

# Get Global Guarantees: On the Probabilistic Nature of Perturbation Robustness

Wenchuan Mu

wenchuan\_mu@sutd.edu.sg

Singapore University of Technology and Design  
Singapore

Kwan Hui Lim

kwanhui@acm.org

Singapore University of Technology and Design  
Singapore

## ABSTRACT

Robustness is a critical requirement for deploying machine learning models in safety-sensitive domains, where even imperceptible input perturbations can lead to hazardous outcomes. However, existing robustness assessment techniques prior to deployment often face a trade-off between computational feasibility and measurement precision, limiting their effectiveness in practice. To address these limitations, we provide a systematic comparative study of prevailing robustness definitions and their corresponding evaluation methodologies. Building on this analysis, we propose *tower robustness*, which is a novel and practical concept setting out from a global perspective. Further, we provide upper and lower bounds of tower robustness, based on hypothesis testing, for quantitative evaluation, enabling more rigorous and efficient pre-deployment assessments. Through empirical investigation, we demonstrate that our approach provides reliable robustness assessments. These findings advance the systematic understanding of robustness and contribute a practical framework for enhancing the safety of machine learning models in safety-critical applications.

## CCS CONCEPTS

• **Theory of computation** → *Adversarial learning*; • **Mathematics of computing** → *Hypothesis testing and confidence interval computation*.

### ACM Reference Format:

Wenchuan Mu and Kwan Hui Lim. 2025. Get Global Guarantees: On the Probabilistic Nature of Perturbation Robustness. In *Proceedings of the 34th ACM International Conference on Information and Knowledge Management (CIKM '25)*, November 10–14, 2025, Seoul, Republic of Korea. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3746252.3761039>

## 1 INTRODUCTION

Machine learning has transformed numerous fields, enabling breakthroughs in image recognition [9, 12], natural language processing [10, 43], and autonomous systems [11]. Particularly in safety-critical domains, such as self-driving vehicles [17, 41], medical diagnosis [30, 38], and industrial control systems [2], the reliability and robustness of neural networks are paramount. In these contexts,

even minor perturbations in input data can lead to severe consequences, highlighting the need for rigorous robustness evaluation prior to deployment.

Model robustness evaluation is well-studied but remains challenging. Formal verification is one established approach, but it often suffers from incomplete problem formulations [14], failing to capture the full range of real-world perturbations. In addition, verification techniques are typically computationally expensive, rendering them impractical for large-scale models [46]. These limitations call for evaluation methods that are both more comprehensive and computationally feasible.

To address these gaps, we explore probabilistic robustness assessment, which aims to estimate a model's failure probability under adversarial perturbations. While existing probabilistic methods frequently rely on approximations, such approximations risk overlooking critical adversarial instances [1], potentially overestimating model robustness.

**Contribution 1:** In this work, we propose a new concept *tower robustness*, a probabilistic measure of robustness from a global perspective (Section 3). Following this, we come up with its upper and lower bounds (TEB-U/TEB-L) that leverage exact binomial tests to provide statistical guarantees (Section 4). Tower robustness enables precise quantification of failure probabilities, yielding an interpretable and actionable robustness measure. By grounding robustness assessment in exact statistical principles, our approach offers a more realistic characterisation of model behaviour in real-world deployments.

**Contribution 2:** To support adoption, we release an open-source implementation at <https://github.com/cestwc/tower-robustness>, providing researchers and practitioners with a user-friendly and statistically rigorous evaluation tool. We further demonstrate the practicality of our approach through empirical evaluations on several neural network models.

## 2 PRELIMINARIES

This section discusses the prerequisite knowledge. We first introduce the notations of each quantity. After that, we review the essentials of statistical learning, e.g., data sampling, objective functions, and neural networks.

### 2.1 Notations

We use the simple letter styles as given in Table 1 to distinguish different types of mathematical objects. Note that the letter  $P$  is reserved for probability, and  $E$  for expectation. The notation (e.g., for variables) is generally insensitive to whether an object is a scalar or a vector. Subscripts or superscripts in variable names are simply part of the name, not as indices. In fact, we do not index into vector

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

CIKM '25, November 10–14, 2025, Seoul, Republic of Korea

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-2040-6/2025/11

<https://doi.org/10.1145/3746252.3761039>

**Table 1: Mathematical object letter styles (e.g., using letter A)**

A variable or variable function	$a$
A random variable or random variable function	$\mathbf{a}$
A set	$\mathbb{A}$
A distribution	$\mathcal{A}$
A function whose value is 1 if $a$ is true, 0 otherwise	$\mathbf{1}_a$

variables or sets throughout this manuscript. Also, lowercase Greek letters are always used for scalars.

## 2.2 Data, Distribution, and Samples

In the following, we briefly review the probability and statistics background of this work. Let  $D$  denote a joint probability distribution for random variables  $\mathbf{x}$  and  $\mathbf{y}$ , such that for some positive integer  $n$ , a sample from  $D$  could look like

$$\{(x_1, y_1), (x_2, y_2), \dots, (x_n, y_n)\} \quad (1)$$

where  $n$  is the sample size. In practice, we are usually given the sample. For example, MNIST [19] dataset could be seen as a sample drawn from the distribution of handwritten digits<sup>1</sup>. If we further look into the set of MNIST training images and testing images, they can be considered two samples drawn from this distribution where  $n = 60,000$  for the training set and  $n = 10,000$  for the testing set.

Let  $\mathbb{X}$  be the outcome space of  $\mathbf{x}$ , and  $\mathbb{Y}$  be the outcome space of  $\mathbf{y}$ . Intuitively, in the MNIST context,  $\mathbb{X}$  would be the set containing all  $28 \times 28$  pixel greyscale images, and  $\mathbb{Y}$  would be a set of integers from 0 to 9.

Often, we are interested in the expected value of some function  $f(\mathbf{x}, \mathbf{y})$ , denoted by  $E(f(\mathbf{x}, \mathbf{y}))$  and computed as follows.

$$E[f(\mathbf{x}, \mathbf{y})] = \int_{\mathbb{Y}} \int_{\mathbb{X}} f(x, y) P(\mathbf{x} = x, \mathbf{y} = y) dx dy \quad (2)$$

In practice, if a sample is large enough, the law of large numbers ensures that the sample average converges to the true expected value, i.e., as  $n \rightarrow \infty$

$$E[f(\mathbf{x}, \mathbf{y})] \approx \frac{1}{n} \sum_{i=1}^n f(x_i, y_i). \quad (3)$$

## 2.3 Learning and Neural Models

Let us denote a function  $h : \mathbb{X} \rightarrow \mathbb{Y}$ . Essentially,  $\mathbb{X}$  is considered an input feature space, and  $\mathbb{Y}$  is considered a label space, while a model function  $h$  predicts a label from a given feature. A good model makes correct predictions, i.e., we want to maximise

$$E[1_{h(\mathbf{x})=\mathbf{y}}]. \quad (4)$$

Its practical measurement, however, often relies on a finite set. According to approximation (3), we can rewrite expression (4) as expression (5). Indeed, expression (5) denotes the commonly used test accuracy metric when the set is the testing set.

$$\frac{1}{n} \sum_{i=1}^n 1_{h(x_i)=y_i} \quad (5)$$

<sup>1</sup>To view MNIST as an unbiased sample, we might as well say it is from a distribution of American employees' and high school students' handwritten digits.

The process of actively optimising the model is referred to as learning (or training). Learning is complicated and may be done empirically through various forms [6, 8, 29, 34]. This is primarily because we only have access to a finite training set of data that theoretically does not suffice to calculate the objective function, i.e. expression (4) or (5). Consequently, a common practice is to define an optimisation problem (typically the objective function) that is fully measurable from the training set. For instance, the empirical risk minimisation (ERM [34]) method is defined as follows, which is one of the most straightforward learning strategies.

*Definition 2.1 (Empirical risk minimisation).* Define a density function  $g : \mathbb{X} \times \mathbb{Y} \rightarrow [0, 1]$  such that  $\forall x \in \mathbb{X}, 1 = \sum_{y \in \mathbb{Y}} g(x, y)$ . Supposedly, this density reflects how likely a prediction on  $x$  will be  $y$ , i.e.,  $g(x, y) \equiv P(\hat{\mathbf{y}} = y \mid \mathbf{x} = x)$ , where  $\hat{\mathbf{y}}$  denotes the random variable of prediction. In this way, expression (4) can be rewritten as  $E_{\mathbf{x}}[\langle P(\mathbf{y} \mid \mathbf{x}), P(\hat{\mathbf{y}} \mid \mathbf{x}) \rangle]$ , where  $\langle \cdot, \cdot \rangle$  stands for inner product. Then, instead of directly maximising this inner product, ERM minimises the Kullback-Leibler divergence (KL divergence) between these two conditional probabilities. When it comes to the training sample with size  $n$ , the term to be minimised is as follows.

$$-\frac{1}{n} \sum_{i=1}^n \log g(x_i, y_i) \quad (6)$$

Obtaining  $h$  from  $g$  can be straightforward, e.g., one common practice is to take  $h(x) := \arg \max_y g(x, y)$ .

