & \text{if } \mathbf{x}_i = \mathbf{z}_i \\ \frac{1-\alpha}{K}, & \text{if } \mathbf{x}_i \in \{0, \dots, K\}, \mathbf{x}_i \neq \mathbf{z}_i, \end{cases}$$

with \mathbf{x}_i being the i^{th} dimension of \mathbf{x} . $\rho_{\alpha,K}$ is similarly defined for a perturbed input \mathbf{z}' . They guarantee that (1) holds if we define

$$d(\mu, \rho) = \|\mathbf{z}' - \mathbf{z}\|_0, \quad f(p) = \mathcal{F}_{\alpha,K,d}(\max(p, 1-p)). \quad (2)$$

In words, the smoothing distribution is such that each dimension is independently perturbed to one of the other K values uniformly at random with probability $1 - \alpha$. $\mathcal{F}_{\alpha,K,d}(p)$ is a combinatorial function defined as the maximum number of dimensions—out of d total—by which $\mu_{\alpha,K}$ and $\rho_{\alpha,K}$ can differ such that a set with measure p under $\mu_{\alpha,K}$ is guaranteed to have measure at least $\frac{1}{2}$ under $\rho_{\alpha,K}$. Lee et al. (2019) prove that this value depends only on $\|\mathbf{z}' - \mathbf{z}\|_0$.

Finally, Dvijotham et al. (2020) consider a more general form of randomized smoothing that doesn't require strict assumptions on the distributions but is still able to provide similar guarantees.

Example 3 (Generic bound). *Given any two smoothing distributions μ, ρ , we have the generic randomized smoothing robustness certificate, ensuring that (1) holds with definitions*

$$d(\mu, \rho) = \text{KL}(\rho \parallel \mu), \quad f(p) = -\frac{1}{2} \log(4p(1-p)). \quad (3)$$

Randomized smoothing in practice For deep classifiers, the expectation G cannot be computed exactly, so we must resort to Monte Carlo approximation. This is done by drawing samples from μ and using these to construct a high-probability bound on G for certification. More precisely, this bound should be a *lower* bound on G when the hard prediction $g = 1$ and an *upper* bound otherwise; this ensures in both cases that we under-certify the true robustness of the classifier g . The procedure is shown in Algorithm 2 in Appendix A. These estimates can then be plugged into a randomized smoothing robustness guarantee to provide a high probability certified robustness bound for the classifier.

4. Pointwise Data Poisoning Robustness

We now present the main contributions of this paper: we first describe a generic strategy for applying randomized smoothing to certify a prediction function against arbitrary classes of data poisoning attacks. We then propose a specific implementation of said strategy to certify a classifier against label-flipping attacks. We show how this approach can be made tractable using linear least-squares classification, and we use the Chernoff inequality to analytically bound the relevant probabilities for the randomized smoothing certificate. Notably, although we are employing a randomized approach, the final algorithm does not use any random sampling, but rather relies upon a convex optimization problem to compute the certified robustness.

General data poisoning robustness We begin by noting that in prior work, randomized smoothing was applied at test time with the function $\phi : \mathcal{X} \rightarrow \{0, 1\}$ being a (potentially deep) classifier that we wish to smooth. However, there is no requirement that the function ϕ be a classifier at all; the theory holds for any binary-valued function. Instead of treating ϕ as a trained classifier, we consider ϕ to be *an arbitrary learning algorithm* which takes as input a training dataset $\{x_i, y_i\}_{i=1}^n \in (\mathcal{X} \times \{0, 1\})^n$ and additional test points without corresponding labels, which we aim to predict.³ In other words, the combined goal of ϕ is to first train a classifier and then predict the label of the new example. Thus, we consider test time outputs to be a function of both the test time input and the training data that produced the classifier. This perspective allows us to reason about how changes to training data affect the classifier at test time, reminiscent of work on influence functions of deep neural networks (Koh & Liang, 2017; Yeh et al., 2018).

This immediately suggests our protocol for pointwise robustness to general data poisoning attacks: randomize over the elements of the input to which we desire certified robustness, rather than over the test-time input to be classified. For example, to induce robustness to backdoor attacks, we could randomly add noise to the training points and/or their labels. Analogous to previous applications of randomized smoothing, if the majority vote of the classifiers trained with these randomly perturbed inputs has a large margin, it will confer a degree of robustness within an appropriately-defined radius of adversarially perturbed training data (as defined in Equation 1).

Specific application to label-flipping robustness To demonstrate the effectiveness of our proposed strategy, we now present a specific implementation, providing an algo-

³Note that our algorithm does not require access to the test data to do the necessary precomputation. We present it here as such merely to give an intuitive idea of the procedure.

rithm for tractable linear classification which is pointwise-certifiably robust to label-flipping attacks. When applying randomized smoothing in this setting, we randomize over the labels in the training set as described above—a suitably large margin in the majority vote will therefore result in pointwise robustness to adversarial label flips. In this scenario, the adversarial “radius” is defined as number of labels on which two training sets differ.

