



**HAL**  
open science

# Riemannian data-dependent randomized smoothing for neural networks certification

Pol Labarbarie, Hatem Hajri, Marc Arnaudon

► **To cite this version:**

Pol Labarbarie, Hatem Hajri, Marc Arnaudon. Riemannian data-dependent randomized smoothing for neural networks certification. New Frontiers in Adversarial Machine Learning in the Thirty-ninth International Conference on Machine Learning (ICML), Jul 2022, Baltimore (MA), United States. hal-04264005

**HAL Id: hal-04264005**

**<https://hal.science/hal-04264005>**

Submitted on 29 Oct 2023

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Copyright

---

# Riemannian data-dependent randomized smoothing for neural networks certification

---

Pol Labarbarie<sup>1</sup> Hatem Hajri<sup>1</sup> Marc Arnaudon<sup>2</sup>

## Abstract

Certification of neural networks is an important and challenging problem that has been attracting the attention of the machine learning community since few years. In this paper, we focus on randomized smoothing (RS) which is considered as the state-of-the-art method to obtain certifiably robust neural networks. In particular, a new data-dependent RS technique called ANCER introduced recently can be used to certify ellipses with orthogonal axis near each input data of the neural network. In this work, we remark that ANCER is not invariant under rotation of input data and propose a new rotationally-invariant formulation of it which can certify ellipses without constraints on their axis. Our approach called Riemannian Data Dependant Randomized Smoothing (RDDRS) relies on information geometry techniques on the manifold of covariance matrices and can certify bigger regions than ANCER based on our experiments on the MNIST dataset. Code and models are available at <https://github.com/PLabarbarie/RDDRS>.

## 1. Introduction

Making neural networks robust against adversarial examples (Szegedy et al., 2013; Biggio et al., 2013) or proving that they are has been the focus of many works. Robustification methods can be split into two main categories. First, empirical defenses empirically improve the resilience of neural networks but has no theoretical guarantee. The most and only successful known empirical defense technique is adversarial training (Madry et al., 2017). Second, certified defense techniques try to provide a guarantee that an adversary does not exist in a certain neighborhood around a given

<sup>1</sup>Institut de Recherche Technologique SystemX, Palaiseau, France <sup>2</sup>Univ. Bordeaux, CNRS, Bordeaux INP, IMB, UMR 5251, Talence, France. Correspondence to: Pol Labarbarie <pol.labarbarie@irt-systemx.fr>.

input (Huang et al., 2017; Raghunathan et al., 2018; Bunel et al., 2018; Lomuscio & Maganti, 2017; Tjeng et al., 2017; Weng et al., 2018; Lecuyer et al., 2019; Cohen et al., 2019). Despite the progress on these methods, the certified regions are still meaningless compared with the human perception.

In this paper we focus on the Randomized smoothing (RS) technique (Lecuyer et al., 2019; Cohen et al., 2019) which, despite its simplicity, has proved the state-of-the-art to certify neural networks on small neighborhoods of inputs. Given a base classifier  $F$ , RS constructs a new classifier  $G = F * \mathcal{N}(0, \sigma^2 I_d)$ . The smoothed classifier is certified robust:  $G(x) = G(x + \delta)$  for all  $\delta$  such as  $\|\delta\|_2 < R(x, \sigma)$  with

$$R(x, \sigma) = \frac{\sigma^2}{2} (\Phi^{-1}(p_A) - \Phi^{-1}(p_B)) \quad (1)$$

where  $p_A = \mathbb{E}_Z[F_{c_A}(x + Z)] = G(x)_{c_A}$  is the probability of the top class  $c_A$ ,  $p_B = \max_{c \neq c_A} \mathbb{E}_Z[F_c(x + Z)] = G(x)_{c_B}$  is the probability of the runner-up class  $c_B$ ,  $Z \sim \mathcal{N}(0, \sigma^2 I_d)$  and  $\Phi^{-1}$  is the inverse cumulative distribution function of a standard Gaussian distribution. Several extensions and improvements of RS have been proposed subsequently. For instance, by adversarially training the base classifier (Salman et al., 2019), regularization (Zhai et al., 2020), general adjustments for training routines (Zhai et al., 2020) and by extending RS to more general distributions beyond the Gaussian case (Yang et al., 2020).

