
Globally-Robust Neural Networks

Klas Leino¹ Zifan Wang¹ Matt Fredrikson¹

Abstract

The threat of adversarial examples has motivated work on training *certifiably robust* neural networks, to facilitate efficient verification of *local robustness* at inference time. We formalize a notion of *global robustness*, which captures the operational properties of on-line local robustness certification while yielding a natural learning objective for robust training. We show that widely-used architectures can be easily adapted to this objective by incorporating efficient global Lipschitz bounds into the network, yielding certifiably-robust models *by construction* that achieve state-of-the-art verifiable and clean accuracy. Notably, this approach requires significantly less time and memory than recent certifiable training methods, and leads to negligible costs when certifying points on-line; for example, our evaluation shows that it is possible to train a large tiny-imagenet model in a matter of hours. We posit that this is possible using inexpensive global bounds—despite prior suggestions that tighter local bounds are needed for good performance—because these models are trained to achieve tighter global bounds. Namely, we prove that the maximum achievable verifiable accuracy for a given dataset is not improved by using a local bound. An implementation of our approach is available on GitHub².

1. Introduction

We consider the problem of training neural networks that are robust to input perturbations with bounded ℓ_p norm. Precisely, given an input point, x , network, F , and norm bound, ϵ , this means that F makes the same prediction on all points within the ℓ_p -ball of radius ϵ centered at x .

This problem is significant as deep neural networks have been shown to be vulnerable to *adversarial examples* (Paper-

¹Carnegie Mellon University, Pittsburgh, Pennsylvania, USA. Correspondence to: Klas Leino <kleino@cs.cmu.edu>.

²Code made publicly available at <https://github.com/klasleino/gloro>

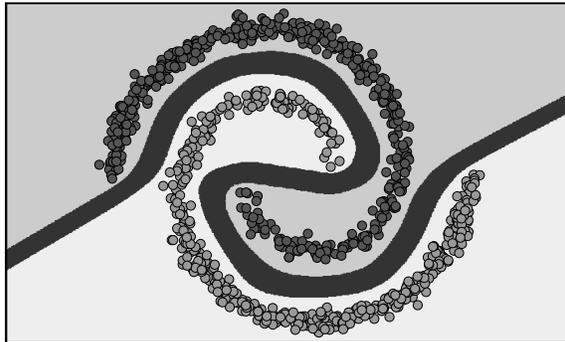


Figure 1: Illustration of global robustness. The model abstains from predicting on the margin between the classes (dark gray), which has width at least ϵ .

not et al., 2016; Szegedy et al., 2014), wherein perturbations are chosen to deliberately cause misclassification. While numerous heuristic solutions have been proposed to address this problem, without providing rigorous *certificates* of robustness for new test points, these solutions are often shown to be ineffective by subsequent adaptive attacks (Carlini & Wagner, 2017). Thus, this paper focuses on training methods that produce models whose robust predictions can be efficiently certified (Wong et al., 2018; Lee et al., 2020; Tsuzuku et al., 2018).

We begin by introducing a notion of *global robustness* for classification models (Section 2.1), which requires that classifiers maintain a separation of width at least ϵ in feature space between any pair of regions that are assigned different prediction labels. This separation means that there are certain inputs on which a globally-robust classifier must refuse to give a prediction, instead signaling that a violation has occurred (see Figure 1). While requiring that the model fail in some cases may at first appear to be a hindrance, we note that in operational terms this is no different than composing a model with a routine that only returns predictions when $\epsilon/2$ -local robustness can be certified.

While it is straightforward to construct a globally-robust model in this way via composition with a certification procedure, doing so with most current certification methods leads to severe penalties on performance and utility. Most techniques for verifying local robustness are costly even on small models (Fromherz et al., 2021; Gehr et al., 2018;

Tjeng & Tedrake, 2017; Fischetti & Jo, 2018; Jordan et al., 2019), requiring several orders of magnitude more time than a typical forward pass of a network; on moderately-large CNNs, these techniques either time out after minutes or hours, or simply run out of memory.

One approach to certification that shows promise in this regard uses Lipschitz bounds to efficiently calculate the robustness region around a point (Zhang et al., 2018; Weng et al., 2018a). In particular, when *global* bounds are used with this approach, it is possible to implement the bound computation as a neural network of comparable size to the original (Section 4.2), making on-line certification nearly as efficient as inference. Unfortunately, current training methods do not produce models with sufficiently small global bounds for this to succeed. Recent work (Lee et al., 2020) explored the possibility of training networks with sufficiently small *local* bounds, but the training cost in time and memory remains prohibitive in many cases.

Surprisingly, we find that using global Lipschitz bounds for certification may not be as limiting as previously thought (Huster et al., 2018; Yang et al., 2020). We show that for any set of points that can be robustly-classified using a local Lipschitz bound, there exists a model whose global bound implies the same robust classification (Theorem 3). This motivates a new approach to certifiable training that makes exclusive use of global bounds (Section 2.2). Namely, we construct a globally-robust model that incorporates a global bound in its forward pass to define an additional “violation” class, and use standard training methods.

Focusing on the case of ℓ_2 -certifiable training, we show that this approach yields verified-robust accuracy (VRA) that matches or exceeds the state-of-the-art, while imposing little overhead during training and none during certification. For example, we find that we can achieve 52% VRA with a large robustness radius of $\epsilon = 1.58$ on MNIST, surpassing all prior approaches by at least four percentage points. We also match the state-of-the-art VRA on CIFAR-10, and scale to larger applications than previous work, such as Tiny-Imagenet (see Section 5).

To summarize, we provide a method for training certifiably-robust neural networks that is simple, fast, capable of running with limited memory, and that yields state-of-the-art verified accuracy. We prove that the potential of our approach is not hindered by its simplicity; rather, its simplicity is an asset — our empirical results demonstrate the many benefits it enjoys over more complicated methods.

2. Constructing Globally-Robust Networks

In this section we present our method for constructing globally-robust networks, which we will also refer to as *GloRo Nets*. We begin in Section 2.1 by formally introduc-

ing our notion of *global robustness*, after briefly covering the essential background and notation. We then show how to mathematically construct GloRo Nets in Section 2.2, and prove that our construction is globally robust.

2.1. Global Robustness

Let $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$ be a neural network that categorizes points into m different classes. Let F be the function representing the predictions of f , i.e., $F(x) = \operatorname{argmax}_i \{f_i(x)\}$.

F is said to be ϵ -*locally-robust* at point x if it makes the same prediction on all points in the ϵ -ball centered at x (Definition 1).

Definition 1. (*Local Robustness*) A model, F , is ϵ -locally-robust at point, x , with respect to norm, $\|\cdot\|$, if $\forall x'$,

$$\|x - x'\| \leq \epsilon \implies F(x) = F(x').$$

Most work on robustness verification has focused on this local robustness property; in this work, we present a natural notion of *global robustness*, which captures the operational properties of on-line local robustness certification.

Clearly, local robustness cannot be simultaneously satisfied at every point — unless the model is entirely degenerate, there will always exist points that are arbitrarily close to a decision boundary. Instead, we will introduce a global robustness definition that can be satisfied even on models with non-trivial behavior by using an additional class, \perp , that signals that a point cannot be certified as globally robust. At a high level, we can think of separating each of the classes with a margin of width at least ϵ in which the model always predicts \perp . In order to satisfy global robustness, we require that no two points at distance ϵ from one another are labeled with different non- \perp classes.

More formally, let us define the following relation (\triangleq): we will say that $c_1 \triangleq c_2$ if $c_1 = \perp$ or $c_2 = \perp$ or $c_1 = c_2$. Using this relation, we provide our formal notion of global robustness in Definition 2.

Definition 2. (*Global Robustness*) A model, F , is ϵ -globally-robust, with respect to norm, $\|\cdot\|$, if $\forall x_1, x_2$,

$$\|x_1 - x_2\| \leq \epsilon \implies F(x_1) \triangleq F(x_2).$$

An illustration of global robustness is shown in Figure 1.

2.2. Certified Globally-Robust Networks

Because of the threat posed by adversarial examples, and the elusiveness of such attacks against heuristic defenses (Carlini & Wagner, 2017), there has been a volume of previous work on verifying local robustness on previously unseen points. In this work, we shift our focus to global robustness

directly, resulting in a method for producing models that make predictions that are verifiably robust *by construction*.

Intuitively, we aim to instrument a model with an extra output, \perp , that labels a point as not locally-robust, such that the instrumented model predicts a non- \perp class *only if the point is locally-robust* (with respect to the original model). At a high level, we do this by ensuring that in order to avoid predicting \perp , the maximum output of f must surpass the other outputs by a sufficient margin. While this margin is measured in the output space, we can ensure it is sufficiently large to ensure local robustness by relating the output space to the input space via an upper bound on the model’s Lipschitz constant.

Suppose that K_i is an upper bound on the Lipschitz constant for f_i . I.e., for all x_1, x_2 , Equation 1 holds. Intuitively, K_i bounds the largest possible change in the logit output for class i per unit change in the model’s input.

$$\frac{|f_i(x_1) - f_i(x_2)|}{\|x_1 - x_2\|} \leq K_i \quad (1)$$

Let $y = f(x)$, and let $j = F(x)$, i.e., the class predicted on point x . Let $y_\perp = \max_{i \neq j} \{y_i + (K_i + K_j)\epsilon\}$. We then define the instrumented model, or GloRo Net, \bar{f}^ϵ , as follows: $\bar{f}_i^\epsilon(x) ::= y_i \forall i \in [m]$ and $\bar{f}_\perp^\epsilon(x) ::= y_\perp$; that is, \bar{f}^ϵ concatenates y_\perp with the output of f .

We show that the predictions of this GloRo Net, \bar{F}^ϵ , can be used to certify the predictions of the instrumented model, F : whenever \bar{F}^ϵ predicts a class that is not \perp , the prediction coincides with the prediction of F , and F is guaranteed to be locally robust at that point (Theorem 1).