To formalize this intuition, consider two different assignments of n training labels $Y_1, Y_2 \in \{0, 1\}^n$ which differ on precisely r labels. Let μ (resp. ρ) be the distribution resulting from independently flipping each of the labels in Y_1 (resp. Y_2) with probability q . It is clear that as r increases, $\text{KL}(\mu \parallel \rho)$ will also increase. In fact, it is simple to show (see Appendix B.3 for derivation) that the exact KL divergence between these two distributions is

$$\text{KL}(\mu \parallel \rho) = \text{KL}(\rho \parallel \mu) = r(1 - 2q) \log \left(\frac{1 - q}{q} \right). \quad (4)$$

Plugging in the robustness guarantee (3), we have that $g(\mu, \phi) = g(\rho, \phi)$ so long as

$$r \leq \frac{\log(4p(1 - p))}{2(1 - 2q) \log \left(\frac{q}{1 - q} \right)}, \quad (5)$$

where $p = G(\mu, \phi)$. This implies that for any test point, as long as (5) is satisfied, g ’s prediction (the majority vote weighted by the smoothing distribution) will not change if an adversary corrupts the training set from Y_1 to Y_2 , or indeed to any other training set that differs on at most r labels. We can tune the noise hyperparameter q to achieve the largest possible upper bound in (5); more noise will likely decrease the margin of the majority vote p , but it will also decrease the divergence.

Computing a tight bound This approach has a simple closed form, but the bound is not tight. We can derive a tight bound via a combinatorial approach as in Lee et al. (2019). By precomputing the quantities $\mathcal{F}_{1-q, 1, n}^{-1}(r)$ from Equation (2) for each r , we can simply compare p to each of these and thereby certify robustness to the highest possible number of label flips. This precomputation can be expensive, but it provides a significantly tighter robustness guarantee, certifying approximately twice as many label flips for a given bound on G (See Figure E.3 in the Appendix).

4.1. Efficient implementation via least squares classifiers

There may appear to be one major impracticality of the algorithm proposed in the previous section, if considered naively: treating the function ϕ as an entire training-plus-single-prediction process would require that we train multiple classifiers, over multiple random draws of the labels y ,

Algorithm 1 Randomized smoothing for label-flipping robustness

Input: feature mapping $h : \mathbb{R}^d \rightarrow \mathbb{R}^k$; noise parameter q ; regularization parameter λ ; training set $\{(x_i, y_i) \in \mathbb{R}^d \times \{0, 1\}\}_{i=1}^n$ (with potentially adversarial labels); additional inputs to predict $\{x_j \in \mathbb{R}^d\}_{j=1}^m$.

1. Pre-compute matrix \mathbf{M} ,

$$\mathbf{M} = \mathbf{X} (\mathbf{X}^T \mathbf{X} + \lambda \mathbf{I})^{-1}$$

where $\mathbf{X} \equiv h(x_{1:n})$.

for $j = 1, \dots, m$ **do**

1. Compute vector $\alpha^j = \mathbf{M}h(x_j)^T$.

2. Compute optimal Chernoff parameter t via Newton’s method

$$t^* = \arg \min_t \left\{ t/2 + \sum_{i:y_i=1} \log \left(q + (1 - q)e^{-t\alpha^i} \right) + \sum_{i:y_i=0} \log \left((1 - q) + qe^{-t\alpha^i} \right) \right\}$$

and let $p^* = \max(1 - B_{|t^*|}, 1/2)$ where $B_{|t^*|}$ is the Chernoff bound (6) evaluated at $|t^*|$.

Output: Prediction $\hat{y}_j = \mathbf{1} \{t^* \geq 0\}$ and certification that prediction will remain constant for up to r training label flips, where

$$r = \left\lfloor \frac{\log(4p^*(1 - p^*))}{2(1 - 2q) \log \left(\frac{q}{1 - q} \right)} \right\rfloor.$$

end for

all to make a prediction on a single example. In this section, we describe a sequence of tools we employ to restrict the architecture and training process in a manner that drastically reduces this cost, bringing it in line with the traditional cost of standard classification. The full procedure, with all the parts described below, can be found in Algorithm 1.

Linear least-squares classification The fundamental simplification we make in this work is to restrict the “training” of the classifier ϕ to be done via the solution of a least-squares problem. Given the training set $\{x_i, y_i\}_{i=1}^n$, we assume that there exists some feature mapping $h : \mathbb{R}^d \rightarrow \mathbb{R}^k$ (where $k < n$). If existing linear features are not available, this could instead consist of deep features learned from a similar task—the transferability of such features is well documented (Donahue et al., 2014; Bo et al., 2010; Yosinski et al., 2014)—or features could be learned in an unsupervised fashion on $x_{1:n}$ (learning the features from poisoned labels could degrade performance). Given this

feature mapping, let $\mathbf{X} = h(x_{1:n}) \in \mathbb{R}^{n \times k}$ be the training point features and let $\mathbf{y} = y_{1:n} \in \{0, 1\}^n$ be the labels. Our training process consists of finding the least-squares fit to the training data, i.e., we find parameters $\hat{\beta} \in \mathbb{R}^k$ via the normal equation $\hat{\beta} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{y}$ and then we make a prediction on the new example via the linear function $h(x_{n+1})\hat{\beta}$. Although it may seem odd to fit a classification task with least-squares loss, binary classification with linear regression is equivalent to Fisher’s linear discriminant (Mika, 2003) and works quite well in practice.