One limitation of the original RS work (Cohen et al., 2019) is that it only certifies isotropic regions. Properly extending this approach to anisotropic domains is a first important challenge. Enlarging at most the certified regions is a second ongoing challenge for all the proposed approaches. To generalise RS beyond fixed variance noise, (Alfarra et al., 2020) proposed data-dependent randomised smoothing (DDRS) which additionally maximises the radius (1) over the parameter  $\sigma$ . This idea has been improved in (Eiras et al., 2021) which proposes ANCER; an anisotropic diagonal certification method that performs sample-wise (*i.e.* per sample in the test set) region volume maximization. They generalized (1) for ellipsoid regions *i.e.* they proved that for  $\Sigma$  a non degenerate covariance matrix,  $G(x) = G(x + \delta)$  for all  $\delta$

such as  $\|\delta\|_{\Sigma,2} < r(x, \Sigma)$  with

$$r(x, \Sigma) = \Phi^{-1}(p_A) - \Phi^{-1}(p_B) \quad (2)$$

and where  $G = F * \mathcal{N}(0, \Sigma)$ .

**Motivation and contribution.** The main motivation of our work is that ANCER is not resilient to rotations of input data. Indeed, in a configuration where the input data are concentrated around one canonical axis, ANCER may give good results. However, rotating all the points will dramatically impact the method since it always seeks for axis parallel to the canonical ones which will not be optimal after rotation. We propose an approach based on the use of information geometry which adequately fixes this issue by design as depicted by Figure 1. To summarize our approach, we consider a general Gaussian noise  $\mathcal{N}(0, C)$  with covariance  $C$  and for each input  $x$  compute the corresponding certification radius  $R(x, C)$ . Next, we maximize  $R(x, C)$  over  $C$  on the manifold of covariance matrices through Riemannian optimisation and finally smooth the original classifier with the noise  $\mathcal{N}(0, C)$ . This methodology generalises ANCER which assumes that  $C$  is a diagonal matrix and optimizes its diagonal in the Euclidean space. Through experiments on MNIST (LeCun et al., 1998), we show that our algorithm outperforms the previous works (Alfarra et al., 2020; Eiras et al., 2021) and obtain new state-of-the-art certified accuracy for the DDRS technique.

**Organisation of the paper.** In Section 2, we review existing DDRS techniques. In Section 3, we present our approach named Riemannian DDRS. In Section 4, we experiment our method and show its advantage with respect to the state-of-the-art. Finally, in Section 5, we discuss limitations of our method and RS techniques in general and outline some challenges related to them.

**Notations.** We denote by  $F : \mathbb{R}^d \rightarrow \mathcal{P}(\mathcal{Y})$  a base classifier typically a neural network, where  $\mathcal{P}(\mathcal{Y})$  is a probability simplex over  $K$  classes.  $\Phi^{-1}$  is the inverse cumulative distribution function of a standard Gaussian distribution. The  $\ell_p$ -ball is defined with respect to the  $\|\cdot\|_p$  norm ( $p \geq 1$ ) and the  $\ell_p^A$ -ellipsoid is defined with respect to  $\|\cdot\|_{A,p}$ . For  $p \in \{1, 2\}$ ,  $\|\cdot\|_{A,p}$  is the composite norm defined with respect to a positive definite matrix  $A$  and a vector  $u$  as  $\|A^{-1/p}u\|_p$ .

## 2. Data-dependent randomized smoothing

DDRS (Alfarra et al., 2020) aims to improve the choice of the hyperparameter  $\sigma$  in order to obtain bigger certified balls. This parameter is supposed to be fixed by the original RS method. Concretely, DDRS solves the optimisation problem

$$\sigma_x^* = \arg \max_{\sigma} R(x, \sigma) \quad (3)$$

using gradient ascent and defines the smoothed classifier  $G$  by means of the sampling noise  $\mathcal{N}(0, \sigma_x^{*2}I)$  for each input  $x$ . This approach showed better certified isotropic regions per points. Next, (Eiras et al., 2021) generalized randomized smoothing to certify anisotropic regions, by pairing it with a generalization of the data-dependent method to maximize the certified region at each input point. Following a Lipschitz argument of (Salman et al., 2019; Jordan & Dimakis, 2020), they developed a methodology to compare anisotropic regions with isotropic ones. They proposed the following optimisation problem