Another key aspect of learning, which is orthogonal to the choice of the objective function, is the model architecture. While the optimisation process (e.g., minimising empirical risk) is general, the way we represent the model function  $h$  can vary significantly, and some representations may perform better than others in practice [12]. A common and effective approach is to use a parametrised model, such as a neural network. In a simplified view, a neural network consists of (1) a fixed structural design that defines the topology of the network (e.g., number of layers, connections), and (2) variable parameters such as numerical weights and biases. During training, the optimisation algorithm searches over the parameter space, while the structure remains unchanged.

*Definition 2.2 (Neural Network Models).* If  $h$  is a neural model, we can re-express it as  $h(x) := \arg \max_y g(x, \mathbf{w}, y)$ , where  $g : \mathbb{X} \times \mathbb{W} \times \mathbb{Y} \rightarrow [0, 1]$  and  $\forall x \in \mathbb{X}, \forall \mathbf{w} \in \mathbb{W}, 1 = \sum_{y \in \mathbb{Y}} g(x, \mathbf{w}, y)$ . Set  $\mathbb{W}$  is the entire search space of parameters. Generally, the (sub)gradients  $\partial g / \partial \mathbf{w}$  and  $\partial g / \partial x$  are both available.

## 2.4 Adversarial Examples

Maximising the probability of making correct predictions is generally a sound and intuitively reasonable objective. Yet, recent research has raised important questions about whether this objective is sufficient on its own. One concern arises from the context of adversarial examples. Inputs that have been slightly perturbed in ways often imperceptible to humans, yet cause a model's prediction to flip from correct to incorrect. This could be a security problem in practice, e.g., a subtly altered stop sign may be misread by an autonomous vehicle as a speed limit sign.

Although it remains an open question whether the adversarial example (AE) issue can be solved inherently by optimising expression (4), empirical evidence shows a troubling trade-off. Some models with high test accuracy perform poorly under adversarial perturbations, while others with better adversarial example performance often suffer from lower test accuracy. Here, we describe the formal definitions of adversarial examples for subsequent analysis.

**Definition 2.3 (Adversarial Example).** For real number  $p \geq 1$ ,  $\epsilon \in [0, 1]$ , given any input  $x$  from space  $\mathbb{X}$ , we have that  $x'$  a (measurable)  $(p, \epsilon)$ -neighbour of  $x$  if and only if  $x' \in \mathbb{X} \wedge \|x - x'\|_p \leq \epsilon$ . For any model  $h : \mathbb{X} \rightarrow \mathbb{Y}$ , for any paired  $(x, y) \in \mathbb{X} \times \mathbb{Y}$ , we have that a  $(p, \epsilon)$ -neighbour  $x'$  of  $x$  is also a  $(p, \epsilon)$ -AE of  $x$  (regarding this  $h$  and this  $y$ ) if and only if  $h(x) = y \wedge h(x') \neq y$ . Additionally, given  $(x, y)$ , let us define a condition  $c_{AE}(x, y; p, \epsilon)$  or simply  $c_{AE}$  as there exists some  $(p, \epsilon)$ -AE of  $(x, y)$ . Formally,

$$c_{AE}(x, y; p, \epsilon, h) := (h(x) = y) \wedge \exists x'. x' \in \mathbb{X} \wedge \|x - x'\|_p \leq \epsilon \wedge h(x') \neq y \quad (7)$$

Definition 2.3 is often the most intuitive and widely accepted way to define adversarial examples, as it aligns well with our understanding of small, imperceptible perturbations causing misclassification. However, it introduces a subtle issue, *i.e.*, what if the original input is misclassified, yet there exists a neighbour that is correctly classified? Does that imply we should treat the original input as the “adversarial” point from the perspective of the neighbour?

Despite this ambiguity, Definition 2.3 is not flawed. We adopt this definition as a foundation and demonstrate that different contexts of robustness are given in Table 2.

### 3 TOWER ROBUSTNESS

In this section, we present tower robustness, which is a sound probabilistic quantity of the perturbation robustness concept. First, we describe the essence of robustness and revisit various observation strategies for robustness to see why the concept of robustness has not been captured well enough. After that, we propose tower robustness based on the tower law and formally show why this is a reasonable measure of robustness.

#### 3.1 Robustness, What Are We Actually Observing?

While avoiding adversarial examples is commonly considered the core of robustness<sup>2</sup>, the quantitative definition of robustness remains an open problem. Empirically, there exist a few metrics to measure robustness as demonstrated in Table 2, *e.g.*, adversarial attack failure rate, but there has never been an agreement on which one truly represents robustness.

In the following, we aim to set this issue in a theoretical framework. We start from this often overlooked question, *i.e.*, what exactly do we aim to maximise when we ask for greater robustness?

Note that in this study, we do not focus on specific optimisation strategies, *e.g.*, adversarial training. Instead, we focus on the objective functions on the evaluation side, *e.g.*, test accuracy for

classification. In other words, we would like to measure the robustness of models and would not concern ourselves with how these models have been built. Let us now walk through how each robustness metric is currently defined *de facto* in practice.

**Adversarial robustness and its bounds.** The overall adversarial robustness of a model can be expressed as  $\mathbb{E}[1_{c_{AR}}]$ , or equivalently,  $P(c_{AR})$ . Here,  $P(c_{AR})$  represents the probability that no adversarial examples exist in the neighbourhood of a randomly chosen input. Computing this probability directly is intractable, because even simple high-dimensional data like an MNIST [19] image can have over  $10^{200}$  neighbours.

A practical approach is to adversarially attack each testing case by heuristically searching for adversarial examples within a limited number of steps, rather than exhaustively checking all neighbours. Under this method, if no adversarial example is found within the search steps (*i.e.*,  $\neg c_{AA}$ ), we can infer  $\neg c_{AR}$ . However, if the attack times out (*i.e.*,  $c_{AA}$  holds), we gain no information about  $c_{AR}$ . In other words,  $P(c_{AA})$  serves as an upper bound for  $P(c_{AR})$ . The issue is that  $P(c_{AA})$  depends on the specific attack method used, which may not be fully independent of the model itself, especially in cases like adversarial training where the attack procedure influences the model. As a result, ranking models based on  $P(c_{AA})$  might not reflect their ranking in terms of the adversarial robustness,  $P(c_{AR})$ .

An alternative strategy is to define a condition  $c_{DC}$  such that  $c_{DC} \rightarrow c_{AR}$ , providing a lower bound for  $P(c_{AR})$ . Unlike  $c_{AA}$ , the condition  $c_{DC}$  does not rely on any attack process, but it is known to be heavily influenced by the model choices. When  $c_{DC}$  does not hold, it remains possible that  $c_{AR}$  still holds, but we have no way to confirm this.

It might be imaginably possible that we can simultaneously lift both the lower and upper bounds on  $P(c_{AR})$  to effectively improve robustness. In practice, however, a substantial gap often remains between these bounds, and it is rare to find models that make meaningful progress on both fronts at once. As a result, these quantities offer limited insight into true robustness.

**Probabilistic robustness.** Probabilistic robustness was originally proposed to address the problem that high-accuracy models usually lack robustness guarantees [31, 46]. To this end, the focused event is modified from  $c_{AR}$  to  $c_{PR}$ , *i.e.*, from whether there are no adversarial examples to whether there are few enough adversarial examples. However, this also makes the definition less interpretable, *e.g.*, it would read like maximising “the probability of misclassification upon input perturbation is less than a tolerance level” [31]. Furthermore, the choice of this tolerance level ( $\kappa$  as in Table 2) introduces a new parameter that not only requires justification but aggravates the complicatedness of robustness evaluation.

Worse, this  $P(c_{PR})$  form itself is not verifiable through a finite number of samples, and the practical solution to provide such a probabilistic robustness guarantee must read like maximising “the probability of that the probability that the probability of misclassification upon input perturbation is less than a tolerance level is greater than or equal to a significance level” [46]. While this is a reasonable choice out of practicality, it introduces yet another new parameter ( $\alpha$  as in Table 2), which makes the evaluation even more confusing. These quantities risk becoming abstractions that obscure more than they reveal: approximations that gesture toward

<sup>2</sup>There are other loosely related concepts in the name of robustness, *e.g.*, generalisation robustness or compression robustness, which are orthogonal to our focus.