The real advantage of the least-squares approach is that it reduces the prediction to a linear function of \mathbf{y} , and thus randomizing over the labels is straightforward. Specifically, letting

$$\alpha = \mathbf{X} (\mathbf{X}^T \mathbf{X})^{-1} h(x_{n+1})^T,$$

the prediction $h(x_{n+1})\hat{\beta}$ can be equivalently given by $\alpha^T \mathbf{y}$ (this is effectively the kernel representation of the linear classifier). Thus, we can simply compute α one time and then randomly sample many different sets of labels in order to build a standard randomized smoothing bound. Further, we can pre-compute just the $\mathbf{X} (\mathbf{X}^T \mathbf{X})^{-1}$ term and reuse it for each test point.

ℓ_2 regularization for better conditioning It is unlikely to be the case that the training points are well-behaved for linear classification in the feature space. To address this, we instead solve an ℓ_2 regularized version of least-squares. This is a common tool for solving systems with ill-conditioned or random design matrices (Hsu et al., 2014; Suggala et al., 2018). Luckily, there still exists a pre-computable closed-form solution to this problem, whereby we instead solve

$$\alpha = \mathbf{X} (\mathbf{X}^T \mathbf{X} + \lambda \mathbf{I})^{-1} h(x_{n+1})^T.$$

The other parts of our algorithm remain unchanged. Following results in Suggala et al. (2018), we set the regularization parameter $\lambda = (1 + q) \frac{\hat{\sigma}^2 k}{2n} \kappa(\mathbf{X}^T \mathbf{X})$ for all our experiments, where $\hat{\sigma}^2 = \frac{\|\mathbf{y} - \mathbf{X}\hat{\beta}_{OLS}\|_2^2}{n-k}$ is an estimate of the variance (Dicker, 2014) and $\kappa(\cdot)$ is the condition number.

Efficient tail bounds via the Chernoff inequality Even more compelling, due to the linear structure of this prediction, we can forego a sampling-based approach entirely and directly bound the tail probabilities using Chernoff bounds. Because the underlying binary prediction function ϕ will output the label 1 for the test point whenever $\alpha^T \mathbf{y} \geq 1/2$ and 0 otherwise, we can derive an analytical upper bound on the probability that g predicts one label or the other via the Chernoff bound. By upper bounding the probability of the *opposite* prediction, we simultaneously derive a lower bound on p which can be plugged in to (5) to determine the classifier’s robustness. Concretely, we can upper bound the

probability that the classifier outputs the label 0 by

$$\begin{aligned} P(\alpha^T \mathbf{y} \leq 1/2) &\leq \min_{t>0} \left\{ e^{t/2} \prod_{i=1}^n E[e^{-t\alpha_i y_i}] \right\} \\ &= \min_{t>0} \left\{ e^{t/2} \prod_{i=1}^n q e^{-t\alpha_i(1-y_i)} + (1-q)e^{-t\alpha_i y_i} \right\}. \end{aligned} \quad (6)$$

Conversely, the probability that the classifier outputs the label 1 is upper bounded by (6) but evaluated at $-t$. Thus, we can solve the minimization problem unconstrained over t , and then let the sign of t dictate which label to predict and the value of t determine the bound. The objective (6) is log-convex in t and can be easily solved by Newton’s method. Note that in some cases, neither Chernoff upper bound will be less than $1/2$, meaning we cannot determine the true value of g . In these cases, we simply define the classifier’s prediction to be determined by the sign of t . While we can’t guarantee that this classification will match the true majority vote, our algorithm will certify a robustness to 0 flips, so the guarantee is still valid. We avoid abstaining so as to assess our classifier’s non-robust accuracy.

The key property we emphasize is that, unlike previous randomized smoothing applications, the final algorithm involves *no randomness whatsoever*. Instead, the probabilities are bounded directly via the Chernoff bound, without any need for Monte Carlo approximation. Thus, the method is able to generate *truly certifiable* robust predictions using approximately the same complexity as traditional predictions.

5. Experiments

Following Koh & Liang (2017) and Steinhardt et al. (2017), we perform experiments on MNIST 1/7, the IMDB review sentiment dataset (Maas et al., 2011), and the Dogfish binary classification challenge taken from ImageNet. We run additional experiments on multi-class MNIST and CIFAR10. For each dataset and each noise level q we report the *certified test set accuracy* at r training label flips. That is, for each possible number of flips r , we plot the fraction of the test set that was both correctly classified and certified to not change under at least r flips.