Theorem 1. *If $\bar{F}^\epsilon(x) \neq \perp$, then $\bar{F}^\epsilon(x) = F(x)$ and F is ϵ -locally-robust at x .*

The proof of Theorem 1 is given in Appendix A.1.

Note that in this formulation, we assume that the predicted class, j , will decrease by the maximum amount within the ϵ -ball, while all other classes increase by their respective maximum amounts. This is a conservative assumption that guarantees local robustness; however, in practice, we can dispose of this assumption by instead calculating the Lipschitz constant of the margin by which the logit of the predicted class surpasses the other logits, $y_j - y_i$. The details of this tighter variant are presented in Appendix A.2, along with the corresponding correctness proof.

Notice that the GloRo Net, \bar{F}^ϵ , will always predict \perp on points that lie directly on the decision boundary. Moreover, any point that is within ϵ of the decision boundary will also be labeled as \perp by \bar{F}^ϵ . From this, it is perhaps clear that GloRo Nets achieve global robustness (Theorem 2).

Theorem 2. *$\bar{F}^{\epsilon/2}$ is ϵ -globally-robust.*

The proof of Theorem 2 is given in Appendix A.3.

3. Revisiting the Global Lipschitz Constant

The global Lipschitz constant gives a bound on the maximum rate of change in the network’s output over the entire input space. For the purpose of certifying robustness, it suffices to bound the maximum rate of change in the network’s output over any pair of points *within the ϵ -ball* centered at the point being certified, i.e., the *local* Lipschitz constant. Recent work has explored methods for obtaining upper bounds on the local Lipschitz constant (Weng et al., 2018a; Zhang et al., 2018; Lee et al., 2020); the construction of GloRo Nets given in Section 2 remains correct whether K represents a global or a local Lipschitz constant.

The advantage to using a local bound is, of course, that we may expect tighter bounds; after all, the local Lipschitz constant is no larger than the global Lipschitz constant. However, using a local bound also has its drawbacks. First, a local bound is typically more expensive to compute. In particular, a local bound always requires more memory, as each instance has its own bound, hence the required memory grows with the batch size. This in turn reduces the amount of parallelism that can be exploited when using a local bound, reducing the model’s throughput.

Furthermore, because the local Lipschitz constant is different for every point, it must be computed every time the network sees a new point. By contrast, the global bound can be computed in advance, meaning that verification via the global bound is essentially free. This makes the global bound advantageous, assuming that it can be effectively leveraged for verification.

It may seem initially that a local bound would have greater prospects for successful certification. *Local* Lipschitzness is sufficient for robustly classifying well-separated data (Yang et al., 2020), that is, global Lipschitzness is not necessary; meanwhile, global bounds on typical networks have been found to be prohibitively large (Weng et al., 2018a), while local bounds on in-distribution points may tend to be smaller on the same networks. However, the potential disadvantages of a global bound become less clear if the model is specifically trained to have a small global Lipschitz constant.

For example, GloRo Nets using a global Lipschitz constant will be penalized for incorrect predictions if the global Lipschitz constant is not sufficiently small to verify its predictions; therefore, the loss actively discourages any unnecessary steepness in the network function. In practice, this natural regularization of the global Lipschitz constant may serve to make the steepness of the network function more uniform, such that the global Lipschitz constant will be similar to the local Lipschitz constant.

We show that this is possible in theory, in that for any network for which local robustness can be verified on some set of points using the local Lipschitz constant, there exists a

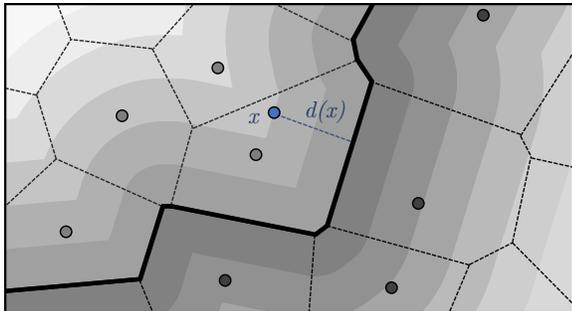


Figure 2: Illustration of a function, g , constructed to satisfy Theorem 3. The points in S are shown in light and dark gray, with different shades indicating different labels. The Voronoi tessellation is outlined in black, and the faces belonging to the decision boundary are highlighted in bold. The level curves of g are shown in various shades of gray and correspond to points, x , at some fixed distance, $d(x)$, from the decision boundary.

model on which the same points can be certified using the global Lipschitz constant (Theorem 3). This suggests that if training is successful, our approach has the same potential using a global bound as using a local bound.

Theorem 3. *Let f be a binary classifier that predicts $1 \iff f(x) > 0$. Let $K_L(x, \epsilon)$ be the local Lipschitz constant of f at point x with radius ϵ .*

Suppose that for some finite set of points, S , $\forall x \in S$, $|f(x)| > \epsilon K_L(x, \epsilon)$, i.e., all points in S can be verified via the local Lipschitz constant.

Then there exists a classifier, g , with global Lipschitz constant K_G , such that $\forall x \in S$, (1) g makes the same predictions as f on S , and (2) $|g(x)| > \epsilon K_G$, i.e., all points in S can be verified via the global Lipschitz constant.

Theorem 3 is stated for binary classifiers, though the result holds for categorical classifiers as well. Details on the categorical case and the proof of Theorem 3 can be found in Appendix A.4; however we provide the intuition behind the construction here. The proof relies on the following lemma, which states that among locally-robust points, points that are classified differently from one another are 2ϵ -separated. The proof of Lemma 1 can be found in Appendix A.5.

Lemma 1. *Suppose that for some classifier, F , and some set of points, S , $\forall x \in S$, F is ϵ -locally-robust at x . Then $\forall x_1, x_2 \in S$ such that $F(x_1) \neq F(x_2)$, $\|x_1 - x_2\| > 2\epsilon$.*

In the proof of Theorem 3, we construct a function, g , whose output on point x increases linearly with x 's minimum distance to any face in the Voronoi tessellation of S that separates points in S with different labels. An illustration with an example of g is shown in Figure A.1. Notably, the local Lipschitz constant of g is everywhere the same as the global

constant.

The function, g , produced in this construction is not piecewise linear; however, g can be approximated with arbitrary precision by a piecewise linear function — in particular, it can be approximated sufficiently well to satisfy Theorem 3 with a finite number of regions — suggesting that such a function is implementable by a ReLU network.

However, we note that while Theorem 3 suggests that networks exist *in principle* on which it is possible to use the global Lipschitz constant to certify 2ϵ -separated data, it may be that such networks are not easily obtainable via training. Furthermore, as raised by Huster et al. (2018), an additional potential difficulty in using the global bound for certification is the estimation of the global bound itself. Methods used for determining an upper bound on the global Lipschitz constant, such as the method presented later in Section 4.2, may provide a loose upper bound that is insufficient for verification even when the true bound would suffice. Nevertheless, our evaluation in Section 5 shows that in practice the global bound *can* be used effectively for certification (Section 5.1), and that the bounds obtained on the models trained with our approach are far tighter than those obtained on standard models (Section 5.3).

4. Implementation

In this section we describe how GloRo Nets can be trained and implemented. Section 4.1 covers training, as well as the loss functions used in our evaluation, and Section 4.2 provides detail on how we compute an upper bound on the global Lipschitz constant.

4.1. Training

Crucially, because certification is intrinsically captured by the GloRo Net's predictions — specifically, the \perp class represents inputs that cannot be verified to be locally robust — a standard learning objective for a GloRo Net corresponds to a robust objective on the original model that was instrumented. That is, we can train a GloRo Net by simply appending a zero to the one-hot encodings of the original data labels (signifying that \perp is never the correct label), and then optimizing a standard classification learning objective, e.g., cross-entropy. Using this approach, \bar{F}^ϵ will be penalized for incorrectly predicting each point, x , unless x is both predicted correctly and F is ϵ -locally-robust at x .

While the above approach is sufficient for training models with good VRAs both in theory and in practice, we find that the resulting VRA can be further improved using a loss inspired by TRADES (Zhang et al., 2019), which balances separately the goals of making *correct* predictions, and making *robust* predictions. Recent work (Yang et al., 2020) has shown that TRADES is effective at controlling

the local Lipschitz constant of networks. While TRADES is implemented using adversarial perturbations, which provide an under-approximation of the robust error, GloRo Nets naturally lend themselves to a variant of TRADES that uses an *over-approximation*, given by Definition 3.

Definition 3. (*TRADES Loss for GloRo Nets*) Given network function, f , loss function \mathcal{L} , and parameter, λ , the TRADES loss (\mathcal{L}_T) on labeled point, (x, y) , is

$$\mathcal{L}_T(x, y) = \mathcal{L}(f(x), y) + \lambda \mathcal{L}\left(\bar{f}^\epsilon(x), \text{softmax}(f(x))\right)$$

Intuitively, \mathcal{L}_T combines the loss of the original model’s predictions with the loss of the corresponding GloRo Net under the assumption that the class predicted by the underlying model is correct.

4.2. Bounding the Global Lipschitz Constant

There has been a great deal of work on calculating upper bounds on the Lipschitz constants of neural networks (see Section 6 for a discussion). Our implementation uses the fact that the product of the spectral norm of each of the individual layers of a feed-forward network provides an upper bound on the Lipschitz constant of the entire network (Szegedy et al., 2014). That is, if the output at class i of a neural network can be decomposed into a series of k linear transformations, i.e., $f_i = h^k \circ h^{k-1} \circ \dots \circ h^1$, then Equation 2 holds (where $\|\cdot\|$ is the spectral norm).

$$K_i \leq \prod_{j=1}^k \|h^j\| \quad (2)$$

In the case of a CNN consisting of convolutional layers, dense layers, and ReLU activations, we use 1 for the spectral norm of each of the ReLU layers, and we use the power method (Farnia et al., 2019; Gouk et al., 2021) to compute the spectral norm of the convolutional and dense layers. Gouk et al. also give a procedure for bounding the spectral norm of skip connections and batch normalization layers, enabling this approach on ResNet architectures. For more complicated networks, there is a growing body of work on computing layer-wise Lipschitz bounds for various types of layers that are commonly used in neural networks (Zou et al., 2019; Fazlyab et al., 2019; Sedghi et al., 2019; Singla & Feizi, 2019; Miyato et al., 2018).