**Table 2: Various contexts of robustness. In any case, we are given a pair of input and label  $(x, y)$  from  $\mathbb{X} \times \mathbb{Y}$ , and  $h$  is the model.**

Quantity	Event formula	Computable	Meaning
Adversarial robustness	$c_{AR}(x, y; p, \epsilon, h) \doteq (h(x) = y) \wedge \neg \exists x'. x' \in \mathbb{X} \wedge \ x - x'\ _p \leq \epsilon \wedge h(x') \neq y$	Intractable	$c_{AR}$ is true if there does not exist any adversarial example to this input-label pair, and false if there do.
Adversarial attack failure rate	$c_{AF}(x, y; p, \epsilon, h, f_{\text{attack}}) \doteq \neg \exists x'. x' \in f_{\text{attack}}(x, y, p, \epsilon, h) \wedge \ x - x'\ _p \leq \epsilon \wedge h(x') \neq y$ , where $f_{\text{attack}}(x, y, p, \epsilon, h) \subset \mathbb{X}$	Yes	$c_{AF}$ being false means we successfully find at least one adversarial example. $c_{AF}$ being true means no adversarial example has been found yet, but we have no idea whether one will be found.
Adversarial accuracy	$c_{AA}(x, y; p, \epsilon, h, f_{\text{attack}}) \doteq (h(x) = y) \wedge \neg \exists x'. x' \in f_{\text{attack}}(x, y, p, \epsilon, h) \wedge \ x - x'\ _p \leq \epsilon \wedge h(x') \neq y$	Yes	$c_{AA}$ indicates the prediction on $x$ is correct, and no adversarial example has been found yet (again, we do not know if one will be found). Otherwise, we get $\vdash \neg c_{AA}$ .
Deterministic certified robust accuracy	To find $c_{DC}(x, y; p, \epsilon, h)$ such that $\vdash c_{DC}(x, y; p, \epsilon, h) \rightarrow c_{AR}(x, y; p, \epsilon, h)$ .	Yes, but not for all $x$	Try to let a sufficient condition of $c_{AR}$ hold. Then, if $c_{DC}$ holds, we know that no adversarial exists around the input. But if $c_{DC}$ does not hold, we do not get any implication.
Probabilistic robustness	$c_{PR}(x, y; p, \epsilon, h, \kappa) \doteq P_{\text{Uniform}, \ x' - x\ _p \leq \epsilon} (h(x') \neq y) \leq \kappa$ , for $0 < \kappa < 1/2$	Unverifiable from samples	$c_{PR}$ hold if and only if the proportion of adversarial examples among neighbours is capped by a real number $\kappa$ . Predicting the original $x$ correctly is not emphasised.
Probabilistic robust accuracy	$c_{PRA}(x, y; p, \epsilon, h, \kappa, \alpha) \doteq P(P_{\text{Uniform}, \ x' - x\ _p \leq \epsilon} (h(x') \neq y) > \kappa) > \alpha$ , for $0 < \kappa < 1/2, 0 < \alpha < 1$ .	Yes	When $c_{PRA}$ holds, the probability of $c_{PR} \neq 1$ is small, such that it is capped at a real value $\alpha$ . When $c_{PRA}$ does not hold, it just means that we are not informed by this inequality.

the underlying phenomenon without fully capturing it. We argue that pursuing a clearer and more principled definition of robustness remains an essential direction.

*Robustness might not be just about rates.* As demonstrated in Table 2, nearly every event involves a complex combination of multiple factors. Although it is straightforward to compute an average across all test cases, this does not necessarily mean that the events are well-defined in a probabilistic sense.

A potential blind spot is that robustness might not have to be represented as a rate (like all in Table 2). Intuitively, the measurable quantity of robustness could be understood as an estimator of a deeper property. Let us take a look at the accuracy counterpart, *i.e.*, accuracy is not just a rate, but an unbiased estimator of a probability, with clear statistical properties like expectation, variance, and concentration bounds (*e.g.*, Hoeffding’s inequality). In the next step, we will find this property for robustness and present our definition.

### 3.2 The Probability of Misprediction Given Any Neighbourhood

In the following, we introduce tower robustness, a property that underpins measurable notions of robustness in a meaningful and structured way. First, we state its formal definition. Then, using the tower law (also known as the law of total expectation), we show how this property encapsulates robustness and connects to measurable quantities. Finally, we offer an interpretation from a convolutional perspective to aid in understanding its implications.

*Definition 3.1 (tower robustness).* For tower robustness, we ask, given a random neighbourhood, what is the probability of making correct predictions? Formally, we are given a joint distribution of  $\mathbf{x}, \mathbf{y}$ , where  $\mathbf{x}$  has support  $\mathbb{X}$  and  $\mathbf{y}$  has support  $\mathbb{Y}$ . For any model  $h : \mathbb{X} \rightarrow \mathbb{Y}$ , its  $(p, \epsilon)$ -tower robustness is defined as

$$P_{\|x' - x\|_p \leq \epsilon} (h(x') = y), \quad (8)$$

or equivalently

$$E_{\|x' - x\|_p \leq \epsilon} [1_{h(x')=y}]. \quad (9)$$

Intuitively, the tower robustness (TR) of a model under a given distribution and specified  $(p, \epsilon)$  would be a unique scalar. A higher TR suggests a more robust model at the given setting. Note that TR itself, as defined in Definition 3.1, is computable using finite

samples, but may be estimated with a variance. Thus, we propose computable bounds of TR in Section 4.

*3.2.1 Derivation of tower robustness and link to other robustness quantities.* To understand the rationale of tower robustness, we start from the sense of deterministic certified robust accuracy (DC accuracy). According to Table 2, DC accuracy can be written as  $\sum_i^n 1_{c_{DC,i}}/n$ , representing the proportion of test cases for which a deterministic certification is achieved. While DC accuracy is typically viewed as unrelated to probabilistic behaviour unless explicitly noted, we will demonstrate that it is, in fact, fundamentally probabilistic in nature.

*The probabilistic nature in deterministic robustness.* The term probabilistic robustness, when used in a deterministic context like a recent systematisation of knowledge [23], is often ambiguous. This ambiguity arises from that there are two ways of choosing random variables. The first way uses  $\mathbf{x}, \mathbf{y}$  as random variable. That is, for each given pair of data, the condition  $c_{DC}$  is a (deterministic) function value, and we aim to get

$$P(c_{DC}(\mathbf{x}, \mathbf{y}, p, \epsilon, h)). \quad (10)$$

This is often overlooked, because we often resort to its unbiased maximum likelihood estimation (MLE), *i.e.*, DC accuracy. The second way uses  $c_{DC}$  as random variable. This particularly works when given  $\mathbf{x}, \mathbf{y}$ , we are not able to determine  $c_{DC}$  and have to guess whether it is true, *e.g.*, when robustness verification is not complete but already acquires some knowledge  $c_{\text{known}}$  about this testing case. This works within the neighbourhood of individual given  $\mathbf{x}, \mathbf{y}$ , rather than across samples. Formally, this could be expressed as  $P(c_{DC} \mid c_{\text{known}}(\mathbf{x}, \mathbf{y}, p, \epsilon, h))$ .

In this study, we leverage the first way of choosing random variables, *i.e.*, expression (10), to help understanding tower robustness. Essentially, expression (10) represents the probability of achieving deterministic certification for a random input-label pair from the distribution. Furthermore, expression (10) is a sound lower bound of  $P(c_{AR}(\mathbf{x}, \mathbf{y}, p, \epsilon, h))$ , *i.e.*, adversarial robustness.

For a long time,  $P(c_{AR}(\mathbf{x}, \mathbf{y}, p, \epsilon, h))$  has been seen well capturing robustness. However, this measure comes with notable limitations. First, its value is theoretically proven to be bounded within a low range, primarily because excluding all adversarial examples is an inherently strict requirement. As a result, relying on

$P(c_{AR}(\mathbf{x}, \mathbf{y}, p, \epsilon, h))$  to quantify robustness can be misleading. Intuitively, it is akin to signing up for an extremely difficult chess match and then scoring poorly. Merely attempting a hard challenge does not demonstrate mastery. More importantly, when encountering a new random input-label pair, one cannot determine in advance whether adversarial examples are present. This uncertainty means that robustness, in this context, remains probabilistic. Worse yet, due to the typically low probability value, such a probabilistic estimate provides little practical confidence. To intuitively understand this, let us look at the following game example.

#### Example: Expected Payout from a Game

Imagine you play a carnival game where you are randomly assigned to either Machine A (40% chance) or Machine B (60% chance). Machine A pays \$10 with a 50% chance and \$0 otherwise, while Machine B pays \$5 with an 80% chance and \$0 otherwise. To find the expected payout, we calculate the expected value for each machine: for Machine A, it is  $0.5 \times 10 + 0.5 \times 0 = 5$ , and for Machine B it is  $0.8 \times 5 + 0.2 \times 0 = 4$ . Then, using the tower law, the overall expected payout is  $0.4 \times 5 + 0.6 \times 4 = 2 + 2.4 = \$4.40$ .

In an adversarial robustness scenario, consider a system like “Machine A” that, with a 50% chance, shows zero adversarial examples, and with the remaining 50%, shows an unknown proportion of adversarial examples, potentially many, say less than  $T\%$ . The expected adversarial example payout is then  $0.5 \times 0 + 0.5 \times T = 5T\%$ . Like in the carnival game, it is the  $5T\%$  that matters practically. With this idea, let us review the tower law formally.

**Definition 3.2 (tower law).** Let  $\mathbf{b}$  be an integrable random variable, i.e.,  $E[|\mathbf{b}|] < \infty$ . As  $E(\mathbf{b} \mid \mathbf{a})$  is a random variable and a function of  $\mathbf{a}$ , we can take its expected value, and it can be shown that

$$E[E(\mathbf{b} \mid \mathbf{a})] = E[\mathbf{b}] \quad (11)$$

**THEOREM 3.3.** We are given a joint distribution of  $\mathbf{x}, \mathbf{y}$ , where  $\mathbf{x}$  has support  $\mathbb{X}$  and  $\mathbf{y}$  has support  $\mathbb{Y}$ . For any model  $h : \mathbb{X} \rightarrow \mathbb{Y}$ , its  $(p, \epsilon)$ -tower robustness can be equivalently written as

$$E \left[ P \left( h(\mathbf{x}') = \mathbf{y} \mid \|\mathbf{x}' - \mathbf{x}\|_p \leq \epsilon \right) \right] \quad (12)$$

**PROOF.** We apply the tower law to expression (9). First, we can see that the premise holds, i.e.,  $1_{(*)}$  returns either 0 or 1, which makes it an integrable function. Then, we confirm that for random variables  $\mathbf{x}'$  and  $\mathbf{y}$ ,  $\|\mathbf{x}' - \mathbf{x}\|_p \leq \epsilon$  and  $1_{h(\mathbf{x}')=\mathbf{y}}$  are indeed a function of  $(\mathbf{x}, \mathbf{y})$ . Then, we get

$$E \left[ E \left[ 1_{h(\mathbf{x}')=\mathbf{y}} \mid \|\mathbf{x}' - \mathbf{x}\|_p \leq \epsilon \right] \right]. \quad (13)$$

Since for any event  $c$ , we have  $P(c) = E[1_c]$ , we get our proof.  $\square$

Intuitively, Theorem 3.3 explains the tower robustness, i.e., we first randomly sample a neighbourhood (around an input-label pair) from the distribution, then we randomly pick a input-label pair from this neighbourhood, we ask, what will be the probability that we make a correct prediction on this picked pair? Next, we show the connection between the tower robustness and  $P(c_{DC}(\mathbf{x}, \mathbf{y}, p, \epsilon, h))$ .