As mentioned, our classifier suffers an additional linear cost in the number of training points due to the kernel representation $\alpha^T \mathbf{y}$. For most datasets there was no discernible difference in the time required to certify an input via our technique versus neural network classification. For larger training sets such as CIFAR10, especially when doing pairwise comparisons for the multi-class case, the algorithm is embarrassingly parallel; this parallelism brings runtime back in line with standard classification.

For binary classification, one could technically achieve a certified accuracy of 50% at $r = \infty$ (or 10% for MNIST

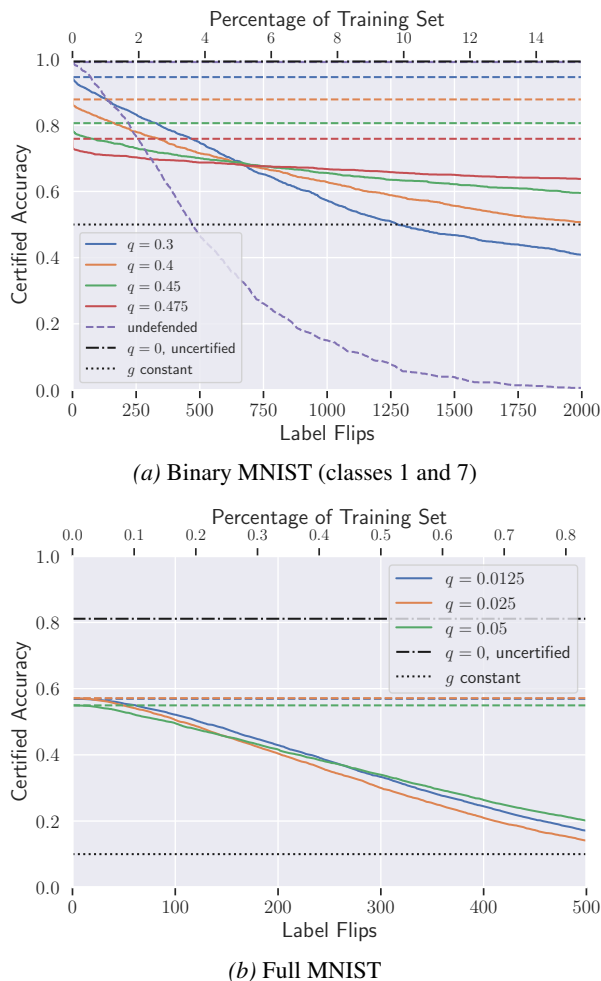


Figure 1. MNIST 1/7 ($n = 13007$, top) and full MNIST ($n = 60000$, bottom) test set certified accuracy to adversarial label flips as q is varied. The bottom axis represents the number of adversarial label flips to which *each individual prediction* is robust, while the top axis is the same value expressed as a percentage of the training set size. The solid lines represent certified accuracy; dashed lines of the same color are the overall non-robust accuracy of each classifier. The black dotted line is the (infinitely robust) performance of a constant classifier, while the black dash-dot line is the (uncertified) performance of our classifier with no label noise.

or CIFAR10) by letting g be constant—a constant classifier would be infinitely robust. Though not a very meaningful baseline, we include the accuracy of such a classifier in our plots (black dotted line) as a reference. We also evaluated our classifier with $q = 0$ (black dash-dot line); this cannot certify robustness, but it indicates the quality of the features.

To properly justify the need for such certified defenses, and to get a sense of the scale of our certifications, we generated label-flipping attacks against the undefended binary MNIST and Dogfish models. Following previous work, the undefended models were implemented as convolutional neural

networks, trained on the clean data, with all but the top layer frozen—this is equivalent to multinomial logistic regression on the learned features. For each test point we recorded how many flips were required to change the network’s prediction. This number serves as an upper bound for the robustness of the network on that test point, but we note that our attacks were quite rudimentary and could almost certainly be improved upon to tighten this upper bound. Appendix C.1 contains the details of our attack implementations. Finally, we implemented attacks on our own defense to derive an empirical upper bound and found that it reasonably tracks our lower bound. Plots and details of this attack can be found in Appendix C.2.

In all plots, the solid lines represent certified accuracy (except for the undefended classifier, which is an upper bound), while the dashed lines of the same color are the overall non-robust accuracy of each classifier.

Results on MNIST The MNIST 1/7 dataset (LeCun et al., 1998) consists of just the classes 1 and 7, totalling 13,007 training points and 2,163 test points. We trained a simple convolutional neural network on the other eight MNIST digits to learn a 50-dimensional feature embedding and then calculated Chernoff bounds for G as described in Section 4.1. Figure 1a displays the certified accuracy on the test set for varying probabilities q . As in prior work on randomized smoothing, the noise parameter q balances a trade-off; as q increases, the required margin $|G - \frac{1}{2}|$ to certify a given number of flips decreases. On the other hand, this results in more noisy training labels, which reduces the margin and therefore results in lower robustness and often lower accuracy. Figure 1b depicts the certified accuracy for the full MNIST test set—see Appendix B for derivations of the bounds and optimization algorithm in the multi-class case. In addition to this being a significantly more difficult classification task, our classifier could not rely on features learned from other handwritten digits; instead, we extracted the top 30 components with ICA (Hyvarinen, 1999) independently of the labels. Despite the lack of fine-tuned features, our algorithm still achieves significant certified accuracy under a large number of adversarial label flips.

