



Neural Network Robustness as a Verification Property: A Principled Case Study

Marco Casadio^{1(✉)}, Ekaterina Komendantskaya¹, Matthew L. Daggitt¹, Wen Kokke²,
Guy Katz³, Guy Amir³, and Idan Refaeli³

¹ Heriot-Watt University, Edinburgh, UK
{mc248, ek19, md2006}@hw.ac.uk

² University of Strathclyde, Glasgow, UK
wen.kokke@strath.ac.uk

³ The Hebrew University of Jerusalem, Jerusalem, Israel
{guykatz, guyam, idan0610}@cs.huji.ac.il

Abstract. Neural networks are very successful at detecting patterns in noisy data, and have become the technology of choice in many fields. However, their usefulness is hampered by their susceptibility to *adversarial attacks*. Recently, many methods for measuring and improving a network’s robustness to adversarial perturbations have been proposed, and this growing body of research has given rise to numerous explicit or implicit notions of robustness. Connections between these notions are often subtle, and a systematic comparison between them is missing in the literature. In this paper we begin addressing this gap, by setting up general principles for the empirical analysis and evaluation of a network’s robustness as a mathematical property—during the network’s training phase, its verification, and after its deployment. We then apply these principles and conduct a case study that showcases the practical benefits of our general approach.

Keywords: Neural Networks · Adversarial Training · Robustness · Verification

1 Introduction

Safety and security are critical for many complex systems that use deep neural networks (DNNs). Unfortunately, due to the opacity of DNNs, these properties are difficult to ensure. Perhaps the most famous instance of this problem is guaranteeing the robustness of DNN-based systems against *adversarial attacks* [5, 17]. Intuitively, a neural network is ϵ -ball robust around a particular input if, when you move no more than ϵ away from that input in the input space, the output does not change much; or, alternatively, the classification decision that the network gives does not change. Even highly accurate DNNs will often display only low robustness, and so measuring and improving the adversarial robustness of DNNs has received significant attention by both the machine learning and verification communities [7, 8, 15].

As a result, neural network verification often follows a *continuous verification cycle* [9], which involves retraining neural networks with a given *verification property* in mind, as Fig. 1 shows. More generally, such training can be regarded as a way to

impose a formal specification on a DNN; and so, apart from improving its robustness, it may also contribute to the network’s explainability, and facilitate its verification. Due to the high level of interest in adversarial robustness, numerous approaches have been proposed for performing such retraining in recent years, each with its own specific details. However it is quite unclear what are the benefits that each approach offers, from a verification point of view.

The primary goal of this case-study paper is to introduce a more holistic methodology, which puts the verification property in the centre of the development cycle, and in turn permits a principled analysis of how this property influences both training and verification practices. In particular, we analyse the verification properties that implicitly or explicitly arise from the most prominent families of training techniques: *data augmentation* [14], *adversarial training* [5,10], *Lipschitz robustness training* [1,12], and *training with logical constraints* [4,20]. We study the effect of each of these properties on verifying the DNN in question.

In Sect. 2, we start with the forward direction of the continuous verification cycle, and show how the above training methods give rise to logical properties of *classification robustness* (CR), *strong classification robustness* (SCR), *standard robustness* (SR) and *Lipschitz robustness* (LR). In Sect. 4, we trace the opposite direction of the cycle, i.e. show how and when the verifier failure in proving these properties can be mitigated. However Sect. 3 first gives an auxiliary logical link for making this step. Given a robustness property as a logical formula, we can use it not just in verification, but also in attack or property accuracy measurements. We take property-driven attacks as a valuable tool in our study, both in training and in evaluation. Section 4 makes the underlying assumption that verification requires retraining: it shows that the verifier’s success ranges only 0–1.5% for an accurate baseline network. We show how our logical understanding of robustness properties empowers us in property-driven training and in verification. We first give abstract arguments why certain properties are stronger than others or incomparable; and then we use training, attacks and the verifier Marabou to confirm them empirically. Sections 5 and 6 add other general considerations for setting up the continuous verification loop and conclude the paper.




Fig. 1. Continuous Verification Cycle

2 Existing Training Techniques and Definitions of Robustness

Data Augmentation is a straightforward method for improving robustness via training [14]. It is applicable to any transformation of the input (e.g. addition of noise, translation, rotation, scaling) that leaves the output label unchanged. To make the network robust against such a transformation, one augments the dataset with instances sampled via the transformation.