The power method may need several iterations to converge; however, we can reduce the number of iterations required at each training step by persisting the state of the power method iterates across steps. While this optimization may not guarantee an upper bound, this fact is inconsequential so long as we still obtain a model that can be certified with a true upper bound; this is actually not unreasonable to expect, presuming the underlying model parameters do

not change too quickly. With a small number of iterations, the additional memory required to compute the Lipschitz constant via this method is approximately the same as to run the network on a single instance.

At test time, the power method must be run to convergence; however, after training, the global Lipschitz bound will remain unchanged and therefore it can be computed in advance.

ℓ_∞ Bounds. While in this work, we focus on the ℓ_2 norm, the ideas presented in Section 2 can be applied to other norms, including the ℓ_∞ norm. However, we find that the analogue of the approximation of the global Lipschitz bound given by Equation 2 is loose in ℓ_∞ space. Meanwhile, a large volume of prior work applies ℓ_∞ -specific certification strategies that proven effective for ℓ_∞ certification (Zhang et al., 2020; Balunovic & Vechev, 2020; Goyal et al., 2019).

5. Evaluation

In this section, we present an empirical evaluation of our method. We first compare GloRo Nets with several certified training methods from the recent literature in Section 5.1. We also report the training cost, in terms of per-epoch time and peak memory usage, required to train and certify the robustness of our method compared with other competitive approaches (Section 5.2). We end by demonstrating the relative tightness of the estimated Lipschitz bounds for GloRo Nets in Section 5.3. We compare against the KW (Wong et al., 2018) and BCP (Lee et al., 2020) certified training algorithms, which prior work (Lee et al., 2020; Croce et al., 2019) reported to achieve the best verified accuracy on MNIST (LeCun et al., 2010), CIFAR-10 (Krizhevsky, 2009) and Tiny-Imagenet (Le & Yang, 2015) relative to other certified training methods for ℓ_2 robustness.

We train GloRo nets to certify robustness against ℓ_2 perturbations within an ϵ -neighborhood of 0.3 and 1.58 for MNIST and 36/255 for CIFAR-10 and Tiny-Imagenet (these are the ℓ_2 norm bounds that have been commonly used in the previous literature). For each model, we report the *clean accuracy*, i.e., the accuracy without verification on non-adversarial inputs, the *PGD accuracy*, i.e., the accuracy under adversarial perturbations found via the PGD attack (Madry et al., 2018), and the *verified-robust accuracy* (VRA), i.e., the fraction of points that are both correctly classified *and* certified as robust. For KW and BCP, we report the corresponding best VRAs from the original respective papers when possible, but measure training and certification costs on our hardware for an equal comparison. We run the PGD attacks using ART (Nicolae et al., 2019) on our models and on any of the models from the prior work for which PGD accuracy is not reported. When training BCP models for MNIST with $\epsilon = 0.3$, we found a differ-

Globally-Robust Neural Networks

<i>method</i>	Model	Clean (%)	PGD (%)	VRA(%)	Sec./epoch	# Epochs	Mem. (MB)
MNIST ($\epsilon = 0.3$)							
Standard	2C2F	99.2	96.9	0.0	0.4	100	0.7
GloRo	2C2F	98.4	96.9	94.6	3.2	500	0.9
GloRo-T	2C2F	98.7	97.4	94.6	3.3	500	0.9
KW	2C2F	98.9	97.8	94.0	66.9	100	20.2
BCP	2C2F	93.4	89.5	84.7	44.8	300	12.6
MNIST ($\epsilon = 1.58$)							
Standard	4C3F	99.0	45.4	0.0	1.1	42*	1.6
GloRo	4C3F	92.9	68.9	50.1	9.0	300	1.7
GloRo-T	4C3F	92.8	67.0	51.9	7.2	300	1.7
KW	4C3F	88.1	67.9	44.5	138.1	60	84.0
BCP	4C3F	92.4	65.8	47.9	43.4	60	12.6
CIFAR-10 ($\epsilon = 36/255$)							
Standard	6C2F	85.7	31.9	0.0	2.58	115*	2.9
GloRo	6C2F	70.7	63.8	49.3	6.5	800	3.0
GloRo-T	6C2F	67.9	61.3	51.0	7.1	800	3.0
KW	6C2F	60.1	56.2	50.9	516.8	60	100.9
BCP	6C2F	65.7	60.8	51.3	47.5	200	12.7
Tiny-Imagenet ($\epsilon = 36/255$)							
Standard	7C1F	35.9	19.4	0.0	89.2	58*	16.4
GloRo	7C1F	31.3	28.2	13.2	107.2	250	31.5
GloRo-T	7C1F	27.4	25.6	15.6	112.3	500	31.5
KW	-	-	-	-	-	-	-
BCP ³	-	-	-	-	-	-	-

(a)

<i>method</i>	Model	Time (sec.)	Mem. (MB)
GloRo	6C2F	0.8	5.9
KW	6C2F	2515.6	1437.5
BCP	6C2F	5.8	19.1

(b)

<i>method</i>	global UB	global LB	local LB
MNIST ($\epsilon = 1.58$)			
Standard	$5.4 \cdot 10^4$	$1.4 \cdot 10^2$	17.1
GloRo	3.2	3.0	2.1
CIFAR-10 ($\epsilon = 36/255$)			
Standard	$1.2 \cdot 10^7$	$1.1 \cdot 10^3$	96.2
GloRo	18.9	11.4	6.2
Tiny-Imagenet ($\epsilon = 36/255$)			
Standard	$2.2 \cdot 10^7$	$3.6 \cdot 10^2$	40.7
GloRo	7.7	3.3	1.5

(c)

Figure 3: **(a)** Certifiable training evaluation results on benchmark datasets. Best results highlighted in bold. “GloRo-T” indicates that TRADES loss (Definition 3) was used rather than crossentropy. A * superscript on the number of epochs denotes that an early-stop callback was used to determine convergence. **(b)** Certification timing and memory usage results on CIFAR-10 ($\epsilon = 36/255$). **(c)** Upper and lower bounds on the global and average local Lipschitz constant. In (a) and (b), peak GPU Memory usage is calculated per-instance by dividing the total measurement by the training or certification batch size.

ent set of hyperparameters that outperforms those given by Lee et al.. Further details on the precise hyperparameters used for training, attacks, the process for obtaining these parameters, and the network architectures are provided in Appendix B.

5.1. Verified Accuracy

We first compare the VRA obtained by GloRo Nets to the VRA achieved by prior approaches. KW and BCP have been found to achieve the best VRA on the datasets commonly used in the previous literature (Lee et al., 2020; Croce et al., 2019). Figure 3a gives the best VRA achieved by standard training, GloRo Nets (with both TRADES and crossentropy loss), KW, and BCP on several benchmark datasets and architectures. In accordance with prior work, we also include the clean accuracy and the PGD accuracy. Whereas the VRA gives a lower bound on the number of correctly-classified points that are locally robust, the PGD accuracy serves as an upper bound on the same quantity.

We find that GloRo Nets (those trained with TRADES in particular) consistently outperform or match the state-of-the-art

VRA. On MNIST, GloRo Nets always outperform all previous approaches with both ℓ_2 bounds commonly used in prior work ($\epsilon = 0.3$ and $\epsilon = 1.58$). When $\epsilon = 0.3$, the VRA begins to approach the clean accuracy of the standard-trained model; for this bound GloRo Nets outperform the previous best VRA (achieved by KW) by over half a percentage point, accounting for roughly 15% of the gap between the VRA of KW and the clean accuracy of the standard model. For $\epsilon = 1.58$, GloRo Nets improve upon the previous best VRA (achieved by BCP) by approximately 4 percentage points. On CIFAR-10, GloRo Nets come within a fraction of a percentage point from the best VRA (achieved by BCP), and slightly outperform the VRA of KW, showing that they remain highly competitive with the state-of-the-art.

Furthermore, we find that GloRo Nets may be the first approach to scale to training a non-degenerate certifiable model on Tiny-Imagenet³ (15.6% VRA); other approaches have been reported to either produce degenerate models

³ Lee et al. (2020) report achieving a VRA of 20.08% on Tiny-Imagenet. However, we are unable to reproduce this result with their code, and are currently in contact with the authors attempting to resolve the issue. More details are provided in Appendix C.

or run out of memory (Lee et al., 2020). We believe this is an important stride towards increasing the scalability of robustness certification to more practical applications.

As pointed out in Section 4.1, we find that TRADES loss gives an additional small boost to the accuracy of GloRo Nets, making it a practical training enhancement.

5.2. Training and Certification Cost

A key advantage to GloRo Nets over prior approaches is their ability to achieve state-of-the-art VRA (see Section 5.1) using a global Lipschitz bound. As discussed in Section 3, this confers performance benefits — both at train and test time — over using a local bound (e.g., BCP), or other expensive approaches (e.g., KW).

Figure 3a shows the cost of each approach both in time per epoch and in memory during training (results given for CIFAR-10). All timings were taken on a machine using a Geforce RTX 3080 accelerator, 64 GB memory, and Intel i9 10850K CPU, with the exception of those for the KW (Wong et al., 2018) method, which were taken on a Titan RTX card for toolkit compatibility reasons. Appendix D provides more details on how memory usage was measured. Because different batch sizes were used to train and evaluate each model, we control for this by reporting the memory used *per instance in each batch*. The cost for standard training is included for comparison.