See Figure E.1 in the Appendix for the effect of ℓ_2 regularization for the binary case. At a moderate cost to non-robust accuracy, the regularization results in substantially higher certified accuracy at almost all radii. We observed that regularization did not make a large difference for the multi-class case, possibly due to the inaccuracy of the residual term in the noise estimate.

Results on CIFAR10 To further demonstrate the effectiveness of our classifier with unsupervised features, we used SimCLR (Chen et al., 2020) to learn unsupervised features for CIFAR10. We used PCA to reduce the features to

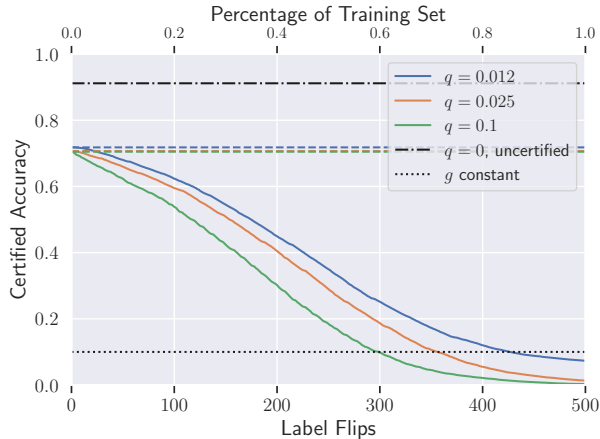


Figure 2. CIFAR10 ($n = 50000$) test set certified accuracy to adversarial label flips as q is varied.

128 dimensions to reduce overfitting. Figure 2 shows the results: our classifier with $q = 0.12$ achieves 50% certified accuracy up to 175 labels flips (recall there are ten classes, not two) and decays gracefully. Further, the classifier maintains better than random chance certified accuracy up to 427 label flips, which is approximately 1% of the training set.

Because the “votes” are changed by flipping so few labels, high values of q reduce the models’ predictions to almost pure chance—this means we are unable to achieve the margins necessary to certify a large number of flips. We therefore found that smaller levels of noise achieved higher certified test accuracy. This suggests that the more susceptible the original, non-robust classifier is to label flips, the lower q should be set for the corresponding randomized classifier.

For much smaller values of q , slight differences did not decrease the non-robust accuracy—they did however have a large effect on certified robustness. This indicates that the sign of t^* is relatively stable, but the margin of G is much less so. This same pattern was observed with the IMDB and Dogfish datasets. We used a high-precision arithmetic library (Johansson et al., 2013) to achieve the necessary lower bounds, but the precision required for non-vacuous bounds grew extremely fast for $q < 10^{-4}$; optimizing (6) quickly became too computationally expensive.

Results on Dogfish The Dogfish dataset contains images from the ImageNet dog and fish synsets, 900 training points and 300 test points from each. We trained a ResNet-50 (He et al., 2016) on the standard ImageNet training set but removed all images labeled dog or fish. Our pre-trained network therefore learned meaningful image features but had no features specific to either class. We again used PCA to reduce the feature space dimensionality. Figure 3 displays the results of our poisoning attack along with our certified

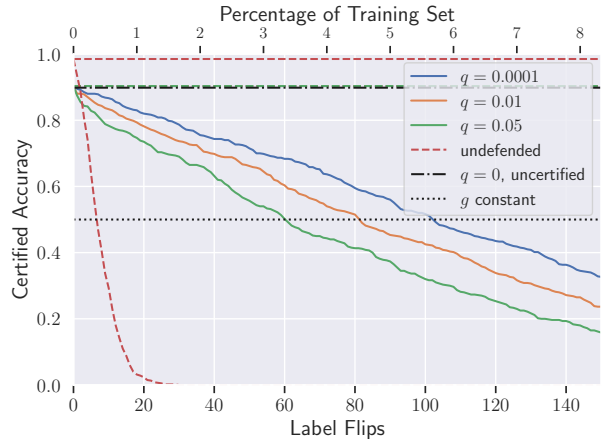


Figure 3. Dogfish ($n = 1800$) test set certified accuracy to adversarial label flips as q is varied.

defense. Under the undefended model, more than 99% of the test points can be successfully attacked with no more than 23 label flips, whereas our model with $q = 10^{-4}$ can certifiably correctly classify 81.3% of the test points under the same threat model. It would take more than four times as many flips—more than 5% of the training set—for *each test point individually* to reduce our classifier to less than 50% certified accuracy.

Here we observe the same pattern, where reducing q does not have a large effect on non-robust accuracy but does increase robustness significantly. This provides further evidence for the hypothesis that more complex datasets/classifiers are more susceptible to attacks and should be smoothed with less label noise.