More formally, given a neural network $N : \mathbb{R}^n \rightarrow \mathbb{R}^m$, the goal of data augmentation is to ensure *classification robustness*, which is defined as follows. Given a training

dataset input-output pair $(\hat{\mathbf{x}}, \mathbf{y})$ and a distance metric $|\cdot - \cdot|$, for all inputs \mathbf{x} within the ϵ -ball distance of $\hat{\mathbf{x}}$, we say that N is *classification-robust* if class \mathbf{y} has the largest score in output $N(\mathbf{x})$.

Definition 1 (Classification robustness).

$$CR(\epsilon, \hat{\mathbf{x}}) \triangleq \forall \mathbf{x} : |\mathbf{x} - \hat{\mathbf{x}}| \leq \epsilon \Rightarrow \arg \max N(\mathbf{x}) = \mathbf{y}$$

In order to apply data augmentation, an engineer needs to specify: **c1.** the value of ϵ , i.e. the admissible range of perturbations; **c2.** the distance metric, which is determined according to the admissible geometric perturbations; and **c3.** the sampling method used to produce the perturbed inputs (e.g., random sampling, adversarial attacks, generative algorithm, prior knowledge of images).

Classification robustness is straightforward, but does not account for the possibility of having “uncertain” images in the dataset, for which a small perturbation ideally should change the class. For datasets that contain a significant number of such images, attaining this kind of training could lead to a significant reduction in accuracy.

Adversarial training is a current state-of-the-art method to robustify a neural network. Whereas standard training tries to minimise loss between the predicted value, $f(\hat{\mathbf{x}})$, and the true value, \mathbf{y} , for each entry $(\hat{\mathbf{x}}, \mathbf{y})$ in the training dataset, adversarial training minimises the loss with respect to the worst-case perturbation of each sample in the training dataset. It therefore replaces the standard training objective $\mathcal{L}(\hat{\mathbf{x}}, \mathbf{y})$ with: $\max_{\mathbf{x}: |\mathbf{x} - \hat{\mathbf{x}}| \leq \epsilon} \mathcal{L}(\mathbf{x}, \mathbf{y})$. Algorithmic solutions to the maximisation problem that find the worst-case perturbation has been the subject of several papers. The earliest suggestion was the Fast Gradient Sign Method (FGSM) algorithm introduced by [5]:

$$\text{FGSM}(\hat{\mathbf{x}}) = \hat{\mathbf{x}} + \epsilon \cdot \text{sign}(\nabla_{\mathbf{x}} \mathcal{L}(\mathbf{x}, \mathbf{y}))$$

However, modern adversarial training methods usual rely on some variant of the Projected Gradient Descent (PGD) algorithm [11] which iterates FGSM:

$$\text{PGD}_0(\hat{\mathbf{x}}) = \hat{\mathbf{x}}; \quad \text{PGD}_{t+1}(\hat{\mathbf{x}}) = \text{PGD}_t(\text{FGSM}(\hat{\mathbf{x}}))$$

It has been empirically observed that neural networks trained using this family of methods exhibit greater robustness at the expense of an increased generalisation error [10, 18, 21], which is frequently referred to as the *accuracy-robustness trade-off* for neural networks (although this effect has been observed to disappear as the size of the training dataset grows [13]).

In logical terms what is this procedure trying to train for? Let us assume that there’s some maximum distance, δ , that it is acceptable for the output to be perturbed given the size of perturbations in the input. This leads us to the following definition, where $\|\cdot - \cdot\|$ is a suitable distance function over the output space:

Definition 2 (Standard robustness).

$$SR(\epsilon, \delta, \hat{\mathbf{x}}) \triangleq \forall \mathbf{x} : |\mathbf{x} - \hat{\mathbf{x}}| \leq \epsilon \Rightarrow \|f(\mathbf{x}) - f(\hat{\mathbf{x}})\| \leq \delta$$

We note that, just as with data augmentation, choices **c1–c3** are still there to be made, although the sampling methods are usually given by special-purpose FGSM/PGD heuristics based on computing the loss function gradients.

Training for Lipschitz Robustness. More recently, a third competing definition of robustness has been proposed: Lipschitz robustness [2]. Inspired by the well-established concept of Lipschitz continuity, Lipschitz robustness asserts that the distance between the original output and the perturbed output is at most a constant L times the change in the distance between the inputs.

