



Perspective

A useful taxonomy for adversarial robustness of Neural Networks

Leslie N Smith*

Naval Center for Applied Research in Artificial Intelligence, U.S Naval Research Laboratory,
Washington, D.C 20375, USAReceived: 17 June, 2020
Accepted: 04 August, 2020
Published: 05 August, 2020*Corresponding author: Leslie N Smith, Ph.D, Naval Center for Applied Research in Artificial Intelligence, U.S Naval Research Laboratory, Washington, D.C 20375, USA, E-mail: leslie.smith@nrl.navy.mil

Keywords: Adversarial examples; Adversarial robustness; Computer vision; Machine learning; Neural networks

<https://www.peertechz.com>

Abstract

Adversarial attacks and defenses are currently active areas of research for the deep learning community. A recent review paper divided the defense approaches into three categories; gradient masking, robust optimization, and adversarial example detection. We divide gradient masking and robust optimization differently: (1) increasing intra-class compactness and inter-class separation of the feature vectors improves adversarial robustness, and (2) marginalization or removal of non-robust image features also improves adversarial robustness. By reframing these topics differently, we provide a fresh perspective that provides insight into the underlying factors that enable training more robust networks and can help inspire novel solutions. In addition, there are several papers in the literature of adversarial defenses that claim there is a cost for adversarial robustness, or a trade-off between robustness and accuracy but, under this proposed taxonomy, we hypothesize that this is not universal. We follow this up with several challenges to the deep learning research community that builds on the connections and insights in this paper.

Introduction

With advances in machine learning technology over the past decade, the use of deep neural networks has had great success in computer vision, speech recognition, robotics, and other applications. Along with these remarkable improvements in performance, the recognition of vulnerabilities has also increased. As applications of deep neural networks are increasingly being deployed, the security needs of these applications have come to the foreground, especially for safety-required applications (i.e., self-driving vehicles) and adversarial domains where attacks must be anticipated, such as defense applications.

A recent paper provides a comprehensive review of adversarial attacks and defenses [1] and provides a taxonomy for both the adversarial attacks and defenses. Pulling on the past literature, this review paper defines adversarial examples as “inputs to machine learning models that an attacker intentionally designed to cause the model to make mistakes”. Here, we present a new perspective on adversarial defenses that we believe can provide clarity and inspire novel defenses to adversarial attacks.

The taxonomy of adversarial defense in Xu, et al. [1] consists of three categories: gradient masking, robust optimization, and adversarial detection. Gradient masking includes input data preprocessing (i.e., jpeg compression [2]), thermometer encoding [3], adversarial logit pairing [4], defensive distillation [5], randomization of the deep neural network models (i.e., randomly choosing a model from a set of models [6]) or using dropout [7,8]), and the use of generative models (i.e., PixelDefend [9] and Defense-GAN [10]). The theme of this diverse set of defenses is to make it more difficult to create adversarial examples and attacks but Athalye, et al. [11], demonstrate that gradient masking techniques are ineffective.

The second category in this taxonomy is called robust optimization, and it includes the popular defense method of adversarial training [12], regularization methods that minimize the effects of small perturbations of the input (i.e., Jacobian regularization [13]), and provable defenses (i.e., Reluplex algorithm [14]). Adversarial training is a form of data augmentation where adversarial examples are added to or replace the benign training data. Adversarial training is an important defense discussed in the literature, and variations have been proposed, such as ensemble adversarial training where the adversarial examples are computed from a set

of pretrained classifiers [6]. Robust optimization includes methods for making deep neural networks behave more robustly to the presence of adversarial perturbations in the input, which is the primary focus of our taxonomy in Section 2.

The third category in this review paper is to detect the presence of adversarial examples in the input in order to protect trained classifiers. That is, one can design a separate model to classify if a sample is benign or adversarial. Carlini and Wagner [15], rigorously demonstrate that the properties of adversarial examples are not easy to detect.

For our purposes, we consider adversarial robustness to include all approaches for training networks to improve that network's performance on adversarial examples. We focus primarily on category 2 of the above taxonomy but we also include many of the methods in their category 1. We propose this new taxonomy on adversarial robustness to provide insight to the underlying factors that enable training more robust networks.