Figure E.2 in the Appendix shows our classifier’s performance with unsupervised features. Because Dogfish is such a small dataset ($n = 1800$), deep unsupervised feature learning techniques were not feasible—we instead learned overcomplete features on 16x16 image patches using RICA (Le, 2013).

Results on IMDB Figure 4 plots the result of our randomized smoothing procedure on the IMDB review sentiment dataset. This dataset contains 25,000 training examples and 25,000 test examples, evenly split between “positive” and “negative”. To extract the features we applied the Google News pre-trained Word2Vec to all the words in each review and averaged them. This feature embedding is considerably noisier than that of an image dataset, as most of the words in a review are irrelevant to sentiment classification. Indeed, Steinhardt et al. (2017) also found that the IMDB dataset was much more susceptible to adversarial corruption than images when using bag-of-words features. Consistent with this, we found smaller levels of noise resulted in larger cer-

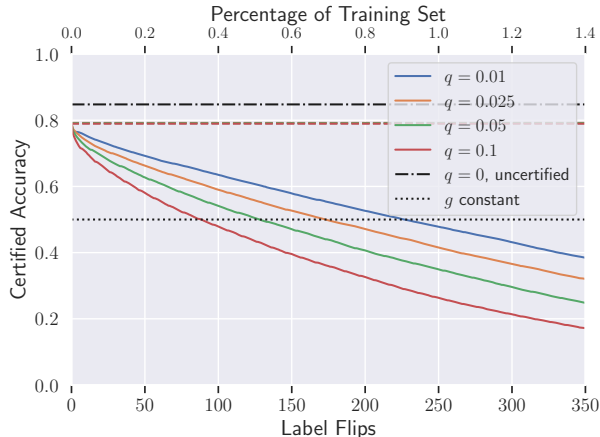


Figure 4. IMDB Review Sentiment ($n = 25000$) test set certified accuracy. The non-robust accuracy slightly decreases as q increases; for $q = 0.01$ the non-robust accuracy is 79.11%, while for $q = 0.1$ it is 78.96%.

tified accuracy. We expect significant improvements could be made with a more refined choice of feature embedding.

6. Conclusion

In this work we presented a unifying view of randomized smoothing, which borrows from the literature of differential privacy in order to provide black-box certificates of robustness. Based on the observation that this framework is applicable more broadly than just defenses to adversarial examples, we used it to derive a framework for certified defenses against arbitrary data poisoning attacks—we dub such defenses “pointwise” because they provide a certificate of robustness for each test point.

We next implemented this protocol as a specific classifier which is robust to a strong class of label-flipping attacks, where an adversary can flip labels to target each test point individually. This contrasts with previous data poisoning defenses which have typically only considered an adversary who wishes to degrade the classifier’s accuracy on the test distribution as a whole. Finally, we offered a tractable algorithm for evaluating this classifier which, despite being rooted in randomization, can be computed with no Monte Carlo sampling whatsoever, resulting in a truly certifiably robust classifier. This work represents the first classification algorithm that is pointwise-certifiably robust to *any type* of data poisoning attack; we anticipate many possible new directions within this framework.

A particular strength of this framework is when we specifically care about robustly classifying each input individually. Compared to traditional robust classification, this technique is superior for determining who receives a coveted resource

(a loan, parole, etc.) or for making some other sensitive classification, as it provides a guarantee for each individual. Other works only ensure that they correctly classify some fraction p of the population, which is often not acceptable as that still leaves the $1 - p$ fraction who could be misclassified, with no indication of which ones belong to the test set.

There are several avenues for improvements to this line of work. Most immediately, our protocol could be implemented with other types of smoothing distributions applied to the training data, such as randomizing over the input data or features, to derive specific algorithms that are pointwise-certifiably robust to other types of data poisoning attacks. Additionally, the method for learning the input features in an unsupervised, semi-supervised, or self-supervised manner could be improved. Finally, we hope that our defense to this threat model will inspire the development of more powerful (e.g., pointwise) train-time attacks, against which future defenses can be evaluated.

Acknowledgements

We thank Guang-He Lee for sharing his code on discrete robustness certificates, Adarsh Prasad for several helpful sources on robust linear regression and classification, and Alnur Ali for pointing us to Dicker (2014) for help choosing the appropriate regularization term. We are grateful to Jeremy Cohen and Adarsh Prasad for helpful discussions and reviewing drafts of this work. E. R. and P. R. acknowledge the support of DARPA via HR00112020006.

References

- Athalye, A., Carlini, N., and Wagner, D. Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018*, July 2018.
- Barreno, M., Nelson, B., Sears, R., Joseph, A. D., and Tygar, J. D. Can machine learning be secure? In *Proceedings of the 2006 ACM Symposium on Information, Computer and Communications Security*, pp. 16–25, New York, NY, USA, 2006.
- Bhatia, K., Jain, P., and Kar, P. Robust regression via hard thresholding. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1*, pp. 721–729, Cambridge, MA, USA, 2015.
- Biggio, B., Nelson, B., and Laskov, P. Support vector machines under adversarial label noise. In *Proceedings of the Asian Conference on Machine Learning*, volume 20, pp. 97–112, South Garden Hotels and Resorts, Taoyuan, Taiwan, 14–15 Nov 2011.