Definition 3 (Lipschitz robustness).

$$LR(\epsilon, L, \hat{\mathbf{x}}) \triangleq \forall \mathbf{x} : |\mathbf{x} - \hat{\mathbf{x}}| \leq \epsilon \Rightarrow \|f(\mathbf{x}) - f(\hat{\mathbf{x}})\| \leq L|\mathbf{x} - \hat{\mathbf{x}}|$$

As will be discussed in Sect. 4, this is a stronger requirement than standard robustness. Techniques for training for Lipschitz robustness include formulating it as a semi-definite programming optimisation problem [12] or including a projection step that restricts the weight matrices to those with suitable Lipschitz constants [6].

Training with Logical Constraints. Logically, this discussion leads one to ask whether a more general approach to constraint formulation may exist, and several attempts in the literature addressed this research question [4, 20], by proposing methods that can translate a first-order logical formula C into a *constraint loss function* \mathcal{L}_C . The loss function penalises the network when outputs do not satisfy a given Boolean constraint, and universal quantification is handled by a choice of sampling method. Our standard loss function \mathcal{L} is substituted with:

$$\mathcal{L}^*(\hat{\mathbf{x}}, \mathbf{y}) = \alpha \mathcal{L}(\hat{\mathbf{x}}, \mathbf{y}) + \beta \mathcal{L}_C(\hat{\mathbf{x}}, \mathbf{y}) \quad (1)$$

where weights α and β control the balance between the standard and constraint loss.

This method looks deceptively as a generalisation of previous approaches. However, even given suitable choices for **c1–c3**, classification robustness cannot be modelled via a constraint loss in the DL2 [4] framework, as *argmax* is not differentiable. Instead, [4] defines an alternative constraint, which we call *strong classification robustness*:

Definition 4 (Strong classification robustness).

$$SCR(\epsilon, \eta, \hat{\mathbf{x}}) \triangleq \forall \mathbf{x} : |\mathbf{x} - \hat{\mathbf{x}}| \leq \epsilon \Rightarrow f(\mathbf{x}) \geq \eta$$

which looks only at the prediction of the true class and checks whether it is greater than some value η (chosen to be 0.52 in their work).

We note that sometimes, the constraints (and therefore the derived loss functions) refer to the true label \mathbf{y} rather than the current output of the network $f(\hat{\mathbf{x}})$, e.g. $\forall \mathbf{x} : |\mathbf{x} - \hat{\mathbf{x}}| \leq \epsilon \Rightarrow |f(\mathbf{x}) - \mathbf{y}| \leq \delta$. This leads to scenarios where a network that is robust around $\hat{\mathbf{x}}$ but gives the wrong prediction, being penalised by \mathcal{L}_C which on paper is designed to maximise robustness. Essentially \mathcal{L}_C is trying to maximise both accuracy and constraint adherence concurrently. Instead, we argue that to preserve the intended semantics of α and β it is important to instead compare against the current output of the network. Of course, this does not work for SCR because deriving the most popular class from the output $f(\hat{\mathbf{x}})$ requires the $\arg \max$ operator—the very function that SCR seeks to avoid using. This is another argument why (S)CR should be avoided if possible.

3 Robustness in Evaluation, Attack and Verification

Given a particular definition of robustness, a natural question is how to quantify how close a given network is to satisfying it. We argue that there are three different measures that one should be interested in: 1. Does the constraint hold? This is a binary measure, and the answer is either true or false. 2. If the constraint does not hold, how easy is it for an attacker to find a violation? 3. If the constraint does not hold, how often does the average user encounter a violation? Based on these measures, we define three concrete metrics: *constraint satisfaction*, *constraint security*, and *constraint accuracy*.¹

Let \mathcal{X} be the training dataset, $\mathbb{B}(\hat{\mathbf{x}}, \epsilon) \triangleq \{\mathbf{x} \in \mathbb{R}^n \mid |\mathbf{x} - \hat{\mathbf{x}}| \leq \epsilon\}$ be the ϵ -ball around $\hat{\mathbf{x}}$ and P be the right-hand side of the implication in each of the definitions of robustness. Let \mathbb{I}_ϕ be the standard indicator function which is 1 if constraint $\phi(\mathbf{x})$ holds and 0 otherwise. The *constraint satisfaction* metric measures the proportion of the (finite) training dataset for which the constraint holds.