In addition, there are several papers in the literature of adversarial attacks and defenses that claim there is a cost for robustness, such that greater robustness requires more data [16], larger model complexity [17] and longer training times. Furthermore, there are claims of trade-offs between robustness and accuracy [18,19], and even robustness and simplicity [20]. There appears to be widespread acceptance of these claims as universal. Another motivation of our work is to demonstrate that these claims are appropriate only for a subset of existing methods for training in adversarial robustness.

While there are other taxonomies mentioned in other papers, they offer only well-known factors for dividing approaches. Guo, et al. [2] divide the work in adversarial robustness into model-specific strategies (i.e., adversarial training [12], regularization methods [13]) and model-agnostic methods (i.e., input preprocessing [21]). Zhang, et al. divide adversarial defense into three categories of data preprocessing [2], gradient masking [11], and adversarial training [12]. Here we reframe the category of making networks adversarially robust in order to provide a fresh perspective and inspire novel solutions in a way these other taxonomies do not.

Our taxonomy

There have been several recent papers showing that using metric learning loss functions during training helps in making neural networks more robust to adversarial examples [22-24]. Mustafa, et al. [23], used their own variation of the contrastive center-loss [25], that encourages both intra-class compactness and inter-class separation of the feature vectors or logits, which are the activations from the last hidden layer. The center loss [26], is a loss function that encourages the feature vectors for each class to lie close to each other (i.e., it encourages intraclass compactness) and the contrastive center-loss function is a generalization of it that also encourages interclass separation. We claim that these works imply a general factor for adversarial robustness, which can be stated as:

Category 1: *Increasing intra-class compactness and inter-class separation of the feature vectors improves adversarial robustness.*

There are several other papers that can be categorized under Category 1. Wu and Yu [27], postulate that the training of deep models decreases the average margin while increasing the minimum margin, and recommend increasing the average margin (i.e., the inter-class separation). Galloway, et al. [28], suggest that batch normalization is a cause of adversarial vulnerability. This aligns with Category 1 because batch normalization constrains the magnitude of the feature vectors (i.e., the activations in the next to the last layer, which is input to the fully connected and softmax layers). Hence, batch normalization limits interclass separation and therefore it can increase adversarial vulnerability.

It is particularly interesting to note that the defensive distillation approach [5], utilizes Category 1. Defensive distillation uses two networks and modifies softmax by dividing by a temperature T , such that $\text{softmax}(Z(\theta, x)/T)$, where $Z(\theta, x)$ is the feature vector, x is the input sample, and θ are the network's weights. In a rigorous paper by Carlini and Wagner [29], they describe the mechanism behind defensive distillation, they state "When we train a distilled network at temperature T and then test it at temperature 1, we effectively cause the inputs to the softmax to become larger by a factor of T ." Since the architecture used in defensive distillation does not contain batch normalization, the average magnitude for the feature vectors increases by T , thereby increasing the inter-class separation. Based on their analysis, we hypothesize that the teacher network (even without distillation) will also show signs of robustness and that adding batch normalization to the architecture or using a feature based attack [30], will break the effectiveness of defensive distillation.

Additionally, there are a number of papers in the literature focused on improving generalization (but not robustness) by increasing intra-class compactness and interclass separation of the feature vectors, such as centerloss [26], contrastive center-loss [25] and lifted structures [31], as well as papers that have appeared recently, such as G-Softmax [32] and Softmax dissection [33]. In our context, generalization refers to the ability of the network to classify images unseen during training and is measured by the gap between the training and testing loss. Category 1 implies that these methods will improve both generalization and robustness.

However, improving both generalization and robustness appears to contradict the conjecture of papers in the literature that suggest there is a trade-off between test accuracy and adversarial robustness [19,20]. This implies the existence of at least one other Category of adversarial robustness where this might be true. One possible set of defenses include image preprocessing [2,22] and gradient masking methods (see [34]). Image preprocessing approaches are based on reducing or eliminating "nonrobust" adversarial perturbations in the training images.

Adversarial perturbations were described as "nonrobust features" by Ilyas, et al. [35]. Ilyas, et al. postulate that



machines use all the image features that are discriminatory between classes (assuming the task is classification), even those features that are invisible to humans. Adversarial training [36], specifically includes training images with non-robust features (i.e., adversarial examples) in order for the network to learn to classify examples with non-robust features properly.