- Biggio, B., Fumera, G., and Roli, F. Security evaluation of pattern classifiers under attack. *IEEE Transactions on Knowledge and Data Engineering*, 26(4):984–996, April 2014.
- Bo, L., Ren, X., and Fox, D. Kernel descriptors for visual recognition. In *Advances in Neural Information Processing Systems 23*, pp. 244–252. 2010.
- Carlini, N. and Wagner, D. Adversarial examples are not easily detected: Bypassing ten detection methods. In *Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security*, pp. 3–14, New York, NY, USA, 2017.
- Chen, R. and Paschalidis, I. C. A robust learning approach for regression models based on distributionally robust optimization. *J. Mach. Learn. Res.*, 19(1):517–564, January 2018. ISSN 1532-4435.
- Chen, T., Kornblith, S., Norouzi, M., and Hinton, G. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709*, 2020.
- Chen, X., Liu, C., Li, B., Lu, K., and Song, D. Targeted backdoor attacks on deep learning systems using data poisoning. *arXiv preprint arXiv:1712.05526*, 2017.
- Chen, Y., Caramanis, C., and Mannor, S. Robust sparse regression under adversarial corruption. In *Proceedings of the 30th International Conference on Machine Learning*, volume 28, pp. 774–782, Atlanta, Georgia, USA, 17–19 Jun 2013.
- Cohen, J., Rosenfeld, E., and Kolter, J. Z. Certified adversarial robustness via randomized smoothing. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97, pp. 1310–1320, Long Beach, California, USA, 09–15 Jun 2019.
- Diakonikolas, I., Kong, W., and Stewart, A. Efficient algorithms and lower bounds for robust linear regression. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 2745–2754, Philadelphia, PA, USA, 2019.
- Dicker, L. H. Variance estimation in high-dimensional linear models. *Biometrika*, 101(2):269–284, 03 2014.
- Donahue, J., Jia, Y., Vinyals, O., Hoffman, J., Zhang, N., Tzeng, E., and Darrell, T. Decaf: A deep convolutional activation feature for generic visual recognition. In *Proceedings of the 31st International Conference on Machine Learning*, volume 32, pp. 647–655, Beijing, China, 22–24 Jun 2014.
- Dvijotham, K., Hayes, J., Balle, B., Kolter, Z., Qin, C., Gyorgy, A., Xiao, K., Gowal, S., and Kohli, P. A framework for robustness certification of smoothed classifiers using f-divergences. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020, Conference Track Proceedings*, 2020.
- Goodfellow, I. J., Shlens, J., and Szegedy, C. Explaining and harnessing adversarial examples. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- Hsu, D., Kakade, S. M., and Zhang, T. Random design analysis of ridge regression. *Found. Comput. Math.*, 14(3):569–600, June 2014.
- Hyvarinen, A. Fast and robust fixed-point algorithms for independent component analysis. *IEEE Transactions on Neural Networks*, 10(3):626–634, May 1999.
- Johansson, F. et al. *mpmath: a Python library for arbitrary-precision floating-point arithmetic (version 0.18)*, December 2013. <http://mpmath.org/>.
- Karmalkar, S. and Price, E. Compressed sensing with adversarial sparse noise via ℓ_1 regression. *arXiv preprint arXiv:1809.08055*, 2018.
- Klivans, A., Kothari, P. K., and Meka, R. Efficient algorithms for outlier-robust regression. *arXiv preprint arXiv:1803.03241*, 2018.
- Koh, P. W. and Liang, P. Understanding black-box predictions via influence functions. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, pp. 1885–1894, 2017.
- Koh, P. W., Steinhardt, J., and Liang, P. Stronger data poisoning attacks break data sanitization defenses. *arXiv preprint arXiv:1811.00741*, 2018.
- Le, Q. V. Building high-level features using large scale unsupervised learning. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 8595–8598, May 2013.
- LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, Nov 1998.
- Lecuyer, M., Atlidakis, V., Geambasu, R., Hsu, D., and Jana, S. Certified robustness to adversarial examples with differential privacy. *arXiv preprint arXiv:1802.03471*, 2018.