$$\arg \max_{\Theta_x} r(x, \Theta_x) \sqrt{\prod_i \theta_i^x} \quad (4)$$

under the constraint that  $\min_i \theta_i^x r(x, \Theta_x) \geq r_{iso}^*$  where  $r(x, \Theta_x) = \Phi^{-1}(p_A) - \Phi^{-1}(p_B)$ ,  $r_{iso}^*$  is the maximum isotropic  $l_2$  radius obtained with (3) and  $\Theta_x = \text{diag}(\{\theta_i\}_{i=1}^d)$ . The constraint ensures that the isotropic region obtained by (3) is enclosed in the anisotropic one.

## 3. Riemannian data-dependent randomized smoothing

As outlined previously, we will consider general smoothing noises  $\mathcal{N}(0, C)$  where  $C$  is a covariance matrix, i.e a symmetric positive definite one. For this, we will work on the space  $\mathcal{S}_d^{++}$  of  $d \times d$  covariance matrices which will be equipped with the Riemannian metric known as the affine-invariant metric or Rao–Fisher metric in information geometry (Atkinson & Mitchell, 1981; Skovgaard, 1984):

$$g_Y(A, B) = \text{tr}(Y^{-1}AY^{-1}B)$$

where  $Y \in \mathcal{S}_d^{++}$  and  $A, B \in T_Y\mathcal{S}_d^{++}$ , with  $T_Y\mathcal{S}_d^{++}$  is the tangent space to  $\mathcal{S}_d^{++}$  at  $Y$ , identified with the space of symmetric  $d \times d$  matrices. The Riemannian distance associated with this metric is

$$d^2(Y, Z) = \text{tr}(\log(Y^{-\frac{1}{2}}ZY^{-\frac{1}{2}}))^2$$

for all  $Y, Z \in \mathcal{S}_d^{++}$ . Next, we consider the following optimisation problem on  $\mathcal{S}_d^{++}$ :

$$\arg \max_{C^x} R(C^x) \sqrt{\prod_i \lambda_i^x}$$

such that  $\min_i \lambda_i^x R(C^x) \geq r_{iso}^*$  and  $\lambda_i^x \geq \sigma_x^*$

and where

$$R(C^x) = (\Phi^{-1}(G(C^x)_{c_A}) - \Phi^{-1}(G(C^x)_{c_B})),$$

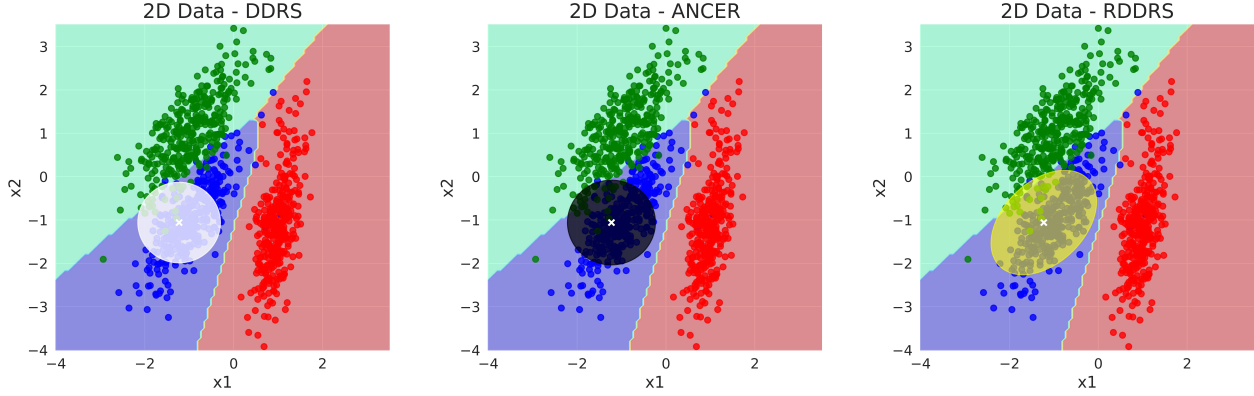


Figure 1. Representation of a 2D classification problem in 3 classes (red, blue, and green). For the same point (white cross) from the test set, we run DDRS, ANCER and RDDRS. We plot the boundaries of the isotropic sampling region obtained with DDRS (white), anisotropic diagonal sampling region obtained with ANCER (black), and general anisotropic sampling region obtained with RDDRS (yellow).