We too believe as described in Ilyas, et al. [35], that humans and machines perform tasks differently. For example, humans are limited in the number of image features they use in making a decision while machines are much less limited. Adversarial examples exist where we expect human performance from a machine. To attain human performance from a machine, we can manually eliminate non-robust features from the training images via preprocessing or make all non-robust image features nondiscriminatory with approaches such as adversarial training.

If we consider the network's training, we realize that as it learns, it averages away the non-discriminatory image features as "nuisance variables". This is analogous to computing the marginal probability by summing or integrating the nuisance variables [34]. Hence, using a bit of inductive reasoning, we hypothesize a second Category for adversarial robustness:

Category 2: Marginalization or removal of non-robust image features improves adversarial robustness.

Many of the papers on adversarial robustness seem to lie within this Category 2, including adversarial training [33] and methods of gradients masking [11]. In addition, we show below with a toy example that there is a tradeoff between accuracy and adversarial robustness [19] for methods that fall under Category 2 (Note: while many of the papers on the trade-off between accuracy and robustness use the adversarial training defense, a similar argument holds for it).

The most obvious way to train a network at human performance levels is to modify the training data to only contain the robust information we want it to use in classification. One extreme way to eliminate non-robust image features is to preprocess the training and test images with an edge detection algorithm to produce binary edge images. These edge images commonly display shape information that humans are able to use to recognize objects. Training a network on edge images results in a highly robust network because all non-visible perturbations have been removed. However, the performance on benign images is reduced due to a decrease in discriminatory information between classes in the edge detection images relative to the original imagery. This example demonstrates the trade-off between accuracy and adversarial robustness. Of course, edge imagery leaves minimal discriminatory information and there is a range of preprocessing that falls on the spectrum between human and machine image features, such as low pass filtering (i.e., DFT [21]), denoising, sparse coding, synthetic imagery, and jpeg compression [2]. Note that it is possible to create examples that can fool even a network trained on edge examples by making large visible changes to the input, but the current definitions of adversarial examples include making small imperceptible changes.

To the best of our knowledge, most of the methods in the literature for attaining adversarial robustness fall under Category 2. The goal of these methods is to marginalize the non-robust features. This explains why training on more data improves the adversarial robustness of deep networks (i.e., increases the likelihood of non-robust features appearing in different classes to be marginalized away as nuisance variables) [37]. This also explains the added adversarial robustness from Jacobian regularization [13], where the loss function trains the network to be invariant to small, non-robust features. It also suggests new methods to obtain adversarial robustness, such as a variant of adversarial training where one adds the same perturbation to images of different classes to make that perturbation non-discriminatory.

Discussion

While we believe that we have presented a few novel connections and insights that we have not seen in the literature, we must still ask if this taxonomy is useful and if so, how.

First, this taxonomy suggests that both robustness and generalization can be improved simultaneously. It clarifies that papers declaring there is a trade-off between robustness and accuracy are misleading because the tradeoff is not universal. We suggest the deep learning community take up the challenge to discover ways to improve both robustness and generalization rather than pursue the current focus of improving robustness at the expense of accuracy. Techniques based on metric learning appear to offer performance improvements in both, and other methods may also exist. Of course, the other side of this challenge is to create new attacks that defeat any new defenses that improve both generalization and robustness.

Second, our paper proposes eliminating non-robust features from the training data so that trained networks learn to only rely on robust image features. But we don't delineate an optimal way to process images to contain only robust features. Obviously binary edge detection images are too extreme as they also eliminate many robust image features. On the other hand, after low pass filtering (i.e., image blurring) non-robust features still remain. The challenge still remains to discover an ideal preprocessing method or combination of methods.

Third, new network training methods can be inspired by our analogy of training to marginalization. For example, data preprocessing and augmentation can insure that non-robust image features are explicitly present in multiple or all classes to insure that the network treats them as non-discriminatory. Similarly, marginalization implies that the training methodology in few-shot meta-learning of changing the tasks every iteration creates more universal features that will be beneficial in transfer learning and perhaps in other scenarios. In addition, the community can investigate better training data combinations that optimally marginalize non-robust features. There is much additional work to be done in this direction to better understand the theoretical and practical aspects of marginalization.