We see that KW is the most expensive approach, requiring tens to hundreds of seconds per epoch and roughly $30\times$ more memory per batch instance than standard training. BCP is less expensive than KW, but still takes nearly one minute per epoch and anywhere between $4\text{--}18\times$ more memory than standard training. Meanwhile, the cost of GloRo Nets is more comparable to standard training than to KW or BCP, taking only a few seconds per epoch, and approximately the same amount of memory, except on Tiny-Imagenet, where it used roughly $2\times$ more memory than standard training. Because of its memory scalability, we were able to use a larger batch size with GloRo Nets. As a result, more epochs were required during training however, this did not outweigh the significant reduction in time per epoch, as the total time for training was still only at most half of the total time for BCP.

Figure 3b shows the cost of each approach both in the time required to certify the entire test set and in the memory used to do so (results given for CIFAR-10). KW is the most expensive in terms of time and memory, followed by BCP. And again, GloRo Nets are far superior in terms of cost, making certified predictions over $7\times$ faster than BCP with less than a third of the memory, and over $3,000\times$ faster than KW. We thus conclude that GloRo Nets are the most scalable state-of-the-art technique for robustness certification.

5.3. Lipschitz Tightness

Theorem 3 demonstrates that a global Lipschitz bound is theoretically sufficient for certifying 2ϵ -separated data. However, as discussed in Section 3, there may be several practical limitations making it difficult to realize a network satisfying Theorem 3; we now assess how these limitations are borne out in practice by examining the Lipschitz bounds that GloRo Nets use for certification.

Weng et al. (2018a) report that an upper bound on the global Lipschitz constant is not capable of certifying robustness for a non-trivial radius, ϵ . While this is true of models produced via standard training, GloRo Nets impose a strong implicit regularization on the global Lipschitz constant. Indeed, Figure 3c shows that the global upper bound is several orders of magnitude smaller on GloRo Nets than on standard networks.

Another potential limitation of using an upper bound of the global Lipschitz constant is the bound itself (Huster et al., 2018). Figure 3c shows that a lower bound of the Global Lipschitz constant, obtained via optimization, reaches an impressive 92% of the upper bound on MNIST, meaning that the upper bound is nearly tight. On CIFAR-10 and Tiny-Imagenet the lower bound reaches approximately 60% and 40% of the upper bound, respectively. However, on a standard model, the lower bound is potentially orders of magnitude looser. These results show there is still room for improvement, though the fact that the bound is tighter for GloRo Nets suggests the objective of the GloRo Net helps by incentivizing parameters for which the upper bound estimate is sufficiently tight for verification. Interestingly though, artificially changing the bound used by the GloRo Net to match the empirical lower bound on CIFAR-10 reveals that the existing looseness accounts for a drop in accuracy of at most 2 percentage points.

Finally, we compare the global upper bound to an empirical lower bound of the local Lipschitz constant. The local lower bound given in Figure 3c reports the mean local Lipschitz constant found via optimization in the ϵ -balls centered at each of the test points. In the construction given for the proof of Theorem 3, the local Lipschitz constant is the same as the global bound at all points. While the results in Figure 3c show that this may not be entirely achieved in practice, the ratio of the local lower bound to the global upper bound is essentially zero in the standard models, compared to 20–66% in the GloRo Nets, establishing that the upper bound is again much tighter for GloRo Nets. Still, this suggests that a reasonably tight estimate of the local bound may yet help improve the VRA of a GloRo Net at runtime, although this is a challenge in its own right. Intriguingly, the fact that GloRo Nets are already competitive in VRA with BCP means that any improvement to the VRA of GloRo Nets obtained via the local bound used by BCP would cause GloRo Nets

to outperform BCP, suggesting that GloRo Nets provide a better objective for certifiably-robust training despite using a looser bound during training.

We provide further discussion of the upper and lower bounds, and details for how the lower bounds were obtained in Appendix E.

6. Related Work

Utilizing the Lipschitz constant to certify robustness has been studied in several instances of prior work. On discovering the existence of adversarial examples, Szegedy et al. (2014) analyzed the sensitivity of neural networks using a global Lipschitz bound, explaining models’ “blind spots” partially in terms of large bounds and suggesting Lipschitz regularization as a potential remedy. Huster et al. (2018) noted the potential limitations of using global bounds computed layer-wise according to Equation 2, and showed experimentally that direct regularization of the Lipschitz constant by penalizing the weight norms of a two-layer network yields subpar results on MNIST. While Theorem 3 does not negate their concern, as it may not always be feasible to compute a tight enough bound using Equation 2, our experimental results show to the contrary that global bounds can suffice to produce models with at least comparable utility to several more expensive, and complicated, techniques. More recently, Yang et al. (2020) showed that robustness and accuracy need not be at odds on common benchmarks when locally-Lipschitz functions are used, and call for further investigation of methods that impose this condition while promoting generalization. Our results show that globally-Lipschitz functions, which bring several practical benefits (Section 3), are a promising direction as well.

Lipschitz constants have been applied previously for fast post-hoc certification (Weng et al., 2018a; Hein & Andriushchenko, 2017; Weng et al., 2018b). While our work relies on similar techniques, our exclusive use of the global bound means that no additional work is needed at inference time. Additionally, we apply this certification only to networks that have been optimized for it.

There has also been prior work seeking to use Lipschitz bounds, or close analogues, during training to promote robustness (Tsuzuku et al., 2018; Raghunathan et al., 2018; Cisse et al., 2017; Cohen et al., 2019; Anil et al., 2019; Pauli et al., 2021; Qin et al., 2019; Finlay & Oberman, 2019; Lee et al., 2020; Gouk et al., 2021; Singla & Feizi, 2019; Farnia et al., 2019). Cisse et al. (2017) introduced Perseval networks, which enforce contractive Lipschitz bounds on all layers by orthonormalizing their weights. Anil et al. (2019) proposed replacing ReLU activations with sorting activations to construct a class of *universal Lipschitz approximators*, that that can approximate any Lipschitz-bounded

function over a given domain, and Cohen et al. (2019) subsequently studied the application to robust training; these advances in architecture design may complement our work, and deserve further study.

The closest work in spirit to ours is Lipschitz Margin Training (LMT) (Tsuzuku et al., 2018), which also uses global Lipschitz bounds to train models that are more certifiably robust. The approach works by constructing a loss that adds $\sqrt{2}\epsilon K_G$ to all logits other than that corresponding to the ground-truth class. Note that this is different from GloRo Nets, which add a *new logit* defined by the *predicted* class at x . In addition to providing different gradients than those of LMT’s loss, our approach avoids penalizing logits corresponding to boundaries distant from x . In practical terms, Lee et al. (2020) showed that LMT yields lower verified accuracy than more recent methods that use local Lipschitz bounds (Lee et al., 2020) or dual networks (Wong et al., 2018), while Section 5.1 shows that our approach can provide greater verified accuracy than either. LMT’s use of global bounds means its cost is comparable to our approach.

More recently, Lee et al. (2020) explored the possibility of training networks against local Lipschitz bounds, motivated by the fact that the global bound may vastly exceed a typical local bound on some networks. They showed that a localized refinement of the global spectral norm of the network offers a reasonable tradeoff of precision for cost, and were able to achieve competitive, and in some cases superior, verified accuracy to prior work. Theorem 3 shows that in principle, the difference in magnitude between local and global bounds may not matter for robust classification. Moreover, while it is true that the bounds computed by Equation 2 may be loose on some models, our experimental results suggest that it is possible in many cases to mitigate this limitation by training against a global bound with the appropriate loss. The advantages of doing so are apparent in the cost of both training and certification, where the additional overhead involved with computing tighter local bounds is an impediment to scalability.

Finally, several other methods have been proposed for training ℓ_2 -certifiable networks that are not based on Lipschitz constants. For example, Wong & Kolter (2018) use an LP-based approach that can be optimized relatively efficiently using a *dual network*, Croce et al. (2019) and Madry et al. (2018) propose training routines based on maximizing the size of the linear regions within a network, and Mirman et al. (2018) propose a method based on abstract interpretation.

7. Conclusion

In this work, we provide a method for training certifiably-robust neural networks that is simple, fast, memory-efficient, and that yields state-of-the-art verified accuracy. Our ap-

proach is particularly efficient because of its effective use of global Lipschitz bounds, and while we prove that the potential of our approach is in theory not limited by the global Lipschitz constant itself, it remains an open question as to whether these bounds can be tightened, or if additional training techniques can help unlock its remaining potential. Finally, we note that if instances arise where a global bound is not sufficient *in practice*, costlier post-hoc certification techniques may be complimentary, as a fall-back.

Acknowledgments. The work described in this paper has been supported by the Software Engineering Institute under its FFRDC Contract No. FA8702-15-D-0002 with the U.S. Department of Defense, as well as DARPA and the Air Force Research Laboratory under agreement number FA8750-15-2-0277.