**COROLLARY 3.4.** *DC accuracy is an unbiased estimator of a lower bound of tower robustness. Formally,*

$$P(c_{DC}(\mathbf{x}, \mathbf{y}, p, \epsilon, h)) \leq E \left[ P \left( h(\mathbf{x}') = \mathbf{y} \mid \|\mathbf{x}' - \mathbf{x}\|_p \leq \epsilon \right) \right] \quad (14)$$

**PROOF.** We have known that  $\sum_i^n 1_{c_{DC,i}}/n$  is the sample mean of  $P(c_{DC}(\mathbf{x}, \mathbf{y}, p, \epsilon, h))$ , such that our focus is to show the inequality above. It is sufficient if we can prove

$$\forall \mathbf{x}, \mathbf{y}, \neg \left( 1_{c_{DC}(\mathbf{x}, \mathbf{y}, p, \epsilon, h)} > P \left( h(\mathbf{x}') = \mathbf{y} \mid \|\mathbf{x}' - \mathbf{x}\|_p \leq \epsilon \right) \right). \quad (15)$$

Since  $1_{(*)}$  returns only 0 or 1, and probabilities are always non-negative, we need only consider the case where the probability is strictly less than 1. In this case, there must be a positive measure of adversarial examples, which implies that  $c_{DC}$  is false. This completes the proof.  $\square$

Intuitively, when trying to maximise the tower robustness, we can still use DC accuracy as a sound lower bound to provide guarantee. Yet, neither probabilistic robustness  $E[c_{PR}(\mathbf{x}, \mathbf{y}, p, \epsilon, h, \kappa)]$  nor probabilistic robust accuracy  $E[c_{PRA}(\mathbf{x}, \mathbf{y}, p, \epsilon, h, \kappa, \alpha)]$  are naturally lower bounds of tower robustness, i.e., they may overestimate.

**3.2.2 Tower robustness from a convolutional perspective.** We now draw a link from tower robustness to commonly adopted correctness, i.e., expression (4). The link between the robustness of a given distribution and the correctness of a convolved distribution has been studied recently [44]. Here, we formally extend this perspective to tower robustness.

As we are sampling  $\mathbf{x}, \mathbf{y}$  (to get its neighbourhood) from the given distribution and then sample within the neighbourhood, it could be understood as adding two random variables. The first random variable is  $\mathbf{x}$ , representing the input from the joint distribution. We denote the second random variable as  $\mathbf{t} \in \mathbb{X}$  and  $\|\mathbf{t}\|_p \leq \epsilon$ . Unsurprisingly, we can re-express  $\mathbf{x}' = \mathbf{x} + \mathbf{t}$ . Theoretically, sampling of  $\mathbf{x}$  and sampling of  $\mathbf{t}$  are independent. Therefore, the probability density function (PDF) of  $\mathbf{x}'$  can be expressed as

$$p_{\mathbf{x}'}(\mathbf{x}') = \int_{\mathbb{X}} p_{\mathbf{x}}(\mathbf{x}) p_{\mathbf{t}}(\mathbf{t} - \mathbf{x}) d\mathbf{x}. \quad (16)$$

Considering the joint distribution  $D$  of  $\mathbf{x}, \mathbf{y}$ , and that  $\mathbf{t}$  is independent of  $\mathbf{y}$  as well, we can also express the conditional distributions in convolution as follows.

$$p_{\mathbf{x}', \mathbf{y}}(\mathbf{x}', \mathbf{y}) = \int_{\mathbb{X}} p_{\mathbf{x}, \mathbf{y}}(\mathbf{x}, \mathbf{y}) p_{\mathbf{t}}(\mathbf{t} - \mathbf{x}) d\mathbf{x} \quad (17)$$

That is,  $(\mathbf{x}', \mathbf{y}) \sim D * p_{\mathbf{t}}$ . In this way, we can rewrite tower robustness from expression (9) to the following form.

$$E_{\mathbf{x}, \mathbf{y} \sim D * p_{\mathbf{t}}} [h(\mathbf{x}) = \mathbf{y}] \quad (18)$$

Note that *only* for expression (18) we use  $\mathbf{x}$  instead of  $\mathbf{x}'$  to denote the perturbed input because we assume the distribution has been changed so we no longer need to add perturbations again. As seen expression (18) is in a highly consistent form with expression (4). In other words, maximising robustness is indeed maximising some special kind of correctness, and this depends on how much the user/researcher would like to “melt” the original given distribution.

## 4 GLOBAL BOUNDS (LOWER AND UPPER) ON TOWER ROBUSTNESS

Tower robustness can be estimated from finite samples without bias, but we still need to control the variance of the estimators. To establish a stricter formal guarantee in practice, we introduce TEB-L, a computable tower robustness Lower bound through Tower(ed) Exact Binomial tests.

In the following, we present TEB-L in a bottom-up style. First, we describe how to test each individual input-label pair from a given distribution. Then, we show how we can combine the test results to provide guarantees.

### 4.1 Binomial Hypothesis Test With Exact Solution

In this step, we are given a pair of known  $x, y$ , and our task is to know whether the proportion of mispredictions in this  $(p, \epsilon)$ -neighbourhood is less than or equal to a real value  $\kappa$  where  $0 < \kappa < 1/2$ . Formally, we need to check if the following condition holds.

$$P(h(x') \neq y \mid \|x - x'\|_p \leq \epsilon) \leq \kappa \quad (19)$$

Observe that the inequality above is exactly  $c_{PR}(x, y; p, \epsilon, h, \kappa)$ , the event formula for probabilistic robustness [31]. There are a few potential solutions to this task, where the shared strategy is to treat  $h(x') \neq y$  as a random variable with a Bernoulli distribution.

From there, Robey et al. [31] use the maximum likelihood estimator from a random sample, which in the Bernoulli case is the sample mean (with sample size 100). This method is fast, but prone to overfitting and high variance, especially with small samples or when the true probability is near 0 or 1. When estimates approach  $\kappa$ , distinguishing their relation to  $\kappa$  becomes statistically unreliable, necessitating more precise estimation to ensure significance.

*Hypothesis testing.* We can use hypothesis testing to better estimate  $c_{PR}$ . To formulate a hypothesis testing problem, we first let  $p_\tau$  denote the true parameter  $p$  in our Bernoulli distribution. As such, we can state the null hypothesis

$$H_0 : p_\tau > \kappa \quad (20)$$

and the alternative hypothesis

$$H_1 : p_\tau \leq \kappa. \quad (21)$$

Then, we statistically determine whether there is enough evidence to reject the null hypothesis in favour of the alternative hypothesis.

For example, suppose we want to test the following one-sided hypothesis:  $H_0 : p_\tau > 0.01$ , and  $H_1 : p_\tau \leq 0.01$ . Assume that we take a sample of 30 binary outcomes and observe 2 successes (i.e., two misprediction cases out of 30). Then, we can determine if this 2 out of 30 is significant.

While hypothesis testing helps preserve statistical significance, not all tests are well-suited for this problem. For example, the Agresti-Coull confidence interval used by Zhang et al. [46] approximates a binomial test, and may not perform well when the sample size is very small or the true proportion  $p$  is close to zero or one, leading to incorrect inferences [1].

*Exact Binomial Test.* We compute the p-value as the cumulative probability as follows, under the binomial distribution with  $p_\tau = \kappa$ .

$$\text{p-value} = P(k > k \mid k \sim \text{Bin}(n, \kappa)) = \sum_{i=0}^k \binom{n}{i} \kappa^i (1 - \kappa)^{n-i} \quad (22)$$

Here,  $n$  represents the sample size, and  $k$  denotes the number of successes. Then, if the obtained p-value exceeds a significance level  $\alpha$ , e.g., 0.05, we fail to reject  $H_0$ . If the obtained p-value is less than or equal to the significance level  $\alpha$ , we can reject  $H_0$  and adopt  $H_1$ .