- Lee, G.-H., Yuan, Y., Chang, S., and Jaakkola, T. S. Tight certificates of adversarial robustness for randomly smoothed classifiers. In *Advances in Neural Information Processing Systems 31*, 2019.
- Li, B., Chen, C., Wang, W., and Carin, L. Certified adversarial robustness with additive gaussian noise. *arXiv preprint arXiv:1809.03113*, 2018.
- Liu, C., Li, B., Vorobeychik, Y., and Oprea, A. Robust linear regression against training data poisoning. In *AISeC 2017 - Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security, co-located with CCS 2017*, 11 2017.
- Liu, T. and Tao, D. Classification with noisy labels by importance reweighting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(3):447–461, March 2016.
- Maas, A. L., Daly, R. E., Pham, P. T., Huang, D., Ng, A. Y., and Potts, C. Learning word vectors for sentiment analysis. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies - Volume 1*, pp. 142–150, Stroudsburg, PA, USA, 2011.
- Mika, S. *Kernel Fisher Discriminants*. Doctoral thesis, Technische Universitt Berlin, Fakultt IV - Elektrotechnik und Informatik, Berlin, 2003. URL <http://dx.doi.org/10.14279/depositonce-575>.
- Muñoz González, L., Biggio, B., Demontis, A., Paudice, A., Wongrassamee, V., Lupu, E. C., and Roli, F. Towards poisoning of deep learning algorithms with back-gradient optimization. In *Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security*, pp. 27–38, New York, NY, USA, 2017.
- Natarajan, N., Dhillon, I. S., Ravikumar, P. K., and Tewari, A. Learning with noisy labels. In *Advances in Neural Information Processing Systems 26*, pp. 1196–1204. 2013.
- Papernot, N., McDaniel, P., Sinha, A., and Wellman, M. P. Sok: Security and privacy in machine learning. In *2018 IEEE European Symposium on Security and Privacy (EuroS P)*, pp. 399–414, April 2018.
- Patrini, G., Rozza, A., Krishna Menon, A., Nock, R., and Qu, L. Making deep neural networks robust to label noise: A loss correction approach. pp. 2233–2241, 07 2017.
- Paudice, A., Muñoz-González, L., and Lupu, E. C. Label sanitization against label flipping poisoning attacks. In *ECML PKDD 2018 Workshops*, pp. 5–15, Cham, 2019.
- Prasad, A., Suggala, A. S., Balakrishnan, S., and Ravikumar, P. Robust estimation via robust gradient estimation. *arXiv preprint arXiv:1802.06485*, 2018.
- Rubinstein, B. I., Nelson, B., Huang, L., Joseph, A. D., Lau, S.-h., Rao, S., Taft, N., and Tygar, J. D. Antidote: Understanding and defending against poisoning of anomaly detectors. In *Proceedings of the 9th ACM SIGCOMM Conference on Internet Measurement*, pp. 1–14, New York, NY, USA, 2009.
- Shafahi, A., Huang, W. R., Najibi, M., Suci, O., Studer, C., Dumitras, T., and Goldstein, T. Poison frogs! targeted clean-label poisoning attacks on neural networks. In *Advances in Neural Information Processing Systems 31*, pp. 6103–6113. 2018.
- Shen, Y. and Sanghavi, S. Learning with bad training data via iterative trimmed loss minimization. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97, pp. 5739–5748, Long Beach, California, USA, 09–15 Jun 2019.
- Steinhardt, J., Koh, P. W., and Liang, P. Certified defenses for data poisoning attacks. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pp. 3520–3532, USA, 2017.
- Suggala, A., Prasad, A., and Ravikumar, P. K. Connecting optimization and regularization paths. In *Advances in Neural Information Processing Systems 31*, pp. 10608–10619. 2018.
- 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.
- Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., and Wojna, Z. Rethinking the inception architecture for computer vision. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- Taheri, R., Javidan, R., Shojafar, M., Pooranian, Z., Miri, A., and Conti, M. On defending against label flipping attacks on malware detection systems. *arXiv preprint arXiv:1908.04473*, 2019.
- Tran, B., Li, J., and Madry, A. Spectral signatures in back-door attacks. In *Advances in Neural Information Processing Systems 31*, pp. 8000–8010. 2018.
- Xiao, H., Xiao, H., and Eckert, C. Adversarial label flips attack on support vector machines. In *Proceedings of the 20th European Conference on Artificial Intelligence*, pp. 870–875, Amsterdam, The Netherlands, The Netherlands, 2012.
- Xiao, H., Biggio, B., Brown, G., Fumera, G., Eckert, C., and Roli, F. Is feature selection secure against training data poisoning? In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37, pp. 1689–1698, Lille, France, 07–09 Jul 2015.

- Yang, C., Wu, Q., Li, H., and Chen, Y. Generative poisoning attack method against neural networks. *arXiv preprint arXiv:1703.01340*, 2017.
- Yeh, C.-K., Kim, J., Yen, I. E.-H., and Ravikumar, P. K. Representer point selection for explaining deep neural networks. In *Advances in Neural Information Processing Systems 31*, pp. 9291–9301. 2018.
- Yosinski, J., Clune, J., Bengio, Y., and Lipson, H. How transferable are features in deep neural networks? In *Advances in Neural Information Processing Systems 27*, pp. 3320–3328. 2014.
- Zhu, C., Huang, W. R., Li, H., Taylor, G., Studer, C., and Goldstein, T. Transferable clean-label poisoning attacks on deep neural nets. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97, pp. 7614–7623, Long Beach, California, USA, 09–15 Jun 2019.