References

- Anil, C., Lucas, J., and Gross, R. Sorting out Lipschitz function approximation. In *ICML*, 2019.
- Balunovic, M. and Vechev, M. Adversarial training and provable defenses: Bridging the gap. In *ICLR*, 2020.
- Carlini, N. and Wagner, D. Adversarial examples are not easily detected: Bypassing ten detection methods. In *ACM Workshop on Artificial Intelligence and Security*, 2017.
- Cisse, M., Bojanowski, P., Grave, E., Dauphin, Y., and Usunier, N. Parseval networks: Improving robustness to adversarial examples. In *ICML*, 2017.
- Cohen, J. E., Huster, T., and Cohen, R. Universal lipschitz approximation in bounded depth neural networks. *arXiv preprint arXiv:1904.04861*, 2019.
- Croce, F., Andriushchenko, M., and Hein, M. Provable robustness of ReLU networks via maximization of linear regions. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2019.
- Farnia, F., Zhang, J., and Tse, D. Generalizable adversarial training via spectral normalization. In *ICLR*, 2019.
- Fazlyab, M., Robey, A., Hassani, H., Morari, M., and Pappas, G. J. Efficient and accurate estimation of lipschitz constants for deep neural networks. In *NIPS*, 2019.
- Finlay, C. and Oberman, A. M. Scaleable input gradient regularization for adversarial robustness. *CoRR*, abs/1905.11468, 2019.
- Fischetti, M. and Jo, J. Deep neural networks and mixed integer linear optimization. *Constraints*, 23(3):296–309, Jul 2018.
- Fromherz, A., Leino, K., Fredrikson, M., Parno, B., and Păsăreanu, C. Fast geometric projections for local robustness certification. In *ICLR*, 2021.
- Gehr, T., Mirman, M., Drachler-Cohen, D., Tsankov, P., Chaudhuri, S., and Vechev, M. AI2: Safety and robustness certification of neural networks with abstract interpretation. In *Symposium on Security and Privacy (S&P)*, 2018.
- Gouk, H., Frank, E., Pfahringer, B., and Cree, M. J. Regularisation of neural networks by enforcing lipschitz continuity. *Machine Learning*, 110(2):393–416, Feb 2021.
- Gowal, S., Dvijotham, K. D., Stanforth, R., Bunel, R., Qin, C., Uesato, J., Arandjelovic, R., Mann, T., and Kohli, P. Scalable verified training for provably robust image classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019.
- Hein, M. and Andriushchenko, M. Formal guarantees on the robustness of a classifier against adversarial manipulation. In *NIPS*, 2017.
- Huster, T., Chiang, C.-Y. J., and Chadha, R. Limitations of the lipschitz constant as a defense against adversarial examples. In *ECML PKDD*, 2018.
- Jordan, M., Lewis, J., and Dimakis, A. G. Provable certificates for adversarial examples: Fitting a ball in the union of polytopes. In *NIPS*, 2019.
- Krizhevsky, A. Learning multiple layers of features from tiny images. Technical report, 2009.
- Le, Y. and Yang, X. Tiny imagenet visual recognition challenge. 2015.
- LeCun, Y., Cortes, C., and Burges, C. MNIST handwritten digit database. 2010. URL <http://yann.lecun.com/exdb/mnist>.
- Lee, S., Lee, J., and Park, S. Lipschitz-certifiable training with a tight outer bound. In *NIPS*, 2020.
- Madry, A., Makelov, A., Schmidt, L., Tsipras, D., and Vladu, A. Towards deep learning models resistant to adversarial attacks. In *ICLR*, 2018.
- Mirman, M., Gehr, T., and Vechev, M. Differentiable abstract interpretation for provably robust neural networks. In *ICML*, 2018.
- Miyato, T., Kataoka, T., Koyama, M., and Yoshida, Y. Spectral normalization for generative adversarial networks. In *ICLR*, 2018.

- Nicolae, M.-I., Sinn, M., Tran, M. N., Buesser, B., Rawat, A., Wistuba, M., Zantedeschi, V., Baracaldo, N., Chen, B., Ludwig, H., Molloy, I. M., and Edwards, B. Adversarial robustness toolbox v1.0.0, 2019.
- Papernot, N., McDaniel, P., Jha, S., Fredrikson, M., Celik, Z. B., and Swami, A. The limitations of deep learning in adversarial settings. In *European Symposium on Security and Privacy (EuroS&P)*, 2016.
- Pauli, P., Koch, A., Berberich, J., Kohler, P., and Allgower, F. Training robust neural networks using lipschitz bounds. *IEEE Control Systems Letters*, 2021.
- Qin, C., Martens, J., Goyal, S., Krishnan, D., Dvijotham, K., Fawzi, A., De, S., Stanforth, R., and Kohli, P. Adversarial robustness through local linearization. In *NIPS*, 2019.
- Raghunathan, A., Steinhardt, J., and Liang, P. Certified defenses against adversarial examples. 2018.
- Sedghi, H., Gupta, V., and Long, P. M. The singular values of convolutional layers. In *ICLR*, 2019.
- Singla, S. and Feizi, S. Bounding singular values of convolution layers. *CoRR*, abs/1911.10258, 2019.
- Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I. J., and Fergus, R. Intriguing properties of neural networks. In *ICLR*, 2014.
- Tjeng, V. and Tedrake, R. Verifying neural networks with mixed integer programming. *CoRR*, abs/1711.07356, 2017.
- Tsuzuku, Y., Sato, I., and Sugiyama, M. Lipschitz-margin training: Scalable certification of perturbation invariance for deep neural networks. In *NIPS*, 2018.
- Weng, T.-W., Zhang, H., Chen, H., Song, Z., Hsieh, C.-J., Boning, D., Dhillon, I. S., and Daniel, L. Towards fast computation of certified robustness for relu networks. In *ICML*, 2018a.
- Weng, T.-W., Zhang, H., Chen, P.-Y., Yi, J., Su, D., Gao, Y., Hsieh, C.-J., and Daniel, L. Evaluating the robustness of neural networks: An extreme value theory approach. In *ICLR*, 2018b.
- Wong, E. and Kolter, J. Z. Provable defenses against adversarial examples via the convex outer adversarial polytope. In *ICML*, 2018.
- Wong, E., Schmidt, F., Metzen, J. H., and Kolter, J. Z. Scaling provable adversarial defenses. In *NIPS*, 2018.
- Yang, Y.-Y., Rashtchian, C., Zhang, H., Salakhutdinov, R. R., and Chaudhuri, K. A closer look at accuracy vs. robustness. In *NIPS*, 2020.
- Zhang, H., Weng, T.-W., Chen, P.-Y., Hsieh, C.-J., and Daniel, L. Efficient neural network robustness certification with general activation functions. In *NIPS*, 2018.
- Zhang, H., Yu, Y., Jiao, J., Xing, E., Ghaoui, L. E., and Jordan, M. Theoretically principled trade-off between robustness and accuracy. In *ICML*, 2019.
- Zhang, H., Chen, H., Xiao, C., Goyal, S., Stanforth, R., Li, B., Boning, D., and Hsieh, C.-J. Towards stable and efficient training of verifiably robust neural networks. In *ICLR*, 2020.
- Zou, D., Balan, R., and Singh, M. On lipschitz bounds of general convolutional neural networks. *IEEE Transactions on Information Theory*, 66(3):1738–1759, 2019.

A. Proofs

A.1. Proof of Theorem 1

Theorem 1. *If $\bar{F}^\epsilon(x) \neq \perp$, then $\bar{F}^\epsilon(x) = F(x)$ and F is ϵ -locally-robust at x .*

Proof. Let $j = F(x)$. Assume that $\bar{F}^\epsilon(x) \neq \perp$; this happens only if one of the outputs of f is greater than $f_\perp^\epsilon(x)$ — from the definition of $f_\perp(x)$, it is clear that only $f_j(x)$ can be greater than $f_\perp^\epsilon(x)$. Therefore $f_j(x) > f_\perp^\epsilon(x)$, and so $\bar{F}^\epsilon(x) = j = F(x)$.

Now assume x' satisfies $\|x - x'\| \leq \epsilon$. Let K_i be an upper bound on the Lipschitz constant of f_i . Then, $\forall i$

$$\begin{aligned} \frac{|f_i(x) - f_i(x')|}{\epsilon} &\leq \frac{|f_i(x) - f_i(x')|}{\|x - x'\|} \leq K_i \\ \implies |f_i(x) - f_i(x')| &\leq K_i \epsilon \end{aligned} \quad (3)$$

We proceed to show that for any such x' , $F(x')$ is also j . In other words, $\forall i \neq j$, $f_i(x') < f_j(x')$. By applying the definition of the Lipschitz constant as in (3), we obtain (4). Next, (5) follows from the fact that $f_\perp^\epsilon(x) = \max_{i \neq j} \{y_i + (K_i + K_j)\epsilon\}$. We then obtain (6) from the fact that $f_j(x) > f_\perp^\epsilon(x)$, as observed above. Finally, we again apply (3) to obtain (7).

$$f_i(x') \leq f_i(x) + |f_i(x) - f_i(x')| \leq f_i(x) + K_i \epsilon \quad (4)$$

$$\leq f_\perp^\epsilon(x) - K_j \epsilon \quad (5)$$

$$< f_j(x) - K_j \epsilon \quad (6)$$

$$\leq f_j(x) - |f_j(x) - f_j(x')| \leq f_j(x') \quad (7)$$

Therefore, $f_i(x') < f_j(x')$, and so $F(x') = j$. This means that F is locally robust at x . \square

A.2. Tighter Bounds for Theorem 1

Note that in the formulation of GloRo Nets given in Section 2.2, we assume that the predicted class, j , will decrease by the maximum amount within the ϵ -ball, while all other classes increase by their respective maximum amounts. This is a conservative assumption that guarantees local robustness; however, in practice, we can dispose of this assumption by instead calculating the Lipschitz constant of the margin by which the logit of the predicted class surpasses the other logits, $f_j - f_i$.

The *margin Lipschitz constant* of f , defined for a pair of classes, $i \neq j$, is given by Definition 4.

Definition 4. *Margin Lipschitz Constant For network, $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$, and classes $i \neq j \in [m]$, K_{ij}^* is an upper bound on the margin Lipschitz constant of f if $\forall x_1, x_2$,*

$$\frac{|f_j(x_1) - f_i(x_1) - (f_j(x_2) - f_i(x_2))|}{\|x_1 - x_2\|} \leq K_{ij}^*$$

We now define a variant of GloRo Nets (Section 2.2) as follows: For input, x , let $j = F(x)$, i.e., j is the label assigned by the underlying model to be instrumented. Define $\bar{f}_i^\epsilon(x) ::= f_i(x)$, and $\bar{f}_\perp^\epsilon(x) ::= \max_{i \neq j} \{f_i(x) + \epsilon K_{ij}^*\}$.

Theorem 4. *Under this variant, if $\bar{F}^\epsilon(x) \neq \perp$, then $\bar{F}^\epsilon(x) = F(x)$ and F is ϵ -locally-robust at x .*

Proof. The proof is similar to the proof of Theorem 1 (Appendix A.1). Let $j = F(x)$. As before, when $\bar{F}^\epsilon(x) \neq \perp$, we see that $\bar{F}^\epsilon(x) = j = F(x)$.