Definition 5 (Constraint satisfaction).

$$\text{CSat}(\mathcal{X}) = \frac{1}{|\mathcal{X}|} \sum_{\hat{\mathbf{x}} \in \mathcal{X}} \mathbb{I}_{\forall \mathbf{x} \in \mathbb{B}(\hat{\mathbf{x}}, \epsilon): P(\mathbf{x})}$$

In contrast, *constraint security* measures the proportion of inputs in the dataset such that an attack A is unable to find an adversarial example for constraint P . In our experiments we use the PGD attack for A , although in general any strong attack can be used.

Definition 6 (Constraint security).

$$\text{CSec}(\mathcal{X}) = \frac{1}{|\mathcal{X}|} \sum_{\hat{\mathbf{x}} \in \mathcal{X}} \mathbb{I}_P(A(\hat{\mathbf{x}}))$$

Finally, *constraint accuracy* estimates the probability of a random user coming across a counter-example to the constraint, usually referred as *I - success rate* in the robustness literature. Let $S(\hat{\mathbf{x}}, n)$ be a set of n elements randomly uniformly sampled from $\mathbb{B}(\hat{\mathbf{x}}, \epsilon)$. Then constraint accuracy is defined as:

Definition 7 (Constraint accuracy).

$$\text{CAcc}(\mathcal{X}) = \frac{1}{|\mathcal{X}|} \sum_{\hat{\mathbf{x}} \in \mathcal{X}} \left(\frac{1}{n} \sum_{\mathbf{x} \in S(\hat{\mathbf{x}}, n)} \mathbb{I}_P(\mathbf{x}) \right)$$

Note that there is no relationship between constraint accuracy and constraint security: an attacker may succeed in finding an adversarial example where random sampling fails and vice-versa. Also note the role of sampling in this discussion and compare it to the discussion of the choice **c3** in Sect. 2. Firstly, sampling procedures affect both training and evaluation of networks. But at the same time, their choice is orthogonal

¹ Our naming scheme differs from [4] who use the term *constraint accuracy* to refer to what we term *constraint security*. In our opinion, the term *constraint accuracy* is less appropriate here than the name *constraint security* given the use of an adversarial attack.

to choosing the verification constraint for which we optimise or evaluate. For example, we measure constraint security with respect to the PGD attack, and this determines the way we sample; but having made that choice still leaves us to decide which constraint, SCR, SR, LR, or other we will be measuring as we sample. Constraint satisfaction is different from constraint security and accuracy, in that it must evaluate constraints over infinite domains rather than merely sampling from them.

Choosing an Evaluation Metric. It is important to note that for all three evaluation metrics, one still has to make a choice for constraint P , namely SR, SCR or LR, as defined in Sect. 2. As constraint security always uses PGD to find input perturbations, the choice of SR, SCR and LR effectively amounts to us making a judgement of what an adversarial perturbation consists of: is it a class change as defined by SCR, or is it a violation of the more nuanced metrics defined by SR and LR? Therefore we will evaluate constraint security on the *SR/SCR/LR constraints* using a *PGD attack*.

For large search spaces in n dimensions, random sampling deployed in constraint accuracy fails to find the trickier adversarial examples, and usually has deceptively high performance: we found 100% and >98% constraint accuracy for SR and SCR, respectively. We will therefore not discuss these experiments in detail.

4 Relative Comparison of Definitions of Robustness

We now compare the strength of the given definitions of robustness using the introduced metrics. For empirical evaluation, we train networks on *FASHION MNIST* (or just *FASHION*) [19] and a modified version of the *GTSRB* [16] datasets consisting, respectively, by 28×28 and 48×48 images belonging to 10 classes. The networks consist of two fully connected layers: the first one having 100 neurons and ReLU as activation function, and the last one having 10 neurons on which we apply a clamp function $[-100, 100]$, because the traditional softmax function is not compatible with constraint verification tools such as Marabou. Taking four different robustness properties for which we optimise while training (Baseline, LR, SR, SCR), gives us 8 different networks to train, evaluate and attack. Generally, all trends we observed for the two data sets were the same, and we put matching graphs in [3] whenever we report a result for one of the data sets. Marabou [8] was used for evaluating constraint satisfaction.