Fourth, the separation of methods for making networks more robust into two Categories implies that methods from



each Category can be productively combined. The combination of methods from each Category should provide different strengths to a network or an ensemble of networks. Combine these with the best methods for each of the categories (see Xu, et al. [38]) and one has an ensemble with the potential to make a solid defense. Unfortunately, the paper with a title “Ensembles of weak defenses are not strong” [38], is misleading because in that paper the authors only tested ensembles of defenses that all fall into a single category, such as detectors or our Category 2 above. He, et al. [38], mention that their “adaptive adversarial examples transfer across several defenses” which might “explain why ensembling is not an effective approach”. It is obvious that each defense in an ensemble must provide strengths that are orthogonal to all the other defenses and an ensemble of many near identical defenses is not useful.

For example, a potential ensemble might include the best adversarial example detector (e.g., Carlini and Wagner [16], found the Bayesian uncertainty estimate of Feinman, et al. [8], to be the strongest of those they tested), as well as a network trained by ensemble adversarial training [6], plus a dropout network that hides the gradient (i.e., Athalye, et al. compare several methods for hiding gradients and found that randomization [7], to be most effective), and networks each from the two Categories in our taxonomy (i.e., one trained with metric learning and another trained on edge detection images, which will force image perturbations to be visible or else they will be eliminated during preprocessing).

We conjecture that a diverse ensemble, with each member offering orthogonal strengths, will be a strictly more powerful defense than any one defense. In this treatise we have stated that each ensemble member described above must possess orthogonal strengths, which might not prove true in practice. However, ablation studies of an ensemble’s members can determine if each member adds to the security of the system. A rigorous analysis of an ensemble’s strengths will also identify its remaining weaknesses and further defense efforts can focus on eliminating these weaknesses.

In addition, we hypothesize in this paper that several of the new methods based on metric learning for improving generalization in the literature [32,33], will also improve robustness. If this is confirmed, there will be numerous other methods in the literature i.e., [25,26], that will improve robustness but have not been demonstrated yet [39].

Conclusions

In this paper we expand the area of adversarial robustness into a taxonomy with two categories; Category 1: increasing intra-class compactness and inter-class separation of the feature vectors improves adversarial robustness, and Category 2: marginalization or removal of nonrobust image features also improves adversarial robustness. This taxonomy permits an understanding of the underlying factors that drive the adversarial robustness of the known methods, and this understanding allows exploring new methods with the same underlying factors.

In addition, we attempt to dispel several potential misunderstandings and set forth several challenges to the deep learning community, such as the discovery of new methods that improve both robustness and generalization. There are also a number of research items left as future work, such as optimal ways to eliminate non-robust features from the training data via preprocessing or to optimally marginalize non-robust features via training.

We also propose that a diverse ensemble of defenses, with each member offering orthogonal strengths, will be a strictly more powerful approach than any one defense. An ensemble of defenses should include all the strongest defenses and should be tested against all of the strongest attacks, in order to find the remaining weaknesses. Then further research on robustness can concentrate on only the remaining holes in the defenses.

We also call on researchers to go further with adversarial defense than is typically done today in the literature. In addition to the challenge of improving both robustness and generalization, researchers can attempt to simultaneously solve multiple other limitations of deep learning, such as reducing the amount of labeled training data needed and creating adaptable networks that learn continuously.

Furthermore, adversarial defenses must go further than working on small imagery such as MNIST and Cifar, which are the most common benchmarks in the adversarial examples literature. The community seems ready to venture into higher resolution imagery of ImageNet and real world imagery, such as satellite imagery.

Eventually, the research and engineering communities will need to investigate adversarial attacks and defenses in the context of safety-required applications (i.e., selfdriving vehicles) and adversarial domains where attacks must be anticipated, such as defense applications. It is only in the context of these applications where complete and secure solutions can be discovered.

Acknowledgement

This work was funded by the Office of Naval Research. The views and conclusions contained in this document are those of the author and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the US Navy.