Now assume x' satisfies $\|x - x'\| \leq \epsilon$. Let K_{ij}^* be an upper bound on the margin Lipschitz constant. Then, $\forall i$

$$|f_j(x) - f_i(x) - (f_j(x') - f_i(x'))| \leq K_{ij}^* \epsilon \quad (8)$$

We proceed to show that for any such x' , $F(x')$ is also j . In other words, $\forall i \neq j$, $f_i(x') < f_j(x')$. By applying (8), we obtain (9). Next, (10) follows from the fact that $f_\perp^\epsilon(x) = \max_{i \neq j} \{f_i(x) + K_{ij}^* \epsilon\}$. We then obtain (11) from the fact that $f_j(x) > \bar{f}_\perp^\epsilon(x)$, as $\bar{F}^\epsilon(x) = j \neq \perp$.

$$\begin{aligned} &\cancel{f_i(x)} + f_j(x) - \cancel{f_i(x)} - f_j(x') + f_i(x') \\ &\leq f_i(x) + |f_j(x) - f_i(x) - (f_j(x') - f_i(x'))| \\ &\leq f_i(x) + K_{ij}^* \epsilon \end{aligned} \quad (9)$$

$$\leq \bar{f}_\perp^\epsilon(x) \quad (10)$$

$$< f_j(x) \quad (11)$$

Rearranging terms, we obtain that $f_i(x') < f_j(x')$. Thus, $F(x') = j$; this means that F is locally robust at x . \square

A.3. Proof of Theorem 2

Theorem 2. *$\bar{F}^{\epsilon/2}$ is ϵ -globally-robust.*

Proof. Assume x_1 and x_2 satisfy $\|x_1 - x_2\| \leq \epsilon$. Let $\bar{F}^{\epsilon/2}(x_1) = c_1$ and $\bar{F}^{\epsilon/2}(x_2) = c_2$.

If $c_1 = \perp$ or $c_2 = \perp$, global robustness is trivially satisfied.

Consider the case where $c_1 \neq \perp$, $c_2 \neq \perp$. Let x' be the midpoint between x_1 and x_2 , i.e., $x' = (x_1 + x_2)/2$. Thus

$$\|x_1 - x'\| = \left\| \frac{x_1 - x_2}{2} \right\| = \frac{\|x_1 - x_2\|}{2} \leq \frac{\epsilon}{2}.$$

By Theorem 1, this implies $F(x') = c_1$. By the same reasoning, $\|x_2 - x'\| \leq \epsilon/2$, implying that $F(x') = c_2$. Thus, $c_1 = c_2$, so global robustness holds. \square

A.4. Proof of Theorem 3

Theorem 3. *Let f be a binary classifier that predicts $1 \iff f(x) > 0$. Let $K_L(x, \epsilon)$ be the local Lipschitz constant of f at point x with radius ϵ .*

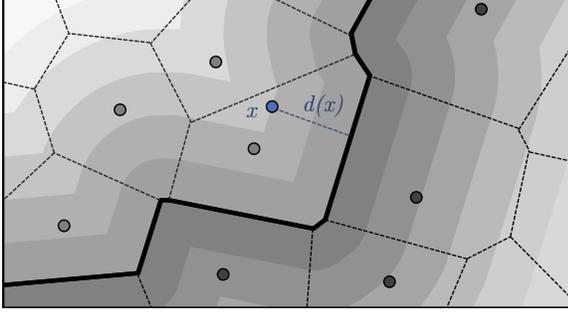


Figure A.1: Illustration of a function, g , constructed to satisfy Theorem 3. The points in S are shown in light and dark gray, with different shades indicating different labels. The Voronoi tessellation is outlined in black, and the faces belonging to the decision boundary are highlighted in bold. The level curves of g are shown in various shades of gray and correspond to points, x , at some fixed distance, $d(x)$, from the decision boundary.

Suppose that for some finite set of points, S , $\forall x \in S$, $|f(x)| > \epsilon K_L(x, \epsilon)$, i.e., all points in S can be verified via the local Lipschitz constant.

Then there exists a classifier, g , with global Lipschitz constant K_G , such that $\forall x \in S$, (1) g makes the same predictions as f on S , and (2) $|g(x)| > \epsilon K_G$, i.e., all points in S can be verified via the global Lipschitz constant.

Proof. Let T be the Voronoi tessellation generated by the points in S . Each Voronoi cell, $C_j \in T$, corresponds to the set of points that are closer to $p_j \in S$ than to any other point in S ; and the face, $F_{ij} \in T$, which separates cells C_i and C_j , corresponds to the set of points that are equidistant from p_i and p_j .

Let $B = \{F_{ij} : \text{sign}(f(p_i)) \neq \text{sign}(f(p_j))\}$, i.e., the set of faces in the Voronoi tessellation that separate points that are classified differently by f (note that B corresponds to the boundary of the 1-nearest-neighbor classifier for the points in S).

Consider a point, x . Let $p_x \in S$ be the closest point in S to x , i.e., the point corresponding to the Voronoi cell containing x . Let $d(x) = \|\text{proj}(x \rightarrow B) - x\|$; that is, $d(x)$ is the minimum distance from x to any point in any of the faces in B . Then define

$$g(x) = \text{sign}(f(p_x)) \frac{d(x)}{\epsilon}$$

First, observe that $g(x) > 0 \iff f(x) > 0$ follows from the fact that $d(x)$ and ϵ are non-negative, thus the sign of $g(x)$ is derived from the sign of $f(x)$.

Next, we show that the global Lipschitz constant of g , K_G ,

is at most $1/\epsilon$, that is, $\forall x_1, x_2$,

$$\frac{|g(x_1) - g(x_2)|}{\|x_1 - x_2\|} \leq \frac{1}{\epsilon}$$

Consider two points, x_1 and x_2 , and let p_1 and p_2 be the points in S corresponding to the respective Voronoi cells of x_1 and x_2 .

First, consider the case where $\text{sign}(f(p_1)) \neq \text{sign}(f(p_2))$, i.e., x_1 and x_2 are on opposite sides of the boundary, B . In this case $|g(x_1) - g(x_2)| = (d(x_1) + d(x_2))/\epsilon$, and thus it suffices to show that $d(x_1) + d(x_2) \leq \|x_1 - x_2\|$.

Assume for the sake of contradiction that $d(x_1) + d(x_2) > \|x_1 - x_2\|$. Note that because x_1 and x_2 belong in Voronoi cells with different classifications from f , the line segment connecting x_1 and x_2 must cross the boundary, B , at some point c . Therefore, $\|x_1 - c\| + \|x_2 - c\| = \|x_1 - x_2\| < d(x_1) + d(x_2)$; without loss of generality, this implies that $\|x_1 - c\| < d(x_1)$. But since $c \in B$, this contradicts that $d(x_1)$ is the minimum distance from x_1 to B . ζ

Next, consider the case where $\text{sign}(f(p_1)) = \text{sign}(f(p_2))$. In this case $|g(x_1) - g(x_2)| = |d(x_1) - d(x_2)|/\epsilon$, and thus, without loss of generality, it suffices to show that $d(x_1) - d(x_2) \leq \|x_1 - x_2\|$.

Assume for the sake of contradiction that $d(x_1) - d(x_2) > \|x_1 - x_2\|$. Thus $d(x_1) > \|x_1 - x_2\| + d(x_2)$. However, this suggests that we can take the path from x_1 to x_2 to B with a smaller total distance than $d(x_1)$, contradicting that $d(x_1)$ is the minimum distance from x_1 to B . ζ

We now show that $\forall p \in S$, $|g(p)| \geq 1$, i.e., $d(p) \geq \epsilon$. In other words, we must show that the distance of any point, $p \in S$ to the boundary, B , is at least ϵ . Consider a point, x , on some face, $F_{ij} \in B$. This point is equidistant from p_i and $p_j \in S$, on which f makes different predictions; and every other point in S is at least as far from x as p_i and p_j . I.e., $\|p_i - x\| = \|p_j - x\| \leq \|p - x\|$, $\forall p \in S$. By the triangle inequality, $2\|p_i - x\| \geq \|p_i - p_j\|$, and by Lemma 1, $\|p_i - p_j\| \geq 2\epsilon$. Thus $\|p - x\| \geq \epsilon$, $\forall p \in S$; therefore every point on the boundary is at least ϵ from $p \in S$.

Putting everything together, we have that $\forall p \in S$, $|g(p)| \geq 1 \geq \epsilon K_G$. \square

Note that while Theorem 3 is stated for binary classifiers, the result holds for categorical classifiers as well. We can modify the construction of g from the above proof in a straightforward way to accommodate categorical classifiers. In the case where there are m different classes, the output of g has m dimensions, each corresponding to a different class. Then, for x in a Voronoi cell corresponding to $p_x \in S$ with label, j , we define $g_j(x) ::= d(x)/\epsilon$ and $g_i(x) ::= 0$ $\forall i \neq j$. We can see that, for all pairs of classes, i and j , the Lipschitz constant of $g_i - g_j$ in this construction is the same

as the Lipschitz constant of g in the above proof, since only one dimension of the output of g changes at once. Thus, we can use the global bound suggested in Appendix A.2 to certify the points in S .

A.5. Proof of Lemma 1

Lemma 1. *Suppose that for some classifier, F , and some set of points, S , $\forall x \in S$, F is ϵ -locally-robust at x . Then $\forall x_1, x_2 \in S$ such that $F(x_1) \neq F(x_2)$, $\|x_1 - x_2\| > 2\epsilon$.*

Proof. Suppose that for some classifier, F , and some set of points, S , $\forall x \in S$, F is ϵ -locally-robust at x . Assume for the sake of contradiction that $\exists x_1, x_2 \in S$ such that $F(x_1) \neq F(x_2)$ but $\|x_1 - x_2\| \leq 2\epsilon$. Consider the midpoint between x_1 and x_2 , $x' = (x_1 + x_2)/2$. Note that

$$\|x' - x_1\| = \frac{\|x_1 - x_2\|}{2} \leq \epsilon$$

Therefore, since F is ϵ -locally-robust at x_1 , $F(x') = F(x_1)$. By the same argument, $F(x') = F(x_2)$. But this contradicts that $F(x_1) \neq F(x_2)$. \square

B. Hyperparameters

In this appendix, we describe hyperparameters used in the training of GloRo Nets to produce the results in Section 5. The full table of hyperparameters is shown in Table 1. We explain each column as follows and discuss how a particular value is selected for each hyperparameter.