$G(C) = \mathbb{E}[F(x + CZ)]$  with  $Z \sim \mathcal{N}(0, I_d)$  and the  $\lambda_i^x$  are the eigenvalues of the matrix  $C^x$ . By using a penalisation method, we instead consider the maximisation problem:

$$\arg \max_{C^x} (R(C^x)P(C^x) + \kappa K(C^x)R(C^x)) \quad (5)$$

such that for all  $i \in \{1, \dots, d\}$ ,  $\lambda_i \geq \sigma_x^*$  and where:

$$P(C^x) = \sqrt{\prod_i \lambda_i^x}, \quad K(C^x) = \min_i \lambda_i^x$$

Define  $H(C^x) = R(C^x)P(C^x) + \kappa K(C^x)R(C^x)$ . Then, we can use Riemannian gradient descent to locally maximise  $H$  on  $\mathcal{S}_d^{++}$  by iterating

$$C_{t+1}^x = \text{Exp}_{C_t^x}(\gamma_t \nabla H(C_t^x)) \quad (6)$$

where  $\gamma_t$  is a step size depending on  $t$  and  $\text{Exp}$  is the exponential map:

$$\text{Exp}_Y(\Delta) = Y^{\frac{1}{2}} \exp(Y^{-\frac{1}{2}} \Delta Y^{-\frac{1}{2}}) Y^{\frac{1}{2}}$$

The update (6) requires computing  $\nabla H$  which is given by:

$$\begin{aligned} \nabla H(C_t^x) &= \nabla R(C_t^x)P(C_t^x) + R(C_t^x)\nabla P(C_t^x) \\ &+ \kappa (\nabla R(C_t^x)K(C_t^x) + R(C_t^x)\nabla K(C_t^x)) \end{aligned}$$

with

$$\nabla R(C_t^x) = \left( \frac{\nabla(G(C_t^x)_{c_A})}{\Phi'(\Phi^{-1}(G(C_t^x)_{c_A}))} - \frac{\nabla(G(C_t^x)_{c_B})}{\Phi'(\Phi^{-1}(G(C_t^x)_{c_B}))} \right)$$

Let us now give the Riemannian gradient of  $G$ . If  $Z \sim \mathcal{N}(0, I_d)$ , we have  $\nabla G(C) = C\mathbb{E}[A]C$ , where:

$$A_{ij} = \frac{1}{2}(Z_j(\nabla F(y))_i + Z_i(\nabla F(y))_j)$$

and  $y := x + CZ$ . Using this expression, we can approximate  $\mathbb{E}[A]$  by Monte-Carlo and then deduce  $\nabla G(C)$ . Similarly,  $\nabla P(C) = CBC$  and  $\nabla K(C) = CDC$  where

$$B_{ij} = \frac{1}{2}((\nabla P(C))_i + (\nabla P(C))_j)$$

and

$$D_{ij} = \frac{1}{2}((\nabla K(C))_i + (\nabla K(C))_j)$$

## 4. Experiments

In this section, we empirically test the performance of our RDDRS framework for neural networks certification and show that it achieves state-of-the-art performance in terms of certified regions on the well-known MNIST dataset. The methods used for comparison are randomized smoothing (RS) as it appeared in (Cohen et al., 2019), DDRS (Alfarra et al., 2020) and ANCER (Eiras et al., 2021). For evaluation, the same procedure used in (Cohen et al., 2019; Salman et al., 2019; Alfarra et al., 2020; Eiras et al., 2021) is followed. We train a classical convolutional neural network on MNIST using the Gaussian data augmentation proposed by (Cohen et al., 2019) with  $\sigma$  varying in  $\{0.12, 0.25, 0.5\}$  and then perform the certification on the whole MNIST test set. To compare between the different tested methods, we plot the approximate certified accuracy curve similar to previous works. This curve is computed for many radius  $R$