4.1 Standard and Lipschitz Robustness

Lipschitz robustness is a strictly stronger constraint than standard robustness, in the sense that when a network satisfies $LR(\epsilon, L)$ then it also satisfies $SR(\epsilon, \epsilon L)$. However, the converse does not hold, as standard robustness does not relate the distances between the inputs and the outputs. Consequently, there are $SR(\epsilon, \delta)$ robust models that are not $LR(\epsilon, L)$ robust for any L , as for any fixed L one can always make the distance $|\mathbf{x} - \tilde{\mathbf{x}}|$ arbitrarily small in order to violate the Lipschitz inequality.

Table 1. Constraint satisfaction results for the Classification, Standard and Lipschitz constraints. These values are calculated over the test set and represented as %.

	FASHION net trained with:				GTSRB net trained with:			
	Baseline	SCR	SR	LR	Baseline	SCR	SR	LR
CR satisfaction	1.5	2.0	2.0	34.0	0.5	1.0	3.0	4.5
SR satisfaction	0.5	1.0	65.8	100.0	0.0	0.0	24.0	97.0
LR satisfaction	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0

Empirical Significance of the Conclusions for Constraint Security. Figure 2 shows an empirical evaluation of this general result. If we train two neural networks, one with the SR, and the other with the LR constraint, then the latter always has higher constraint security against both SR and LR attacks than the former. It also confirms that generally, stronger constraints are harder to obtain: whether a network is trained with SR or LR constraints, it is less robust against an LR attack than against any other attack.

Empirical Significance of the Conclusions for Constraint Satisfaction. Table 1 shows that LR is very difficult to guarantee as a verification property, indeed none of our networks satisfied this constraint for any image in the data set. At the same time, networks trained with LR satisfy the weaker property SR, for 100% and 97% of images – a huge improvement on the negligible percentage of robust images for the baseline network! Therefore, knowing a verification property or mode of attack, one can tailor the training accordingly, and training with stronger constraint gives better results.

4.2 (Strong) Classification Robustness

Strong classification robustness is designed to over-approximate classification robustness whilst providing a logical loss function with a meaningful gradient. We work under the assumption that the last layer of the classification network is a softmax layer, and

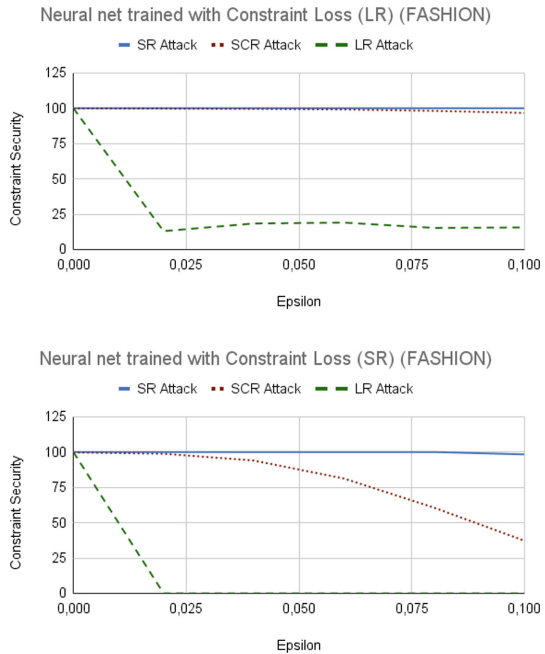


Fig. 2. Experiments that show how the two networks trained with LR and SR constraints perform when evaluated against different definitions of robustness underlying the attack; ϵ measures the attack strength.

therefore the output forms a probability distribution. When $\eta > 0.5$ then any network that satisfies $SCR(\epsilon, \eta)$ also satisfies $CR(\epsilon)$. For $\eta \leq 0.5$ this relationship breaks down as the true class may be assigned a probability greater than η but may still not be the class with the highest probability. We therefore recommended that one only uses value of $\eta > 0.5$ when using strong classification robustness (for example $\eta = 0.52$ in [4]).

Empirical Significance of the Conclusions for Constraint Security.

Because the CR constraint cannot be used within a loss function, we use data augmentation when training to emulate its effect. First, we confirm our assumptions about the relative inefficiency of using data augmentation compared to adversarial training or training with constraints, see Fig. 3. Surprisingly, neural networks trained with data augmentation give worse results than even the baseline network.