*Compare with Agresti-Coull Approximation.* The difference between a binomial test (with exact solution) and its Agresti-Coull approximation lies in the way of calculating the cumulative density function of a binomial distribution. Particularly, the Agresti-Coull approximation estimates  $p_\tau$  using

$$\tilde{p} = \frac{k + z_\alpha^2/2}{n + z_\alpha^2} \quad (23)$$

where  $k$  denotes the number of successes, and  $z_\alpha = \Phi^{-1}(1 - \frac{\alpha}{2})$  is the quantile of a standard normal distribution, e.g., for  $\alpha = 0.05$  at one side,  $z_{0.05} = 1.645$ . A lower bound of the confidence interval is then calculated as

$$\tilde{p} - z_\alpha \sqrt{\frac{\tilde{p}(1 - \tilde{p})}{n + z_\alpha^2}}. \quad (24)$$

If and only if this lower bound is greater than  $\kappa$ , then we can reject  $H_0$  and adopt  $H_1$ .

In the example above, where  $n = 30, k = 2$ , the exact binomial test gives a p-value of 0.9967, which fails to reject. In contrast, the Agresti-Coull confidence interval has a lower bound 0.0153, which exceeds 0.01 and thus leads to rejection of  $H_0$ . This discrepancy illustrates that the Agresti-Coull approximation can overstate evidence against the null in small samples and when the  $p_\tau$  is near 0 or 1, while the exact binomial test remains conservative.

*Both left- and right-tailed tests.* So far, we have been working on a left-tailed hypothesis test, i.e., the null hypothesis is  $p_\tau$  being greater than some value. Intuitively, a left-tailed test that rejects the null hypothesis suggests the probability of event  $a_2$  conditioned on event  $\neg a_1$  is less than  $\alpha$ , where

- Event  $a_1$ :  $c_{PR}(x, y; p, \epsilon, h, \kappa)$  is *indeed* true,
- Event  $a_2$ : Our evidence says that  $c_{PR}(x, y; p, \epsilon, h, \kappa)$  would be true.

Intuitively, the significance level  $\alpha$  also represents the failure rate of our conclusion.

Similar to a left-tailed test, we can meanwhile plan a right-tailed test. When a right-tailed test rejects the null hypothesis, it suggests, the probability of event  $\neg a_2$  conditioned on event  $a_1$  is less than  $\alpha$ .

For every input-label pair  $(x, y)$ , the hypothesis test will reject either side of the tail if there are sufficient samples (if not sufficient, we assume the worst case, i.e., reject the opposite side). The estimation of sample number is well-studied [28, 36], e.g., we could use the following formula

$$n = \left( \frac{z_{1-\alpha} \cdot \sqrt{\kappa(1-\kappa)} + z_\alpha \cdot \sqrt{\hat{p}_\tau(1-\hat{p}_\tau)}}{\kappa - \hat{p}_\tau} \right)^2 \quad (25)$$

Where  $\hat{p}_T$  is an empirically estimated Bernoulli parameter  $p$ . Therefore, we would obtain

$$P(\neg a_2 | a_1) < \alpha, \quad (26)$$

and similarly:

$$P(a_2 | \neg a_1) < \alpha. \quad (27)$$

## 4.2 True Probability Rather Than Observed Events

Next, we use the law of total probability, a special case of tower law, to eliminate the help values  $\kappa$  and  $\alpha$ . We first eliminate  $\alpha$ .

**4.2.1 Eliminating  $\alpha$ .** Here, our goal is to determine the probability of  $c_{PR}(x, y; p, \epsilon, h, \kappa)$  over the distribution. In other words, we want to obtain  $P(a_1)$ . Besides, what we already have is Inequality (26) and (27).

**LEMMA 4.1.** *We are given a joint distribution of  $x, y$ , where  $x$  has support  $\mathbb{X}$  and  $y$  has support  $\mathbb{Y}$ . For any model  $h : \mathbb{X} \rightarrow \mathbb{Y}$ , its  $(p, \epsilon)$ -probabilistic robustness has the following lower and upper bounds.*

$$\begin{aligned} P(c_{PR}(x, y; p, \epsilon, h, \kappa)) &> \frac{P(c_{PRA}(x, y; p, \epsilon, h, \kappa, \alpha)) - \alpha}{1 + \alpha} \\ P(c_{PR}(x, y; p, \epsilon, h, \kappa)) &< \frac{P(c_{PRA}(x, y; p, \epsilon, h, \kappa, \alpha))}{1 - \alpha} \end{aligned} \quad (28)$$

**PROOF.** From the law of total probability, it is obtained that

$$P(a_2) = P(a_2 | a_1)P(a_1) + P(a_2 | \neg a_1)P(\neg a_1). \quad (29)$$

If we further write  $P(\neg a_1) = 1 - P(a_1)$ , and  $P(a_2 | \neg a_1) = 1 - P(\neg a_2 | \neg a_1)$ , we eventually get

$$P(a_1) = \frac{P(a_2) - P(a_2 | \neg a_1)}{1 - P(\neg a_2 | a_1) + P(a_2 | \neg a_1)}. \quad (30)$$

Now that we know that  $0 < P(a_2 | \neg a_1), P(\neg a_2 | a_1) < \alpha$ , we can find the lower and upper limit of  $P(a_1)$  as

$$\frac{P(a_2) - \alpha}{1 + \alpha} < P(a_1) < \frac{P(a_2)}{1 - \alpha}. \quad (31)$$

Recall that the event  $a_1$  is essentially that  $c_{PR}(x, y; p, \epsilon, h, \kappa)$  holds. Also, event  $a_2$  is the observation. This completes the proof.  $\square$

Inequality (28) makes sense because the probabilistic robustness (although itself is not measurable) is still predominantly positively related to the observed probabilistic robust accuracy, and a smaller false positive rate ( $\alpha$ ) also makes these two quantities closer. Note that the usually measurable probabilistic robust accuracy is a maximum likelihood estimator of  $P(c_{PRA}(x, y; p, \epsilon, h, \kappa, \alpha))$ , as it is computing the sample mean given discrete  $\{(x_1, y_1), \dots\}$ .

**4.2.2 Eliminating  $\kappa$ .** Although  $\alpha$  has been eliminated, as noted in prior work [25],  $\kappa$  remains. These extra parameters add confusion and labour to robustness evaluation, e.g., fair comparison is difficult if different  $\kappa$  values are used. Next, we present how to eliminate  $\kappa$ .

In this task, we are given  $P(c_{PR}(x, y; p, \epsilon, h, \kappa))$  or its lower and upper bounds. We can equivalently expand this expression back to the following form using expression (19).

$$P\left(P\left(h(x') \neq y \mid \|x - x'\|_p \leq \epsilon\right) \leq \kappa\right) \quad (32)$$

In the following, we link this form back to tower robustness, through both Markov's Inequality and tower law.

**THEOREM 4.2.** *We are given a joint distribution of  $x, y$ , where  $x$  has support  $\mathbb{X}$  and  $y$  has support  $\mathbb{Y}$ . For any model  $h : \mathbb{X} \rightarrow \mathbb{Y}$ , its  $(p, \epsilon)$ -tower robustness has a computable upper bound as*

$$\frac{\kappa P(c_{PRA}(x, y; p, \epsilon, h, \kappa, \alpha))}{1 - \alpha} - \kappa + 1. \quad (33)$$

**PROOF.** According to Markov's Inequality, for a non-negative function  $f(b) \geq 0$  and a random variable  $b$ ,

$$P(f(b) \geq \kappa) \leq \frac{E[f(b)]}{\kappa}. \quad (34)$$

This gives us

$$P(f(b) < \kappa) \geq 1 - \frac{E[f(b)]}{\kappa}. \quad (35)$$

Observe that  $P(h(x') \neq y \mid \|x - x'\|_p \leq \epsilon)$  is a non-negative function of  $x$ . Thus, substituting it in we get the following form.

$$P\left(P_{\|x-x'\|_p \leq \epsilon}(h(x') \neq y) \leq \kappa\right) \geq 1 - \frac{E[P_{\|x-x'\|_p \leq \epsilon}(h(x') \neq y)]}{\kappa} \quad (36)$$

Shifting the sides, we get

$$\begin{aligned} E\left[P_{\|x-x'\|_p \leq \epsilon}(h(x') \neq y)\right] \\ \geq \kappa \left(1 - P\left(P_{\|x-x'\|_p \leq \epsilon}(h(x') \neq y) \leq \kappa\right)\right) \end{aligned} \quad (37)$$

Multiplying -1 and then adding 1 to both sides, we get

$$\begin{aligned} E\left[P_{\|x-x'\|_p \leq \epsilon}(h(x') = y)\right] \\ \leq 1 - \kappa \left(1 - P\left(P_{\|x-x'\|_p \leq \epsilon}(h(x') \neq y) \leq \kappa\right)\right) \\ \leq 1 - \kappa + \kappa P\left(P_{\|x-x'\|_p \leq \epsilon}(h(x') \neq y) \leq \kappa\right) \end{aligned} \quad (38)$$

We have got the lower and upper limits of the third term on the right-hand side of this inequality. Substituting that in, we get

$$E\left[P_{\|x-x'\|_p \leq \epsilon}(h(x') = y)\right] < \frac{\kappa P(c_{PRA}(x, y; p, \epsilon, h, \kappa, \alpha))}{1 - \alpha} - \kappa + 1. \quad (39)$$

This provides an upper bound on tower robustness of the given model over the given distribution.  $\square$