References

1. Xu H, Ma Y, Liu H, Deb D, Liu H, et al. (2019) Adversarial attacks and defenses in images, graphs and text: A review. [Link: https://bit.ly/31kxaQv](https://bit.ly/31kxaQv)
2. Guo C, Rana M, Cisse M, van der Maaten L (2017) Countering adversarial images using input transformations. [Link: https://bit.ly/31hR7re](https://bit.ly/31hR7re)
3. Buckman J, Roy A, Raffel C, Goodfellow I (2018) Thermometer encoding: One hot way to resist adversarial examples. [Link: https://bit.ly/3i5QQ1t](https://bit.ly/3i5QQ1t)
4. Engstrom L, Ilyas A, Athalye A (2018) Evaluating and understanding the robustness of adversarial logit pairing. [Link: https://bit.ly/3fBm5jq](https://bit.ly/3fBm5jq)
5. Papernot N, McDaniel P, Wu X, Jha S, Swami A (2016) Distillation as a defense to adversarial perturbations against deep neural networks. In



- 2016 IEEE Symposium on Security and Privacy (SP). 582-597. IEEE. [Link: https://bit.ly/31oZeSO](https://bit.ly/31oZeSO)
6. Tramer F, Kurakin A, Papernot N, Goodfellow I, Boneh D, et al. (2017) Ensemble adversarial training: Attacks and defenses. [Link: https://bit.ly/2BXdWrw](https://bit.ly/2BXdWrw)
 7. Dhillon GS, Azizzadenesheli K, Lipton ZC, Bernstein J, Kossaifi J, et al. (2018) Stochastic activation pruning for robust adversarial defense. [Link: https://bit.ly/2DjwhzG](https://bit.ly/2DjwhzG)
 8. Feinman R, Curtin RR, Shintre S, Gardner AB (2017) Detecting adversarial samples from artifacts. [Link: https://bit.ly/311FVd3](https://bit.ly/311FVd3)
 9. Song Y, Kim T, Nowozin S, Ermon S, Kushman N (2017) Pixeldefend: Leveraging generative models to understand and defend against adversarial examples. [Link: https://bit.ly/3i8bw8X](https://bit.ly/3i8bw8X)
 10. Samangouei P, Kabkab M, Chellappa R (2018) Defensegan: Protecting classifiers against adversarial attacks using generative models. [Link: https://bit.ly/33yM6gV](https://bit.ly/33yM6gV)
 11. Athalye A, Carlini N, Wagner D (2018) Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. [Link: https://bit.ly/2Dl8Gie](https://bit.ly/2Dl8Gie)
 12. Goodfellow IJ, Shlens J, Szegedy C (2013) Explaining and harnessing adversarial examples. [Link: https://bit.ly/2Xu8y74](https://bit.ly/2Xu8y74)
 13. Jakubovitz D, Giryas R (2018) Improving dnn robustness to adversarial attacks using jacobian regularization. In Proceedings of the European Conference on Computer Vision (ECCV). 514-529. [Link: https://bit.ly/2Xv0eDV](https://bit.ly/2Xv0eDV)
 14. Carlini N, Katz G, Barrett C, Dill DL (2017) Provably minimally-distorted adversarial examples. [Link: https://bit.ly/2DoHbUO](https://bit.ly/2DoHbUO)
 15. Carlini N, Wagner D (2017) Adversarial examples are not easily detected: Bypassing ten detection methods. In Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security. 3-14. [Link: https://bit.ly/2XuZi2n](https://bit.ly/2XuZi2n)
 16. Schmidt L, Santurkar S, Tsipras D, Talwar K, Madry A (2018) Adversarially robust generalization requires more data. In Advances in Neural Information Processing Systems 5014-5026. [Link: https://bit.ly/33yMfkt](https://bit.ly/33yMfkt)
 17. Madry A, Makelov A, Schmidt L, Tsipras D, Vladu A (2017) Towards deep learning models resistant to adversarial attacks. [Link: https://bit.ly/3i8Vhsb](https://bit.ly/3i8Vhsb)
 18. Tsipras D, Santurkar S, Engstrom L, Turner A, Madry A (2018) Robustness may be at odds with accuracy. [Link: https://bit.ly/30rMifK](https://bit.ly/30rMifK)
 19. Kurakin A, Goodfellow I, Bengio S (2016) Adversarial examples in the physical world. [Link: https://bit.ly/3k71wP4](https://bit.ly/3k71wP4)