Architectures. To denote architectures, we use $c(C, K, S)$ to denote a convolutional layer with C output channels, a kernel of size $K \times K$, and strides of width S . We use SAME padding unless noted otherwise. We use $d(D)$ to denote a dense layer with D output dimensions. We use rectified-linear activations after each layer except the top of the network, and do not include an explicit Softmax activation. Using this notation, the architectures referenced in Section 5 are as shown in the following list.

- 2C2F: $c(16,4,2).c(32,4,2).d(100).d(10)$
- 4C3F: $c(32,3,1).c(32,4,2).c(64,3,1).c(64,4,2).d(512).d(512).d(10)$
- 6C2F: $c(32,3,1).c(32,3,1).(32,4,2).(64,3,1).c(64,3,1).c(64,4,2).d(512).d(10)$
- 7C1F: $c(64,3,1).c(64,3,1).c(64,4,2).c(128,3,1).c(128,4,2).c(256,3,1).(256,4,2).d(200)$

We arrived at these architectures in the following way. 2C2F, 4C3F and 6C2F are used in the prior work (Lee et al., 2020; Croce et al., 2019; Wong & Kolter, 2018) to evaluate the

verifiable robustness, and we used them to facilitate a direct comparison on training cost and verified accuracy. For Tiny-ImageNet, we additionally explored the architecture described in (Lee et al., 2020) for use with that dataset, but found that removing one dense and one convolutional layer (denoted by 7C1F in the list above) produced the same (or better) verified accuracy, but lowered the total training cost.

Data preprocessing. For all datasets, we scaled the features to the range $[0,1]$. On some datasets, we used the following data augmentation, which is denoted by default in Table 1.

```
rotation_range=20
width_shift_range=0.2
height_shift_range=0.2
horizontal_flip=True
shear_range=0.1
zoom_range=0.1
```

Our use of augmentation follows the convention established in prior work (Wong et al., 2018; Lee et al., 2020): we only use it on CIFAR and tiny-imatenet, but not on MNIST.

ϵ scheduling. Prior work has also established a convention of gradually scaling ϵ up to a value that is potentially larger than the one used to measure verified accuracy leads to better results. We made use of the following schemes for accomplishing this.

- No scheduling: we use ‘single’ to denote we ϵ_{train} for all epochs.
- Linear scheduling: we use a string ‘ x, y, e ’ to denote the strategy that at training epoch t , we use $\epsilon_t = x + (y - x) * (t/e)$ if $t \leq e$. When $t > e$, we use the provided ϵ_{train} to keep training the model.
- Logarithmic scheduling: we use ‘log’ to denote that we increase the epsilon with a logarithmic rate from 0 to ϵ_{train} .

We found that scheduling ϵ is often unnecessary when instead scheduling the TRADES parameter λ (discussed later in this section), which appears to be more effective for that loss. To select a scheme for scheduling ϵ , we compared the results of the three options listed above, and selected the one that achieved the highest verified accuracy. If there was no significant difference in this metric, then we instead selected the schedule with the least complexity, assuming the following order: single, (x, y, e), log. When applying (x, y, e) and log, we began the schedule on the first epoch, and ended it on $(\# \text{ epochs})/2$.

Initialization & optimization. In Table 1, `default` refers to the Glorot uniform initialization, given by `tf.keras.initializers.GlorotUniform()`. The string `'ortho'` refers to an orthogonal initialization given by `tf.keras.initializers.Orthogonal()`. To select an initialization, we compared the verified accuracy achieved by either, and selected the one with the highest metric. In the case of a tie, we defaulted to the Glorot uniform initialization. We used the `adam` optimizer to perform gradient descent in all experiments, with the initial learning rate specified in Table 1, and default values for the other hyperparameters ($\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 1e - 07$, `amsgrad` disabled).

Learning rate scheduling. We write `'decay_to_lb'` to denote a schedule that continuously decays the learning rate to `lb` at a negative-exponential rate, starting the decay at $(\#epochs)/2$. To select `lb`, we searched over values $lb \in \{1 \times 10^{-7}, 5 \times 10^{-7}, 1 \times 10^{-6}, 5 \times 10^{-6}\}$, selecting the value that led to the best VRA. We note that for all datasets except Tiny-Imagenet, we used the default initial rate of 1×10^{-3} . On Tiny-Imagenet, we observed that after several epochs at this rate, as well as at 5×10^{-4} , the loss failed to decrease, so again halved it to arrive at 2.5×10^{-4} .

Batch size & epochs. For all experiments, we used minibatches of 256 instances. Because our method does not impose a significant memory overhead, we found that this batch size made effective use of available hardware resources, increasing training time without impacting verified accuracy, when compared to minibatch sizes 128 and 512. Because the learning rate, ϵ , and λ schedules are all based on the total number of epochs, and can have a significant effect on the verified accuracy, we did not monitor convergence to determine when to stop training. Instead, we trained for epochs in the range [100, 1000] in increments of 100, and when verified accuracy failed to increase with more epochs, attempted training with fewer epochs (in increments of 50), stopping the search when the verified accuracy began to decrease again.

Warm-up. Lee et al. (2020) noted improved performance when models were pre-trained for a small number of epochs before optimizing the robust loss. We found that this helped in some cases with GloRo networks as well, in particular on CIFAR-10, where we used the same number of warm-up epochs as prior work.

λ scheduling. When using the TRADES loss described in Section 4, we found that scheduling λ often yielded superior results. We use `'x, y, e'` to denote that at epoch t , we set $\lambda_t = x + (y - x) * (t/e)$ if $t < e$ else $\lambda_t = y$. We write `'x'` to denote we use $\lambda = x$ all the time. To select the final λ , we trained on values in the range [1, 10] in increments of 1,

and on finding the whole number that yielded the best result, attempted to further refine it by increasing in increments 0.1.

Power iteration. As discussed in Section 4, we use power iteration to compute the spectral norm for each layer to find the layer-wise Lipschitz constants. In Table 1, `power_iter` denotes the number of iterations we run for each update during training. We tried values in the set {1, 2, 5, 10}, breaking ties to favor fewer iterations for less training cost. After each epoch, we ran the power iteration until convergence (with tolerance 1×10^{-5}), and all of the verified accuracy results reported in Section 5 are calculated using a global bound based on running power iteration to convergence as well. Since the random variables used in the power iterations are initialized as `tf.Variables`, they are stored in `.h5` files together with the architecture and network parameters. Therefore, one can directly use the converged random variables from the training phase during the test phase.

Search strategy. Because of the number of hyperparameters involved in our evaluation, and limited hardware resources, we did not perform a global grid search over all combinations of hyperparameters discussed here. We plan to do so in future work, as it is possible that results could improve as we may have missed better settings than those explored to produce the numbers reported in our evaluation. Instead, we adopted a greedy strategy, prioritizing the choices that we believed would have the greatest impact on verified accuracy and training cost. In general, we explored parameter choices in the following order: ϵ schedule, λ schedule, # epochs, LR decay, warm-up, initialization, # power iterations, minibatch size.

C. Issues Reproducing BCP Tiny-Imagenet

In order to compare the cost of training GloRo networks to the prior work of Lee et al. (2020) and Wong et al. (2018), we obtained code from the authors' public repositories as noted in the corresponding papers, and attempted to reproduce their reported training results on our hardware. We succeeded in all cases, except for the Tiny-Imagenet results reported by Lee et al. (2020). We note that at the time of writing (commit 85defee from <https://github.com/sungyoon-lee/bcp>), the repository for that work has pre-trained models for MNIST and CIFAR, but none for Tiny-Imagenet. Our inability to reproduce the reported results, or to measure verified accuracy on a pre-trained model, is the reason that we do not include the verified accuracy reported for this method in Table 3a. Below we describe the steps that we took in our attempt to reproduce them.

Data preparation. We obtained and formatted Tiny-Imagenet for use with the BCP code as follows.

- Download the dataset from <http://cs231n.stanford.edu/tiny-imagenet-200.zip>.
- Prepare the validation dataset by creating a folder for each class.
- Modify the paths to the root folders for the training and validation data at [lines 64 and 65](#) in `data_load.py`
- Applied a fix to the labels, based on communication with the authors. We were advised to comment out [line 100](#) in `data_load.py` to make their `data_load.py` work.

Training. While the public repository has scripts for training MNIST and CIFAR models using BCP, there are none for Tiny-Imagenet. However, after reviewing the code in `train_cifar10.py`, we concluded that it is general enough to work for Tiny-Imagenet when given appropriate parameters (namely, `--data tinyimagenet`). We did not pass the argument for `--model`, as the script defaults to the only model built for Tiny-Imagenet ([line 458](#) in `utils.py`) when given Tiny Imagenet as the dataset. We invoked the following command.

```
python train_cifar10.py --warmup 0 \
--batch_size 5 --rampup 50 --opt sgd \
--momentum 0.1 --weight_decay 0.0002 \
--epochs 100 --data tinyimagenet \
--cuda_ids 2 --test_opt_iter 1 \
--test_sniter 1 --niter 1
```

We note that on the machine we used for testing, `--cuda_ids 2` contains a Titan X accelerator, which is the same hardware that the authors reported having used to obtain their results. Additionally, the argument `--test sniter 1` was chosen while attempting to address Issue 2 (below), as we believed that it would reduce memory pressure.