**THEOREM 4.3.** *We are given a joint distribution of  $x, y$ , where  $x$  has support  $\mathbb{X}$  and  $y$  has support  $\mathbb{Y}$ . For any model  $h : \mathbb{X} \rightarrow \mathbb{Y}$ , its  $(p, \epsilon)$ -tower robustness has a computable lower bound as*

$$(1 - \kappa) \frac{P(c_{PRA}(x, y; p, \epsilon, h, \kappa, \alpha)) - \alpha}{1 + \alpha}. \quad (40)$$

**PROOF.** We start from the tower law, and we can get the following inequality

$$\begin{aligned} E\left[P_{\|x-x'\|_p \leq \epsilon}(h(x') \neq y)\right] &\leq \\ \kappa P\left(P_{\|x-x'\|_p \leq \epsilon}(h(x') \neq y) \leq \kappa\right) &+ \\ \kappa^* P\left(P_{\|x-x'\|_p \leq \epsilon}(h(x') \neq y) > \kappa\right) \end{aligned} \quad (41)$$

where  $\kappa^*$  is a known upper bound that  $\kappa^* \geq P_{\|x-x'\|_p \leq \epsilon}(h(x') \neq y)$  on the function's output. A naive  $\kappa^*$  value is 1, which leads to

$$\begin{aligned} E\left[P_{\|x-x'\|_p \leq \epsilon}(h(x') \neq y)\right] &\leq \\ (\kappa - 1)P\left(P_{\|x-x'\|_p \leq \epsilon}(h(x') \neq y) \leq \kappa\right) &+ 1 \end{aligned} \quad (42)$$

Multiplying -1 and then adding 1 to both sides, we get

$$\begin{aligned} & \mathbb{E} \left[ P_{\|x-x'\|_p \leq \epsilon} (h(x') = y) \right] \\ & \geq (1 - \kappa) P \left( P_{\|x-x'\|_p \leq \epsilon} (h(x') \neq y) \leq \kappa \right) \end{aligned} \quad (43)$$

Substituting in a lower bound of the right-hand side, we get

$$\begin{aligned} & \mathbb{E} \left[ P_{\|x-x'\|_p \leq \epsilon} (h(x') = y) \right] \\ & > (1 - \kappa) \frac{P(c_{\text{PRA}}(x, y; p, \epsilon, h, \kappa, \alpha)) - \alpha}{1 + \alpha}. \end{aligned} \quad (44)$$

This completes the proof.  $\square$

Till here, we have successfully eliminated the helper parameters  $\kappa$  and  $\alpha$ . Theorems 4.2 and 4.3 provide the expression of the proposed TEB-U and TEB-L metric, respectively. In practice, we can freely choose convenient values of them and compare the model robustness of, say,  $h_1$  with another model  $h_2$  that is evaluated with different helper parameters. Essentially, we are comparing the lower bound, or the robustness guarantee of each model. And, when one intends to claim a better model robustness, it is reasonable to push the lower bound tighter by using smaller  $\alpha$  and  $\kappa$ . This is reasonable, compared to the existing practice that choosing higher  $\kappa$  and  $\alpha$  inflates the robustness score [45, 46].

Another advantage of using tower robustness, or its lower bound guarantee (TEB-L), to evaluate robustness is that in this unified framework, we can compare the deterministic robustness and probabilistic robustness, *i.e.*, we find a fair comparison method. Furthermore, the proposed lower bound can be used together with deterministic robustness verification, *i.e.*, for those cases that cannot be deterministically verified, we can compute their probabilistic robust accuracy and combine them. Such a combination is not merely a score addition, but is derivable from the governing expression of tower robustness.

A limitation of TEB-L/TEB-U is its tightness and computational cost arising from the fixed sampling scheme. Although the estimated sample size in Equation (25) is reasonable, it may still be too small (preventing rejection of the null hypothesis) or excessively large (reducing efficiency). A potential remedy is the sequential probability ratio test [36], which will be investigated in future work.

Another limitation of TEB-L/TEB-U lies in their distributional dependence. As with all metrics inherently based on MLE from samples, evaluating TEB-L/TEB-U under distributional drift is necessary, even for unseen examples. This direction may be further explored in conjunction with studies on distribution shift [16].

## 5 EMPIRICAL RESULTS

*Setup.* We compute TEB-L and TEB-U for existing neural networks. There are altogether six neural models, trained from ERM [34], PGD adversarial training [24], TRADES [42], MART [37], CVaR [31], and small-box certified training [27].

We evaluate these models on the commonly used MNIST [19] and CIFAR-10 tasks. Typically, MNIST classification is tested within  $(\infty, 1/10)$ -neighbourhood and CIFAR-10  $(\infty, 8/255)$ -neighbourhood, which is also our default setting. For each dataset, we evaluate on the 10,000 images in the testing set.

*Result.* We present our experimental results in Figure 1. Since the validity of tower robustness is formally shown in previous discussions, our experiments are primarily used to help understand and provide visual ideas. From Figure 1, we may first observe that CVaR so far obtains the highest lower bound of tower robustness on CIFAR-10, over 90.6%. The second highest is from PGD adversarial training, but only slightly above 80%. On MNIST, except for ERM, which does not optimise robustness and Small-box, which has concentrated primarily on deterministic robustness, all other models achieve a pretty high guarantee. These results aligned with our intuition and others' results [23]. Again, in a strict sense, this alignment would not lead us to any implication, but simply we are more or less not off track.

What brings us more benefits is to observe the trend when the helper parameters are decreasing. As mentioned in previous discussions, if a model provider would like to get a better certificate of the model's robustness, a fair approach is to squeeze these helper parameters. This lifts the lower bound, as observed in Figure 1, but does not invalidate it. Intuitively, this could be considered as different verification methods in the deterministic robustness context, where each model does not necessarily obtain a certificate in the same approach. In our experiment, we use a default  $\kappa = 1/10$ ,  $\alpha = 1/10$ . When reducing either  $\kappa$  or  $\alpha$ , the lower bound can grow, which overcomes the inflation problem prevalent in probabilistic robustness evaluations. Note that reducing  $\kappa$  turns out to be more effective than reducing  $\alpha$ . Intuitively,  $\kappa$  is more closely related to the proportion of adversarial examples than  $\alpha$  is.

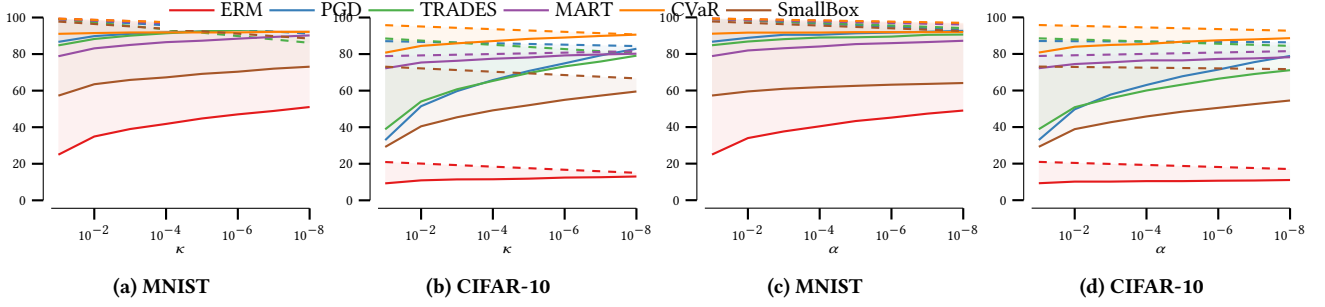
Intriguingly, we can apply the probabilistic robustness verification techniques to the test cases not deterministically verified. For example, Small-Box is a (deterministic) certified training method, which is usually evaluated only for deterministic robustness. While one could indeed subject this model to probabilistic robustness verification, there have been no theoretical guarantees provided for combining both. Our method is the first that put these two verification results together and provide an interpretable quantity. Yet, because the Small-Box model has primarily been tuned towards deterministic verification, its tower robustness lower bound remains one of the lowest on the list, even if we simultaneously enable probabilistic robustness verification.

As for the upper bound of tower robustness, it generally would not be used as a guarantee. But it can be used to confirm the defects of a given model. Its monotonicity is not obvious in expression (33), but we have observed it likely decreases when either  $\kappa$  or  $\alpha$  grows.

## 6 RELATED WORK

*Empirical-based Approaches.* A common and intuitive method for assessing model robustness is the test error rate [5], valued for its speed and versatility in evaluating both  $L^p$  norm and transformation perturbations (*e.g.*, rotation, illumination). Adversarial attacks exemplify this approach, identifying vulnerable samples within defined vicinities [5, 18]. However, this method lacks guarantees, as previously validated robustness may be undermined by stronger attacks, necessitating re-evaluation across models. To mitigate this, [13] proposes using a standardized adversarial test set, enabling efficient and consistent evaluation across models. Although this technique doesn't strictly align with any formal robustness





**Figure 1: Tower robustness guarantees.** Each solid line shows the lower bound, and the dashed line shows the upper bound.

definition and risks overfitting, it offers a practical, rapid method for accurately ranking models.

*Reachability-based algorithms.* take an input and a predefined perturbation to assess whether a given model meets specific safety criteria, such as robustness and boundary adherence. This evaluation of robustness can be achieved by addressing output range analysis problems through computational methods including interval analysis, linear programming, and optimization [21, 22, 35, 40]. Layerwise analysis is a prevalent approach in this context. The K-ReLU technique [33], which utilizes joint relaxation, effectively captures dependencies between inputs to different ReLUs within a layer, thus overcoming the convex barrier imposed by single neuron triangle relaxation and its approximations. Another notable method, CLEVER [39], transforms the robustness analysis into a local Lipschitz constant estimation problem and employs the Extreme Value Theory for efficient evaluation. However, it is important to note that reachability analysis techniques often require Lipschitz continuity over outputs for the target networks, which may limit their applicability in some instances.