As previously discussed, random uniform sampling struggles to find adversarial inputs in large searching spaces. It is logical to expect that using random uniform sampling when training will be less successful than training with sampling that uses FGSM or PGD as heuristics. Indeed, Fig. 3 shows this effect for data augmentation.

One may ask whether the trends just described would be replicated

for more complex architectures of neural networks. In particular, data augmentation is known to require larger networks. By replicating the results with a large, 18-layer convolutional network from [4] (second graph of Fig. 3), we confirm that larger networks handle data augmentation better, and that data augmentation affords improved robustness compared to the baseline. Nevertheless, data augmentation still lags behind all other modes of constraint-driven training, and thus this major trend remains stable across network architectures. The same figure also illustrates our point about the relative strength of SCR compared to CR: a network trained with data augmentation (equivalent to CR) is more prone to SCR attacks than a network trained with the SCR constraint.

Empirical Significance of the Conclusions for Constraint Satisfaction. Although Table 1 confirms that training with a stronger property (SCR) does improve the constraint satisfaction of a weaker property (CR), the effect is an order of magnitude smaller than what we observed for LR and SR. Indeed, the table suggests that training with the




Fig. 3. Experiments that show how adversarial training, training with data augmentation, and training with constraint loss affect standard and classification robustness of networks; ϵ measures the attack strength.

LR constraint gives better results for CR constraint satisfaction. This does not contradict, but does not follow from our theoretical analysis.

4.3 Standard vs Classification Robustness

Given that LR is stronger than SR and SCR is stronger than CR, the obvious question is whether there is a relationship between these two groups. In short, the answer to this question is no. In particular, although the two sets of definitions agree on whether a network is robust around images with high-confidence, they disagree over whether a network is robust around images with low confidence. We illustrate this with an example, comparing SR against CR. We note that a similar analysis holds for any pairing from the two groups.

The key insight is that standard robustness bounds the drop in confidence that a neural network can exhibit after a perturbation, whereas classification robustness does not. Figure 4a shows two hypothetical images from the MNIST dataset. Our network predicts that Fig. 4a has an 85% chance of being a 7. Now consider adding a small perturbation to the image and consider two different scenarios. In the first scenario the output of the network for class 7 decreases from 85% to 83% and therefore the classification stays the same. In the second scenario the output of the network for class 7 decreases from 85% to 45%, and results in the classification changing from 7 to 9. When considering the two definitions, a small change in the output leads to no change in the classification and a large change in the output leads to a change in classification and so robustness and classification robustness both agree with each other.




Fig. 4. Images from the MNIST set

However, now consider Fig. 4b with relatively high uncertainty. In this case the network is (correctly) less sure about the image, only narrowly deciding that it's a 7. Again consider adding a small perturbation. In the first scenario the prediction of the network changes dramatically with the probability of it being a 7 increasing from 51% to 91% but leaves the classification unchanged as 7. In the second scenario the output of the network only changes very slightly, decreasing from 51% to 49% flipping the classification from 7 to 9. Now, the definitions of SR and CR disagree. In the first case, adding a small amount of noise has erroneously massively increased the network's confidence and therefore the SR definition correctly identifies that this is a problem. In contrast CR has no problem with this massive increase in confidence as the chosen output class remains unchanged. Thus, SR and CR agree on low-uncertainty examples, but CR breaks down and gives what we argue are both false positives and false negatives when considering examples with high-uncertainty.

Empirical Significance of the Conclusions for Constraint Security. Our empirical study confirms these general conclusions. Figure 2 shows that depending on the properties of the dataset, SR may not guarantee SCR. The results in Fig. 5 tell us that using the SCR constraint for training does not help to increase defences against SR attacks. A similar picture, but in reverse, can be seen when we optimize for SR but attack with SCR. Table 1 confirms these trends for constraint satisfaction.

5 Other Properties of Robustness Definitions

Table 2. A comparison of the different types of robustness studied in this paper. Top half: general properties. Bottom half: relation to existing machine-learning literature

Definition	Standard robustness	Lipschitz robustness	Classification robustness	Strong class. robustness
Problem domain	General	General	Classification	Classification
Interpretability	Medium	Low	High	Medium
Globally desirable	✓	✓	✗	✗
Has loss functions	✓	✓	✗	✓
Adversarial training	✓	✗	✗	✗
Data augmentation	✗	✗	✓	✗
Logical-constraint training [4]	✓	✓	✗	✓

We finish with a summary of further interesting properties of the four robustness definitions. Table 2 shows a summary of all comparison measures considered in the paper.