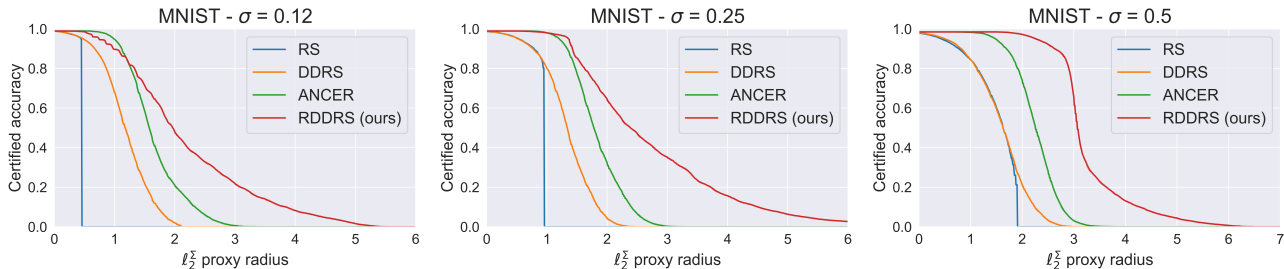


Figure 2. Certified accuracy as a function of the proxy radius  $\ell_2^2$  for several  $\sigma$  used in training for MNIST dataset. We compare our Riemannian framework RDDRS,  $\Sigma = C^x$ , against the standard randomized smoothing (RS),  $\Sigma = \sigma^2 I_d$ , the data-dependent randomized smoothing (DDRS),  $\Sigma = \sigma_x^2 I_d$ , and the anisotropic diagonal data-dependent randomized smoothing (ANCER),  $\Sigma = \Theta^x$ .

and corresponds for each  $R$  to the certified accuracy given by the fraction of the test set images which  $G$  classifies correctly and certifies robust with a radius  $R' \geq R$ . During certification by the state-of-the-art methods, we use the same  $\sigma$  as for training. To run these methods, we closely follow the recommendations of the respective papers. The DDRS optimization of (3) is done using a gradient ascent, a learning rate  $\alpha = 0.0001$ , an initial  $\sigma_0 = \sigma$  and a number of samples  $n = 100$  for Monte-Carlo estimation of  $p_A$  at each iteration. Note that following the previous works there is no need to estimate  $p_B$  even if the number of classes is greater than 2 (see (Cohen et al., 2019)). The number of gradient ascent iterations for DDRS is taken from  $K = 100$  to  $K = 1500$  with a step of 100 and the  $\sigma$  providing the best certified radius is saved. For ANCER, we resolved (4) with different values of the learning rate  $\alpha \in \{0, 04, 0, 4\}$ , a penalization parameter  $\kappa = 2$ , the same number of iterations  $K = 100$  and estimate  $p_A$  with  $n = 100$ . Our method RDDRS is run with different values of the learning rate  $\gamma_t \in \{0, 5, 1, 25\}$ ,  $\kappa = 10^{-6}$ , the same number of iterations  $K = 100$  and a number of samples  $n = 20000$  to estimate  $\mathbb{E}[A]$ . To calculate the exponential map on  $\mathcal{S}_d^{++}$ , we use the Geomstats library (Miolane et al., 2020a;b). After obtaining a first trained model for three values of  $\sigma$  as previously described, we run the four methods RS, DDRS, ANCER and RDDRS and report the final certified accuracy. The obtained results are plotted in Figure 2 and commented in what follows.

**Comments.** The graphics of the certified accuracy obtained for the four methods and the three values of  $\sigma$  demonstrate that our RDDRS method globally outperforms DDRS and ANCER for almost all proxy radius and for all the values of  $\sigma$ . The difference becomes more important in our favour when  $R$  gets bigger, *i.e.*  $R > 3$  and also when  $\sigma$  increases. Only for  $\sigma = 0.12$  and small radius  $R$  our method is slightly less better than ANCER. Moreover, we observe that none of the certification methods except ours can certify a good percentage of the test set at large radii. All of this illustrates that our method has significantly pushed forward the data-dependent smoothing techniques.