 20. Nakkiran P (2019) Adversarial robustness may be at odds with simplicity. [Link: https://bit.ly/2XpZJLi](https://bit.ly/2XpZJLi)
 21. Zhang Z, Jung C, Liang X (2019) Adversarial defense by suppressing high-frequency components. [Link: https://bit.ly/3i6cXoo](https://bit.ly/3i6cXoo)
 22. Pang T, Xu K, Dong Y, Du C, Chen N, et al. (2019) Rethinking softmax cross-entropy loss for adversarial robustness. [Link: https://bit.ly/2DzZWEM](https://bit.ly/2DzZWEM)
 23. Mustafa A, Khan S, Hayat M, Goecke R, Shen J, et al. (2019) Adversarial defense by restricting the hidden space of deep neural networks. [Link: https://bit.ly/30tu1Pj](https://bit.ly/30tu1Pj)
 24. Mao C, Zhong Z, Yang J, Yondrick C, Ray B (2019) Metric learning for adversarial robustness. [Link: https://bit.ly/2Xwl85s](https://bit.ly/2Xwl85s)
 25. Qi C, Su F (2017) Contrastive-center loss for deep neural networks. In 2017 IEEE International Conference on Image Processing (ICIP) 2851-2855. [Link: https://bit.ly/2XvmdY](https://bit.ly/2XvmdY)
 26. Wen Y, Zhang K, Li Z, Qiao Y (2016) A discriminative feature learning approach for deep face recognition. In European conference on computer vision 499-515. [Link: https://bit.ly/2PswWS1](https://bit.ly/2PswWS1)
 27. Wu K, Yu Y (2019) Understanding adversarial robustness: The trade-off between minimum and average margin. [Link: https://bit.ly/2Xu5ebW](https://bit.ly/2Xu5ebW)
 28. Galloway A, Golubeva A, Tanay T, Moussa M, Taylor GW (2019) Batch normalization is a cause of adversarial vulnerability. [Link: https://bit.ly/31hr2Us](https://bit.ly/31hr2Us)
 29. Carlini N, Wagner D (2016) Defensive distillation is not robust to adversarial examples. [Link: https://bit.ly/3kdrdh3](https://bit.ly/3kdrdh3)
 30. Sabour S, Cao Y, Faghri F, Fleet DJ (2015) Adversarial manipulation of deep representations. [Link: https://bit.ly/2XrWbYW](https://bit.ly/2XrWbYW)
 31. Oh Song H, Xiang Y, Jegelka S, Savarese S (2016) Deep metric learning via lifted structured feature embedding. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. 4004-4012. [Link: https://bit.ly/31jHf0e](https://bit.ly/31jHf0e)
 32. Luo Y, Wong Y, Kankanhalli M, Zhao Q (2019) Gsoftmax: Improving intraclass compactness and interclass separability of features. IEEE transactions on neural networks and learning systems. [Link: https://bit.ly/3icnu6](https://bit.ly/3icnu6)
 33. He L, Wang Z, Li Y, Wang S (2019) Softmax dissection: Towards understanding intra-and inter-class objective for embedding learning. [Link: https://bit.ly/3gvAj6A](https://bit.ly/3gvAj6A)
 34. Ilyas A, Santurkar S, Tsipras D, Engstrom L, Tran B, et al. (2019) Adversarial examples are not bugs, they are features. [Link: https://bit.ly/3ic7QTW](https://bit.ly/3ic7QTW)
 35. Szegedy C, Zaremba W, Sutskever I, Bruna J, Erhan D, et al. (2013) Intriguing properties of neural networks. [Link: https://bit.ly/2DdBwBf](https://bit.ly/2DdBwBf)
 36. Goodfellow I, Bengio Y, Courville A (2016) Deep learning. MIT press. [Link: https://bit.ly/3fsgCen](https://bit.ly/3fsgCen)
 37. Sun K, Zhu Z, Lin Z (2019) Towards understanding adversarial examples systematically: Exploring data size, task and model factors. [Link: https://bit.ly/3fz4WGJ](https://bit.ly/3fz4WGJ)
 38. He W, Wei J, Chen X, Carlini N, Song D (2017) Adversarial example defense: Ensembles of weak defenses are not strong. In 11th {USENIX} Workshop on Offensive Technologies. [Link: https://bit.ly/33tM9dx](https://bit.ly/33tM9dx)
 39. Carlini N, Wagner D (2017) Towards evaluating the robustness of neural networks. In 2017 IEEE Symposium on Security and Privacy (SP). IEEE 39-57. [Link: https://bit.ly/311Fwax](https://bit.ly/311Fwax)