Issue 1: Indexing Error. This command quickly leads to a CUDA indexing error reported by Pytorch. We managed to locate the error at [line 173](#) in `BCP.py`. We found that the function `utils.one_hot`, defined at [line 500](#) in `utils.py` takes one positional argument `batch` and another keyword argument `depth=10`. From our understanding, `depth` refers to the size of the one-hot vector used for training, which in the case of Tiny Imagenet should be 200. However, when `utils.one_hot` is called at [line 173](#) in `BCP.py`, the keyword argument `depth` is not provided, which means that the default value of 10 will be used. As a result, when used with Tiny-Imagenet (200 classes), this led to the observed indexing error whenever samples with label > 9 are encountered on executing [line 173](#).

Fixing Issue 1. We modify [line 173](#) in `BCP.py` by passing the appropriate number of classes as the keyword argument `depth` to `utils.one_hot`.

Issue 2: Resource Exhaustion. When employing the above fix to Issue 1, we encountered a CUDA out-of-memory error at the second epoch of BCP training. We tried again on an accelerator with more memory (Titan RTX, 24 GB), but this also failed to resolve the error after several attempts. After some experimentation, we found that the maximum number of classes that would fit in memory with BCP training is 100, which would not allow us to complete the training with Tiny-Imagenet.

Current status. We hope to resolve this issue with the authors, and ultimately include reproducible results for Tiny-Imagenet with BCP in future versions of this paper. At the time of writing, we have engaged in several rounds of communication with the authors, but after several weeks, have not received a response that would confirm that our fix to `utils.one_hot` is appropriate, or an explanation for the resource exhaustion issue. We will continue our attempts to work towards a resolution with the authors until this work has completed peer review.

D. Measuring Memory Usage

In our experiments, we used Tensorflow to train and evaluate standard and GloRo networks, and Pytorch to train and evaluate KW and BCP (since [Wong & Kolter \(2018\)](#) and [Lee et al. \(2020\)](#) implement their respective methods in Pytorch). To measure memory usage, we invoked `tf.contrib.memory_stats.MaxBytesInUse()` at the end of each epoch for standard and GloRo networks, and took the peak active use from `torch.cuda.memory_summary()` at the end of each epoch for KW and BCP.

We note that some differences may arise as a result of differences in memory efficiency between Tensorflow and Pytorch. In particular, Pytorch enables more control over memory management than does Tensorflow. In order to mitigate this difference as much as possible, we did not disable gradient tracking when evaluating certification times and memory usage in Pytorch. While gradient tracking is unnecessary for certification (it is only required for training), Tensorflow does not allow this optimization, so by forgoing it the performance results recorded in [Figure 3b](#) in [Section 5](#) are more comparable across frameworks.

E. Optimizing for Lipschitz Lower Bounds

[Figure 3c](#) in [Section 5](#) gives empirical lower bounds on the global and average local Lipschitz constants on the models trained in our evaluation. We use optimization to obtain

these lower bounds; further details are provided below.

Global Lower Bounds. We use the *margin Lipschitz constant*, K_{ij}^* (Definition 4 in Appendix A.2), which takes a different value for each pair of classes, i and j . To obtain the lower bound we optimize

$$\max_{x_1, x_2} \max_i \left\{ \frac{|f_{j_1}(x_1) - f_i(x_1) - (f_{j_1}(x_2) - f_i(x_2))|}{\|x_1 - x_2\|} \right\}$$

where $j_1 = F(x_1)$. Optimization is performed using Keras' default `adam` optimizer with 7,500 gradient steps. Both x_1 and x_2 are initialized to random points in the test set; we perform this optimization over 100 such initial pairs, and report the maximum value obtained over all initializations.

Local Lower Bounds. We use a variant of the *margin Lipschitz constant* (Definition 4 in Appendix A.2) analogous to the local Lipschitz constant at a point, x_0 , with radius ϵ . To obtain this lower bound we optimize

$$\max_{x_1, x_2} \max_i \left\{ \frac{|f_j(x_1) - f_i(x_1) - (f_j(x_2) - f_i(x_2))|}{\|x_1 - x_2\|} \right\}$$

$$\text{subject to } \|x_1 - x_0\| \leq \epsilon, \|x_2 - x_0\| \leq \epsilon$$

where $j = F(x_0)$. Optimization is performed using Keras' default `adam` optimizer with 5,000 gradient steps. After each gradient step, x_1 and x_2 are projected onto the ϵ -ball centered at x_0 . Both x_1 and x_2 are initialized to random points in the test set, and x_0 is a fixed random point in the test set. We perform this optimization over 100 random choices of x_0 , and report the mean value.

Discussion. In Section 5.3, we observe that the global upper bound is fairly tight on the GloRo Net trained on MNIST, but decreasingly so on CIFAR-10 and Tiny-Imagenet. While this suggests that there is room for improvement in terms of the bounds obtained by GloRo Nets, we make note of two subtleties that may impact these findings.

First, the tightness decreases inversely with the dimensionality of the input. While it is reasonable to conclude that learning tight GloRo Nets in higher dimensions becomes increasingly difficult, it is worth noting that the optimization process described above also becomes more difficult in higher dimensions, meaning that some of the looseness may be attributable to looseness in the *lower* bound, rather than in the upper bound.

Second, the hyperparameters used may have an effect on the tightness of the Lipschitz bound. As seen in Appendix B, different hyperparameters were used on MNIST, CIFAR-10, and Tiny-Imagenet; some of these differences were selected based on impacting training time, which is of greater concern for larger datasets that naturally take longer to train.

Specifically, we note that fewer power iterations were used for CIFAR-10, and even fewer for Tiny-Imagenet. While this is good for expediency, and still produces state-of-the-art VRA, we note that tighter bounds may be learned by putting more computation time into training, in the form of more power iterations (for example). More generally, this speculation suggests that slightly different training strategies, hyperparameters, etc., from the ones used in this work may be sufficient to improve the bounds and the VRA achieved by GloRo Nets. We conclude that future work should further explore this possibility.

Globally-Robust Neural Networks

<i>architecture</i>	<i>dataset</i>	<i>data augmentation</i>	<i>warm-up</i>	<i>batch size</i>	<i># epochs</i>	ϵ_{train}	ϵ_{test}
2C2F GloRo	MNIST	None	0	256	500	0.45	0.3
	<i>initialization</i>	<i>init_lr</i>	<i>lr_decay</i>	<i>loss</i>	ϵ <i>schedule</i>	<i>power_iter</i>	
	default	1e-3	decay_to_1e-6	CE	single	10	
<i>architecture</i>	<i>dataset</i>	<i>data augmentation</i>	<i>warm-up</i>	<i>batch size</i>	<i># epochs</i>	ϵ_{train}	ϵ_{test}
2C2F GloRo-T	MNIST	None	0	256	500	0.45	0.3
	<i>initialization</i>	<i>init_lr</i>	<i>lr_decay</i>	<i>loss</i>	ϵ <i>schedule</i>	<i>power_iter</i>	
	default	1e-3	decay_to_1e-6	0, 2, 500	single	10	
<i>architecture</i>	<i>dataset</i>	<i>data augmentation</i>	<i>warm-up</i>	<i>batch size</i>	<i># epochs</i>	ϵ_{train}	ϵ_{test}
4C3F GloRo	MNIST	None	0	256	300	1.75	1.58
	<i>initialization</i>	<i>init_lr</i>	<i>lr_decay</i>	<i>loss</i>	ϵ <i>schedule</i>	<i>power_iter</i>	
	default	1e-3	decay_to_5e-6	CE	single	10	
<i>architecture</i>	<i>dataset</i>	<i>data augmentation</i>	<i>warm-up</i>	<i>batch size</i>	<i># epochs</i>	ϵ_{train}	ϵ_{test}
4C3F GloRo-T	MNIST	None	0	256	300	1.75	1.58
	<i>initialization</i>	<i>init_lr</i>	<i>lr_decay</i>	<i>loss</i>	ϵ <i>schedule</i>	<i>power_iter</i>	
	default	1e-3	decay_to_5e-6	0, 3, 300	single	10	
<i>architecture</i>	<i>dataset</i>	<i>data augmentation</i>	<i>warm-up</i>	<i>batch size</i>	<i># epochs</i>	ϵ_{train}	ϵ_{test}
6C2F GloRo	CIFAR-10	default	20	256	800	0.1551	0.141
	<i>initialization</i>	<i>init_lr</i>	<i>lr_decay</i>	<i>loss</i>	ϵ <i>schedule</i>	<i>power_iter</i>	
	orthogonal	1e-3	decay_to_1e-6	CE	log	5	
<i>architecture</i>	<i>dataset</i>	<i>data augmentation</i>	<i>warm-up</i>	<i>batch size</i>	<i># epochs</i>	ϵ_{train}	ϵ_{test}
6C2F GloRo-T	CIFAR-10	default	20	256	800	0.1551	0.141
	<i>initialization</i>	<i>init_lr</i>	<i>lr_decay</i>	<i>loss</i>	ϵ <i>schedule</i>	<i>power_iter</i>	
	default	1e-3	decay_to_1e-6	1.2	log	5	
<i>architecture</i>	<i>dataset</i>	<i>data augmentation</i>	<i>warm-up</i>	<i>batch size</i>	<i># epochs</i>	ϵ_{train}	ϵ_{test}
7C1F GloRo	Tiny-Imagenet	default	0	256	250	0.16	0.141
	<i>optimizer</i>	<i>init_lr</i>	<i>lr_decay</i>	<i>loss</i>	ϵ <i>schedule</i>	<i>power_iter</i>	
	default	2.5e-4	decay_to_5e-7	CE	single	5	
<i>architecture</i>	<i>dataset</i>	<i>data augmentation</i>	<i>warm-up</i>	<i>batch size</i>	<i># epochs</i>	ϵ_{train}	ϵ_{test}
7C1F GloRo-T	Tiny-Imagenet	default	0	256	500	0.16	0.141
	<i>initialization</i>	<i>init_lr</i>	<i>lr_decay</i>	<i>loss</i>	ϵ <i>schedule</i>	<i>power_iter</i>	
	default	2.5e-4	decay_to_5e-7	1, 10, 500	single	1	

Table 1: Hyperparameters used for training GloRo Nets and GloRo-T Nets