**Discussion on the parameter  $\kappa$ .** (Eiras et al., 2021) proposed to penalize (4) for empirical considerations. They set the default value of  $\kappa$  to 2. For our formulation (5), we have conducted several runs of RDDRS with different values of  $\kappa$  and have concluded that this factor does not have a significant effect on the method’s convergence. The lower bound condition  $\lambda_i^x \geq \sigma_x^*$  has however a more significant effect on that convergence.

**Runtime discussion.** Our method is slower than the other methods. This essentially comes from the computation of  $\mathbb{E}[A]$  for which we considered a greater number of samples  $n = 20000$ . Indeed, despite the fact that this value is relatively high, it allows our results to be very stable: different runs with the same  $n$ , give almost the same performances.

**Scalability challenges.** Computing the exponential map becomes costly for high dimensional matrices. We believe that, on datasets such as CIFAR10 (Krizhevsky et al.), these computations are feasible and our method is still applicable with reasonable computation resources. However, extending the approach to high-dimensional data such as IMAGENET (Deng et al., 2009) seems to be very challenging. Some geometric approximations of the exponential map could be useful for that.

## 5. Limitations and Conclusion

Randomized smoothing is currently the state-of-the-art method to certify neural networks despite its curse of dimensionality limitation (Kumar et al.). In this paper, we improved the data-dependent RS technique by looking for the best Gaussian covariance matrix to define the smoothed classifier. Our method is less scalable than other approaches but performs significantly better on MNIST which is an important dataset in deep learning. We set up several challenges in the frontier between differential geometry and deep learning to push forward this approach to high-dimensional data. We relied on the Rao-Fisher distance, but other metrics could have been used such as the log-Euclidean metric on the space of covariance matrices (Arsigny et al., 2006). It seems to us an important question to test our method in

terms of cost and also certification regions as a function of the used distance. Regarding complexity, it should be remarked that RS networks are in general slower than classical convolutional neural network. Data dependent smoothing networks are even slower. Also, they are little bit ambiguous; indeed, one network is defined and optimised for each input data. To adversarially attack the smoothed network at a given input, one has to attack the smoothed network at this point; defined with the optimal noise found in this paper. While this procedure is much more computational, it is more robust as it guarantees that adversarial examples can not be found in the certified regions. Finally, extending data dependent approaches beyond the  $L_2$  metric such as to defend against  $L_0$  attacks (see for instance the works (Papernot et al., 2016; Combey et al., 2020; Césaire et al., 2021)) seems to us an important challenge.

## Acknowledgements

This work has been supported by the French government under the France 2030 program, as part of the SystemX Technological Research Institute.