Dataset assumptions concern the distribution of the training data with respect to the data manifold of the true distribution of inputs, and influence evaluation of robustness. For SR and LR it is, at minimum, desirable for the network to be robust over the entire data manifold. In the most domains the shape of the manifold is unknown and therefore it is necessary to approximate it by taking the union of the balls around the inputs in the training dataset. We are not particularly interested about whether the network is robust in regions of the input space that lie off the data manifold, but there is no problem if the network is robust in these regions. Therefore these definitions make no assumptions about the distribution of the training dataset.

This is in contrast to CR and SCR. Rather than requiring that there is only a small change in the output, they require that there is no change to the classification. This is only a desirable constraint when the region being considered does not contain a decision boundary. Consequently when one is training for some form of classification robustness, one is implicitly making the assumption that the training data points lie away from any decision boundaries within the manifold. In practice, most datasets for classification problems assign a single label instead of an entire probability distribution to each input point, and so this assumption is usually valid. However, for datasets that contain input




Fig. 5. Experiments that show how different choices of a constraint loss affect standard robustness of neural networks.

points that may lie close to the decision boundaries, CR and SCR may result in a logically inconsistent specification.

Interpretability. One of the key selling points of training with logical constraints is that, by ensuring that the network obeys understandable constraints, it improves the explainability of the neural network. Each of the robustness constraints encode that “small changes to the input only result in small changes to the output”, but the interpretability of each definition is also important.

All of the definitions share the relatively interpretable ϵ parameter, which measures how large a perturbation from the input is acceptable. Despite the other drawbacks discussed so far, CR is inherently the most interpretable as it has no second parameter. In contrast, SR and SCR require extra parameters, δ and η respectively, which measure the allowable deviation in the output. Their addition makes these models less interpretable.

Finally we argue that, although LR is the most desirable constraint, it is also the least interpretable. Its second parameter L measures the allowable change in the output as a proportion of the allowable change in the input. It therefore requires one to not only have an interpretation of distance for both the input and output spaces, but to be able to relate them. In most domains, this relationship simply does not exist. Consider the MNIST dataset, both the commonly used notion of pixel-wise distance used in the input set, although crude, and the distance between the output distributions are both interpretable. However, the relationship between them is not. For example, what does allowing the distance between the output probability distributions being no more than twice the distance between the images actually mean? This therefore highlights a common trade-off between complexity of the constraint and its interpretability.

6 Conclusions

These case studies have demonstrated the importance of emancipating the study of desirable properties of neural networks from a concrete training method, and studying these properties in an abstract mathematical way. For example, we have discovered that some robustness properties can be ordered by logical strength and some are incomparable. Where ordering is possible, training for a stronger property helps in verifying a weaker property. Some of the stronger properties, such as Lipschitz robustness, are not yet feasible for the modern DNN solvers, such as Marabou [8]. Moreover, we show that the logical strength of the property may not guarantee other desirable properties, such as interpretability. Some of these findings lead to very concrete recommendations, e.g.: it is best to avoid CR and SCR as they may lead to inconsistencies; when using LR and SR, one should use stronger property (LR) for training in order to be successful in verifying a weaker one (SR). In other cases, the distinctions that we make do not give direct prescriptions, but merely discuss the design choices and trade-offs.

This paper also shows that constraint security, a measure intermediate between constraint accuracy and constraint satisfaction, is a useful tool in the context of tuning the continuous verification loop. It is more efficient to measure and can show more nuanced trends than constraint satisfaction. It can be used to tune training parameters and build hypotheses which we ultimately confirm with constraint satisfaction.

We hope that this study will contribute towards establishing a solid methodology for continuous verification, by setting up some common principles to unite verification and machine learning approaches to DNN robustness.

Acknowledgement. Authors acknowledge support of EPSRC grant AISEC EP/T026952/1 and NCSC grant Neural Network Verification: in search of the missing spec.