*Constraint-based Verification.* aims to ensure that a given model adheres to a predefined set of robustness constraints. This process entails defining the constraints and then analyzing the model using techniques such as model checking, testing, or runtime monitoring [3, 14]. Marabou, for instance, translates queries into constraint satisfaction problems using Satisfiability Modulo Theories (SMT)[4]. It is compatible with various network activation functions and topologies, and it carries out high-level reasoning on the network [15]. Another SMT-based approach, Reluplex, employs the simplex method while extending it to accommodate the Rectified Linear Unit (ReLU) activation function [14]. Although constraint-based verification can offer formal guarantees that a model fulfils certain constraints or requirements, challenges such as computational complexity and pessimistic results have hindered its widespread application [46]. Consequently, researchers have proposed probabilistic robustness to address these limitations.

*Probabilistic-Based Approach.* Diverging from the deterministic verification methods previously examined, probabilistic approaches consider the balance between performance and sample complexity in both worst-case and average-case learning situations [7, 20, 31]. Cohen et al. [7] employ randomized smoothing as a reachability method, calculating a lower bound on robustness by adding noise to

input data and observing output behaviour. This lower bound can then be determined by the proportion of perturbed inputs yielding robust outputs. Additionally, Salman et al. [32] develops an adapted attack for smoothed classifiers and utilises adversarial training to enhance the provable robustness of these classifiers. However, these statistical methods primarily concentrate on pixel-level additive perturbations and often require assumptions about target neural networks or input distributions, limiting their broad applicability as they mostly function with  $L^2$  norms. Moreover, Zhang et al. [46] introduces PRoA, which relies on adaptive concentration. Although they employ an approximation method for distribution, it may lead to soundness issues in the future. In contrast, the proposed exact binomial test method retains the benefits of probabilistic approaches while using an exact solution to circumvent soundness problems. Similar to PRoA, our method can address black-box deep learning models without assumptions and is not restricted by transformation attacks (unbounded by  $L^p$  norms).

## 7 CONCLUSION

In this study, we introduce tower robustness, an innovative concept that makes the observation process of perturbation robustness intuitive. Throughout this work, we adopt a global perspective to unify the deterministic and probabilistic robustness, *i.e.*, we focus on the behaviour of an upcoming unseen sample.

We provide computable lower and upper bounds (TEB-L and TEB-U) of tower robustness. These metrics solve the overestimation problem of existing probabilistic robustness metrics by eliminating the excessive degree of freedom. Furthermore, we can combine the proposed bounds with deterministic robustness, and this still provides theoretical guarantees. Besides, TEB-L and TEB-U could be used for other machine learning models as well, not specifically limited to neural networks.

As a follow-up on this work, we plan to conduct a detailed analysis of the computational cost of TEB-L/TEB-U, investigate the tightness of the bounds theoretically, and further broaden our experimental evaluation in scope and depth. In future work, we will also explore incorporating differential privacy as a safeguard against adversarial attacks [26].

**Acknowledgements.** This research is supported in part by the Ministry of Education, Singapore, under its Academic Research Fund Tier 2 (Award No. MOE-T2EP20123-0015). Any opinions, findings and conclusions, or recommendations expressed in this material are those of the authors and do not reflect the views of the Ministry of Education, Singapore.

## GENAI USAGE DISCLOSURE

Generative AI tools were used for minor language editing, such as grammar/spelling checks and minor changes, on original author-written text.

## REFERENCES

- [1] Alan Agresti and Brent A. Coull. 1998. Approximate is Better than "Exact" for Interval Estimation of Binomial Proportions. *The American Statistician* 52, 2 (1998), 119–126. <https://doi.org/10.1080/00031305.1998.10480550> arXiv:<https://doi.org/10.1080/00031305.1998.10480550>
- [2] Abdulrahman Al-Abassi, Hadis Karimipour, Ali Dehghantanha, and Reza M. Parizi. 2020. An Ensemble Deep Learning-Based Cyber-Attack Detection in Industrial Control System. *IEEE Access* 8 (2020), 83965–83973. <https://doi.org/10.1109/ACCESS.2020.2992249>
- [3] Guy Amir, Haoze Wu, Clark Barrett, and Guy Katz. 2021. An SMT-based approach for verifying binarized neural networks. In *Tools and Algorithms for the Construction and Analysis of Systems: 27th International Conference, TACAS 2021, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2021, Luxembourg City, Luxembourg, March 27–April 1, 2021, Proceedings, Part II* 27. Springer, 203–222.
- [4] Clark Barrett and Cesare Tinelli. 2018. *Satisfiability modulo theories*. Springer.
- [5] Nicholas Carlini and David Wagner. 2017. Adversarial Examples Are Not Easily Detected: Bypassing Ten Detection Methods. Association for Computing Machinery, New York, NY, USA, 3–14. <https://doi.org/10.1145/3128572.3140444>
- [6] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A Simple Framework for Contrastive Learning of Visual Representations. In *Proceedings of the 37th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 119)*, Hal Daumé III and Aarti Singh (Eds.). PMLR, 1597–1607. <https://proceedings.mlr.press/v119/chen20j.html>
- [7] Jeremy Cohen, Elan Rosenfeld, and Zico Kolter. 2019. Certified Adversarial Robustness via Randomized Smoothing. In *Proceedings of the 36th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 97)*, Kamalika Chaudhuri and Ruslan Salakhutdinov (Eds.). PMLR, PMLR, 1310–1320. <https://proceedings.mlr.press/v97/cohen19c.html>
- [8] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. 2019. ArcFace: Additive Angular Margin Loss for Deep Face Recognition. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 4685–4694. <https://doi.org/10.1109/CVPR.2019.00482>
- [9] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiuhua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3–7, 2021*. OpenReview.net. <https://openreview.net/forum?id=YicbFdNTTy>
- [10] Luciano Floridi and Massimo Chiriatti. 2020. GPT-3: Its Nature, Scope, Limits, and Consequences. *Minds and Machines* 30, 4 (01 Dec 2020), 681–694. <https://doi.org/10.1007/s11023-020-09548-1>
- [11] Rui Fu, Zuo Zhang, and Li Li. 2016. Using LSTM and GRU neural network methods for traffic flow prediction. In *2016 31st Youth Academic Annual Conference of Chinese Association of Automation (YAC)*. 324–328. <https://doi.org/10.1109/YAC.2016.7804912>
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 770–778. <https://doi.org/10.1109/CVPR.2016.90>
- [13] Dan Hendrycks, Kevin Zhao, Steven Basart, Jacob Steinhardt, and Dawn Song. 2021. Natural Adversarial Examples. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 15257–15266. <https://doi.org/10.1109/CVPR46437.2021.01501>
- [14] Guy Katz, Clark Barrett, David L. Dill, Kyle Julian, and Mykel J. Kochenderfer. 2017. Reluplex: An Efficient SMT Solver for Verifying Deep Neural Networks. In *Computer Aided Verification*, Rupak Majumdar and Viktor Kunčák (Eds.). Springer International Publishing, Cham, 97–117.
- [15] Guy Katz, Derek A. Huang, Duligur Ibeling, Kyle Julian, Christopher Lazarus, Rachel Lim, Parth Shah, Shantanu Thakoor, Haoze Wu, Aleksandar Zeljić, David L. Dill, Mykel J. Kochenderfer, and Clark Barrett. 2019. The Marabou Framework for Verification and Analysis of Deep Neural Networks. In *Computer Aided Verification*, Isil Dillig and Serdar Tasiran (Eds.). Springer International Publishing, Cham, 443–452.
- [16] Pang Wei Koh, Shiori Sagawa, Henrik Marklund, Sang Michael Xie, Marvin Zhang, Akshay Balsubramani, Weihua Hu, Michihiro Yasunaga, Richard Lanus Phillips, Irena Gao, Tony Lee, Etienne David, Ian Stavness, Wei Guo, Berton Earnshaw, Imran Haque, Sara M Beery, Jure Leskovec, Anshul Kundaje, Emma Pierson, Sergey Levine, Chelsea Finn, and Percy Liang. 2021. WILDS: A Benchmark of in-the-Wild Distribution Shifts. In *Proceedings of the 38th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 139)*, Marina Meila and Tong Zhang (Eds.). PMLR, 5637–5664. <https://proceedings.mlr.press/v139/koh21a.html>
- [17] Alexey Kurakin, Ian J. Goodfellow, and Samy Bengio. 2017. Adversarial examples in the physical world. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24–26, 2017, Workshop Track Proceedings*. OpenReview.net. <https://openreview.net/forum?id=HJGU3R0d>
- [18] Alexey Kurakin, Ian J. Goodfellow, and Samy Bengio. 2017. Adversarial Machine Learning at Scale. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24–26, 2017, Conference Track Proceedings*. OpenReview.net. <https://openreview.net/forum?id=Bjm4T4Kgx>
- [19] Yann LeCun, Corinna Cortes, and Chris Burges. 1998. <http://yann.lecun.com/exdb/mnist/>
- [20] Mathias Lecuyer, Vaggelis Atlidakis, Roxana Geambasu, Daniel Hsu, and Suman Jana. 2019. Certified Robustness to Adversarial Examples with Differential Privacy. In *2019 IEEE Symposium on Security and Privacy (SP)*. 656–672. <https://doi.org/10.1109/SP.2019.00044>
- [21] Sungyoon Lee, Jaewook Lee, and Saerom Park. 2020. Lipschitz-Certifiable Training with a Tight Outer Bound. In *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (Eds.), Vol. 33. Curran Associates, Inc., 16891–16902. [https://proceedings.neurips.cc/paper\\_files/paper/2020/file/c46482dd5d39742f0bfd417b492d0e8e-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2020/file/c46482dd5d39742f0bfd417b492d0e8e-Paper.pdf)
- [22] Jianlin Li, Jiangchao Liu, Pengfei Yang, Lijian Chen, Xiaowei Huang, and Lijun Zhang. 2019. Analyzing Deep Neural Networks with Symbolic Propagation: Towards Higher Precision and Faster Verification. In *Static Analysis*, Bor-Yuh Evan Chang (Ed.). Springer International Publishing, Cham, 296–319.
- [23] Linyi Li, Tao Xie, and Bo Li. 2023. SoK: Certified Robustness for Deep Neural Networks. In *2023 IEEE Symposium on Security and Privacy (SP)*. 1289–1310. <https://doi.org/10.1109/SP46215.2023.10179303>
- [24] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. 2018. Towards Deep Learning Models Resistant to Adversarial Attacks. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net. <https://openreview.net/forum?id=rjzIBfZAb>
- [25] Wenchuan Mu and Kwan Hui Lim. 2024. Towards Precise Observations of Neural Model Robustness in Classification. In *Proceedings of the 2024 IEEE/ACM 46th International Conference on Software Engineering: Companion Proceedings (Lisbon, Portugal) (ICSE-Companion '24)*. Association for Computing Machinery, New York, NY, USA, 388–389. <https://doi.org/10.1145/3639478.3643519>
- [26] Wenchuan Mu and Kwan Hui Lim. 2025. Bayesian Privacy Guarantee for User History in Sequential Recommendation Using Randomised Response. In *Proceedings of the 34rd ACM International Conference on Information and Knowledge Management (Seoul, Republic of Korea) (CIKM '25)*. Association for Computing Machinery, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3746252.3760856>
- [27] Mark Niklas Mueller, Franziska Eckert, Marc Fischer, and Martin Vechev. 2023. Certified Training: Small Boxes are All You Need. In *The Eleventh International Conference on Learning Representations*. <https://openreview.net/forum?id=7oFuxtJtUMH>
- [28] Stuart J Pocock. 1977. Group sequential methods in the design and analysis of clinical trials. *Biometrika* 64, 2 (1977), 191–199.
- [29] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. Learning Transferable Visual Models From Natural Language Supervision. In *Proceedings of the 38th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 139)*, Marina Meila and Tong Zhang (Eds.). PMLR, 8748–8763. <https://proceedings.mlr.press/v139/radford21a.html>
- [30] Matthew Ragoza, Joshua Hochuli, Elisa Idrobo, Jocelyn Sunseri, and David Ryan Koes. 2017. Protein-Ligand Scoring with Convolutional Neural Networks. *Journal of Chemical Information and Modeling* 57, 4 (2017), 942–957. <https://doi.org/10.1021/acs.jcim.6b00740> arXiv:<https://doi.org/10.1021/acs.jcim.6b00740> PMID: 28368587
- [31] Alexander Robey, Luiz Chamon, George J. Pappas, and Hamed Hassani. 2022. Probabilistically Robust Learning: Balancing Average and Worst-case Performance. In *Proceedings of the 39th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 162)*, Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (Eds.). PMLR, 18667–18686. <https://proceedings.mlr.press/v162/robey22a.html>
- [32] Hadi Salman, Jerry Li, Ilya Razenshteyn, Pengchuan Zhang, Huan Zhang, Sebastian Bubeck, and Greg Yang. 2019. Provably Robust Deep Learning via Adversarially Trained Smoothed Classifiers. In *Advances in Neural Information Processing Systems*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Eds.), Vol. 32. Curran Associates, Inc. [https://proceedings.neurips.cc/paper\\_files/paper/2019/file/3a24b25a7b092a252166a1641ae953e7-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2019/file/3a24b25a7b092a252166a1641ae953e7-Paper.pdf)
- [33] Gagandeep Singh, Rupanshu Ganvir, Markus Püschel, and Martin Vechev. 2019. Beyond the Single Neuron Convex Barrier for Neural Network Certification. In *Advances in Neural Information Processing Systems*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Eds.), Vol. 32. Curran Associates, Inc. [https://proceedings.neurips.cc/paper\\_files/paper/2019/file/0a9fdabb17feb6ccb7ec405cfb85222c4-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2019/file/0a9fdabb17feb6ccb7ec405cfb85222c4-Paper.pdf)