## References

- Alfarra, M., Bibi, A., Torr, P. H., and Ghanem, B. Data dependent randomized smoothing. *arXiv preprint arXiv:2012.04351*, 2020.
- Arsigny, V., Fillard, P., Pennec, X., and Ayache, N. Log-euclidean metrics for fast and simple calculus on diffusion tensors. *Magnetic Resonance in Medicine*, 56, 2006.
- Atkinson, C. and Mitchell, A. F. S. Rao’s distance measure. *Sankhyā: The Indian Journal of Statistics, Series A (1961-2002)*, 43(3):345–365, 1981. ISSN 0581572X.
- Biggio, B. et al. Evasion attacks against machine learning at test time. In *Joint European conference on machine learning and knowledge discovery in databases*, pp. 387–402. Springer, 2013.
- Bunel, R. R. et al. A unified view of piecewise linear neural network verification. *Advances in Neural Information Processing Systems*, 31, 2018.
- Cohen, J., Rosenfeld, E., and Kolter, Z. Certified adversarial robustness via randomized smoothing. In *International Conference on Machine Learning*, pp. 1310–1320. PMLR, 2019.
- Combey, T., Loison, A., Faucher, M., and Hajri, H. Probabilistic jacobian-based saliency maps attacks. *Machine Learning and Knowledge Extraction*, 2(4):558–578, 2020. ISSN 2504-4990. URL <https://www.mdpi.com/2504-4990/2/4/30>.
- Césaire, M., Schott, L., Hajri, H., Lamprier, S., and Gallinari, P. Stochastic sparse adversarial attacks. In *2021 IEEE 33rd International Conference on Tools with Artificial Intelligence (ICTAI)*, pp. 1247–1254, 2021. doi: 10.1109/ICTAI52525.2021.00198.
- Deng, J. et al. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 248–255, 2009. doi: 10.1109/CVPR.2009.5206848.
- Eiras, F., Alfarra, M., Kumar, M. P., Torr, P. H., Dokania, P. K., Ghanem, B., and Bibi, A. Ancer: Anisotropic certification via sample-wise volume maximization. *arXiv preprint arXiv:2107.04570*, 2021.
- Huang, X., Kwiatkowska, M., Wang, S., and Wu, M. Safety verification of deep neural networks. In *International conference on computer aided verification*, pp. 3–29. Springer, 2017.
- Jordan, M. and Dimakis, A. G. Exactly computing the local lipschitz constant of relu networks. *Advances in Neural Information Processing Systems*, 33:7344–7353, 2020.
- Krizhevsky, A., Nair, V., and Hinton, G. Cifar-10 (canadian institute for advanced research). URL <http://www.cs.toronto.edu/~kriz/cifar.html>.
- Kumar, A., Levine, A., Goldstein, T., and Feizi, S. Curse of dimensionality on randomized smoothing for certifiable robustness. *International Conference on Machine Learning*. URL <https://par.nsf.gov/biblio/10181769>.
- LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Lecuyer, M., Atlidakis, V., Geambasu, R., Hsu, D., and Jana, S. Certified robustness to adversarial examples with differential privacy. In *2019 IEEE Symposium on Security and Privacy (SP)*, pp. 656–672. IEEE, 2019.
- Lomuscio, A. and Maganti, L. An approach to reachability analysis for feed-forward relu neural networks. *arXiv preprint arXiv:1706.07351*, 2017.
- Madry, A., Makelov, A., Schmidt, L., Tsipras, D., and Vladu, A. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017.
- Miolane, N. et al. Geomstats: A python package for riemannian geometry in machine learning. *Journal of Machine Learning Research*, 21(223):1–9, 2020a. URL <http://jmlr.org/papers/v21/19-027.html>.

- Miolane, N. et al. Introduction to Geometric Learning in Python with Geomstats. In Agarwal, M. and al. (eds.), *SciPy 2020 - 19th Python in Science Conference*, pp. 48–57, July 2020b. doi: 10.25080/Majora-342d178e-007. URL <https://hal.inria.fr/hal-02908006>.
- Papernot, N., Mcdaniel, P., Jha, S., Fredrikson, M., Celik, Z. B., and Swami, A. The limitations of deep learning in adversarial settings. *2016 IEEE European Symposium on Security and Privacy (EuroS&P)*, pp. 372–387, 2016.
- Raghunathan, A., Steinhardt, J., and Liang, P. S. Semidefinite relaxations for certifying robustness to adversarial examples. *Advances in Neural Information Processing Systems*, 31, 2018.
- Salman, H., Li, J., Razenshteyn, I., Zhang, P., Zhang, H., Bubeck, S., and Yang, G. Provably robust deep learning via adversarially trained smoothed classifiers. *Advances in Neural Information Processing Systems*, 32, 2019.
- Skovgaard, L. T. A riemannian geometry of the multivariate normal model. *Scandinavian Journal of Statistics*, 11: 211–223, 1984.
- Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., and Fergus, R. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.
- Tjeng, V., Xiao, K., and Tedrake, R. Evaluating robustness of neural networks with mixed integer programming. *arXiv preprint arXiv:1711.07356*, 2017.
- Weng, L. et al. Towards fast computation of certified robustness for relu networks. In *International Conference on Machine Learning*, pp. 5276–5285. PMLR, 2018.
- Yang, G., Duan, T., Hu, J. E., Salman, H., Razenshteyn, I., and Li, J. Randomized smoothing of all shapes and sizes. In *International Conference on Machine Learning*, pp. 10693–10705. PMLR, 2020.
- Zhai, R., Dan, C., He, D., Zhang, H., Gong, B., Ravikumar, P., Hsieh, C.-J., and Wang, L. Macer: Attack-free and scalable robust training via maximizing certified radius. *arXiv preprint arXiv:2001.02378*, 2020.