References

1. Anil, C., Lucas, J., Grosse, R.: Sorting out Lipschitz function approximation. In: International Conference on Machine Learning, pp. 291–301. PMLR (2019)
2. Balan, R., Singh, M., Zou, D.: Lipschitz properties for deep convolutional networks. *Contemp. Math.* **706**, 129–151 (2018)
3. Casadio, M., et al.: Neural network robustness as a mathematical property: a principled case study (2021). <https://github.com/aisec-private/training-with-constraints>
4. Fischer, M., Balunovic, M., Drachler-Cohen, D., Gehr, T., Zhang, C., Vechev, M.T.: DL2: training and querying neural networks with logic. In: Proceedings of the 36th International Conference Machine Learning, ICML 2019, vol. 97, pp. 1931–1941. PMLR (2019)
5. Goodfellow, I.J., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples. In: Bengio, Y., LeCun, Y. (eds.) 3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, 7–9 May 2015, Conference Track Proceedings (2015)
6. Gouk, H., Frank, E., Pfahringer, B., Cree, M.J.: Regularisation of neural networks by enforcing Lipschitz continuity. *Mach. Learn.* **110**(2), 393–416 (2020). <https://doi.org/10.1007/s10994-020-05929-w>
7. Huang, X., Kwiatkowska, M., Wang, S., Wu, M.: Safety verification of deep neural networks. In: Majumdar, R., Kunčák, V. (eds.) CAV 2017. LNCS, vol. 10426, pp. 3–29. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-63387-9_1
8. Katz, G., et al.: The Marabou framework for verification and analysis of deep neural networks. In: Dillig, I., Tasiran, S. (eds.) CAV 2019. LNCS, vol. 11561, pp. 443–452. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-25540-4_26
9. Komendantskaya, E., Kokke, W., Kienitz, D.: Continuous verification of machine learning: a declarative programming approach. In: PPDP 2020: 22nd International Symposium on Principles and Practice of Declarative Programming, Bologna, Italy, 9–10 September 2020, pp. 1:1–1:3. ACM (2020)
10. Madry, A., Makelov, A., Schmidt, L., Tsipras, D., Vladu, A.: Towards deep learning models resistant to adversarial attacks. In: International Conference on Learning Representations (2018)
11. Madry, A., Makelov, A., Schmidt, L., Tsipras, D., Vladu, A.: Towards deep learning models resistant to adversarial attacks (2019)
12. Pauli, P., Koch, A., Berberich, J., Kohler, P., Allgower, F.: Training robust neural networks using Lipschitz bounds. *IEEE Control Syst. Lett.* (2021)
13. Raghuathan, A., Xie, S.M., Yang, F., Duchi, J.C., Liang, P.: Adversarial training can hurt generalization. arXiv preprint [arXiv:1906.06032](https://arxiv.org/abs/1906.06032) (2019)
14. Shorten, C., Khoshgoftaar, T.M.: A survey on image data augmentation for deep learning. *J. Big Data* **6**, 60 (2019)
15. Singh, G., Gehr, T., Püschel, M., Vechev, M.T.: An abstract domain for certifying neural networks. *PACMPL* **3**(POPL), 41:1–41:30 (2019). <https://doi.org/10.1145/3290354>
16. Stallkamp, J., Schlipsing, M., Salmen, J., Igel, C.: The German traffic sign recognition benchmark: a multi-class classification competition. In: IEEE International Joint Conference on Neural Networks, pp. 1453–1460 (2011)

17. Szegedy, C., et al.: Intriguing properties of neural networks. In: 2nd International Conference on Learning Representations, ICLR 2014 (2014)
18. Tsipras, D., Santurkar, S., Engstrom, L., Turner, A., Madry, A.: Robustness may be at odds with accuracy. In: International Conference on Learning Representations (2018)
19. Xiao, H., Rasul, K., Vollgraf, R.: Fashion-MNIST: a novel image dataset for benchmarking machine learning algorithms (2017)
20. Xu, J., Zhang, Z., Friedman, T., Liang, Y., den Broeck, G.V.: A semantic loss function for deep learning with symbolic knowledge. In: Dy, J.G., Krause, A. (eds.) Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, 10–15 July 2018. Proceedings of Machine Learning Research, vol. 80, pp. 5498–5507. PMLR (2018)
21. Zhang, H., Yu, Y., Jiao, J., Xing, E., El Ghaoui, L., Jordan, M.: Theoretically principled trade-off between robustness and accuracy. In: International Conference on Machine Learning, pp. 7472–7482. PMLR (2019)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