- [34] Vladimir Vapnik. 1999. *The nature of statistical learning theory*. Springer science & business media.
- [35] Aladin Virmaux and Kevin Scaman. 2018. Lipschitz regularity of deep neural networks: analysis and efficient estimation. In *Advances in Neural Information Processing Systems*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Eds.), Vol. 31. Curran Associates, Inc. <https://proceedings.neurips.cc/paper/2018/file/d54e99a6c03704e95e6965532dec148b-Paper.pdf>
- [36] A. Wald. 1945. Sequential Tests of Statistical Hypotheses. *The Annals of Mathematical Statistics* 16, 2 (1945), 117 – 186. <https://doi.org/10.1214/aoms/1177731118>
- [37] Yisen Wang, Difan Zou, Jinfeng Yi, James Bailey, Xingjun Ma, and Quanquan Gu. 2020. Improving Adversarial Robustness Requires Revisiting Misclassified Examples. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=rklOg6EFwS>
- [38] Andrew Wen, Mohamed Y Elwazir, Sungrim Moon, and Jungwei Fan. 2020. Adapting and evaluating a deep learning language model for clinical why-question answering. *JAMIA Open* 3, 1 (02 2020), 16–20. <https://doi.org/10.1093/jamiaopen/ooz072> arXiv:<https://academic.oup.com/jamiaopen/article-pdf/3/1/16/33419137/ooz072.pdf>
- [39] Tsui-Wei Weng, Huan Zhang, Pin-Yu Chen, Jinfeng Yi, Dong Su, Yupeng Gao, Cho-Jui Hsieh, and Luca Daniel. 2018. Evaluating the Robustness of Neural Networks: An Extreme Value Theory Approach. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net. <https://openreview.net/forum?id=BkUHlMZ0b>
- [40] Xiaodong Yang, Taylor T Johnson, Hoang-Dung Tran, Tomoya Yamaguchi, Bardh Hoxha, and Danil Prokhorov. 2021. Reachability analysis of deep ReLU neural networks using facet-vertex incidence. In *Proceedings of the 24th International Conference on Hybrid Systems: Computation and Control* (Nashville, Tennessee) (HSCC '21). Association for Computing Machinery, New York, NY, USA, Article 18, 7 pages. <https://doi.org/10.1145/3447928.3456650>
- [41] Wenhao Yu, Chengxiang Zhao, Hong Wang, Jiaxin Liu, Xiaohan Ma, Yingkai Yang, Jun Li, Weida Wang, Xiaosong Hu, and Ding Zhao. 2024. Online legal driving behavior monitoring for self-driving vehicles. *Nature Communications* 15, 1 (09 Jan 2024), 408. <https://doi.org/10.1038/s41467-024-44694-5>
- [42] Hongyang Zhang, Yaodong Yu, Jiantao Jiao, Eric Xing, Laurent El Ghaoui, and Michael Jordan. 2019. Theoretically Principled Trade-off between Robustness and Accuracy. In *Proceedings of the 36th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 97)*, Kamalika Chaudhuri and Ruslan Salakhutdinov (Eds.). PMLR, 7472–7482. <https://proceedings.mlr.press/v97/zhang19p.html>
- [43] Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter J. Liu. 2020. PEGASUS: Pre-training with Extracted Gap-sentences for Abstractive Summarization. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event (Proceedings of Machine Learning Research, Vol. 119)*. PMLR, 11328–11339. <http://proceedings.mlr.press/v119/zhang20ae.html>
- [44] Ruihan Zhang and Jun Sun. 2024. Certified Robust Accuracy of Neural Networks Are Bounded Due to Bayes Errors. In *Computer Aided Verification, Arie Gurfinkel and Vijay Ganesh (Eds.)*. Springer Nature Switzerland, Cham, 352–376.
- [45] Ruihan Zhang and Jun Sun. 2024. How Does Bayes Error Limit Probabilistic Robust Accuracy. *arXiv preprint arXiv:2405.14923* (2024).
- [46] Tianle Zhang, Wenjie Ruan, and Jonathan E. Fieldsend. 2023. PRoA: A Probabilistic Robustness Assessment Against Functional Perturbations. In *Machine Learning and Knowledge Discovery in Databases, Massih-Reza Amini, Stéphane Canu, Asja Fischer, Tias Guns, Petra Kralj Novak, and Grigorios Tsoumakas (Eds.)*. Springer Nature Switzerland, Cham, 154–170.