



# A Hierarchical Feature Constraint to Camouflage Medical Adversarial Attacks

Qingsong Yao<sup>1</sup>, Zecheng He<sup>2</sup>, Yi Lin<sup>3</sup>, Kai Ma<sup>3</sup>, Yefeng Zheng<sup>3</sup>,  
and S. Kevin Zhou<sup>1,4</sup>(✉)

<sup>1</sup> Key Lab of Intelligent Information Processing of Chinese Academy of Sciences (CAS), Institute of Computing Technology, CAS, Beijing 100190, China

yaoqingsong19@mailsucas.edu.cn

<sup>2</sup> Princeton University, Princeton, USA

zechengh@princeton.edu

<sup>3</sup> Tencent Jarvis Lab, Shenzhen, China

<sup>4</sup> Medical Imaging, Robotics, and Analytic Computing Laboratory and Engineering (MIRACLE) School of Biomedical Engineering and Suzhou Institute for Advanced Research, University of Science and Technology of China, Suzhou 215123, China

**Abstract.** Deep neural networks for medical images are extremely vulnerable to adversarial examples (AEs), which poses security concerns on clinical decision-making. Recent findings have shown that existing medical AEs are easy to detect in feature space. To better understand this phenomenon, we thoroughly investigate the characteristic of traditional medical AEs in feature space. Specifically, we first perform a stress test to reveal the vulnerability of medical images and compare them to natural images. Then, we theoretically prove that the existing adversarial attacks manipulate the prediction by continuously optimizing the vulnerable representations in a fixed direction, leading to outlier representations in feature space. Interestingly, we find this vulnerability is a double-edged sword that can be exploited to help hide AEs in the feature space. We propose a novel hierarchical feature constraint (HFC) as an add-on to existing white-box attacks, which encourages hiding the adversarial representation in the normal feature distribution. We evaluate the proposed method on two public medical image datasets, namely Fundoscopy and Chest X-Ray. Experimental results demonstrate the superiority of our HFC as it bypasses an array of state-of-the-art adversarial medical AEs detector more efficiently than competing adaptive attacks. Our code is available at [https://github.com/qsyao/Hierarchical\\_Feature\\_Constraint](https://github.com/qsyao/Hierarchical_Feature_Constraint).

**Keywords:** Adversarial attack · Adversarial defense

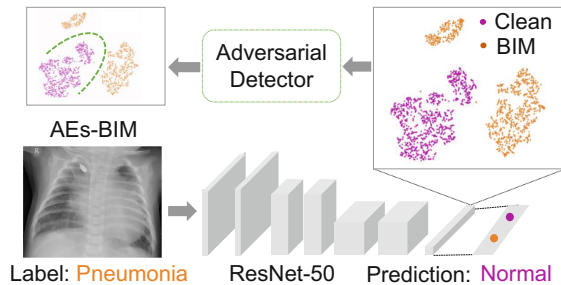
---

**Electronic supplementary material** The online version of this chapter ([https://doi.org/10.1007/978-3-030-87199-4\\_4](https://doi.org/10.1007/978-3-030-87199-4_4)) contains supplementary material, which is available to authorized users.

# 1 Introduction

Deep neural networks (DNNs) are vulnerable to adversarial examples (AEs) [35]. AEs are maliciously generated by adding human-imperceptible perturbations to clean examples, compromising a network to produce the attacker-desired incorrect predictions [6]. The adversarial attack in medical image analysis is disastrous as it can manipulate patients’ disease diagnosis and cause serious subsequent problems. More disturbingly, recent studies have shown that DNNs for medical image analysis, including disease diagnosis [9, 32], organ segmentation [17, 18, 20, 21, 29], and landmark detection [40, 41, 43, 44], are more vulnerable to AEs than natural images.

On the other hand, recent defenses [22, 25] have shown that medical AEs can be easily detected in feature space. We plot the 2D t-SNE [26] in Fig. 1 to illustrate the differences between clean and adversarial features from the penultimate layer of a well-trained pneumonia classifier, revealing that adversarial attacks move the deep representations from the original distribution to extreme outlier positions. As a result, a defender can easily take advantage of this characteristic to distinguish AEs.



**Fig. 1.** We craft AEs by the iterative basic method (BIM) [16] to manipulate the prediction and visualize the penultimate layer’s features of the adversarial and clean examples by 2D t-SNE.

Given this phenomenon, two key questions are investigated in this paper. The first one is: *What causes medical AEs easier to be detected, compared to AEs of natural images?* To better understand this, we conduct both empirical and theoretical analyses. Firstly, we discover that medical features are more vulnerable than natural ones in a stress test, which aims to change the features by the adversarial attack. Then, we theoretically prove that to generate medical AEs, the adversarial features are pushed in a consistent direction towards outlier regions in the feature space where the clean features rarely reside. In consequence, medical adversarial features become outliers.

The second question is: *Can we hide a medical AE from being spotted in the feature space?* A representative adaptive attack selects a guide example and forces the AE feature to be close to the guide feature [33]. However, this does not directly apply to medical AES. Different medical images have different

backgrounds and lesions, making it difficult to manipulate the AE features to be the same as the guide one in all layers within limited perturbation. To hide the adversarial representation at a low price, we propose a novel hierarchical feature constraint (HFC) as an add-on term, which can be plugged into *existing white-box attacks*. HFC first models the normal feature distributions for each layer with a Gaussian mixture model, then encourages the AEs to move to the high-density area by maximizing the log-likelihood of the AE features.

We perform extensive experiments on two public medical diagnosis datasets to validate the effectiveness of HFC. HFC helps medical AE bypass state-of-the-art adversarial detectors, while keeping the perturbations small and unrecognizable to humans. Furthermore, we demonstrate that HFC significantly outperforms other adaptive attacks [1, 2, 33] on manipulating adversarial features. Our experiments support that **the greater vulnerability of medical representations allows an attacker more room for malicious manipulation**.

Overall, we highlight the following contributions:

- We investigate the feature space of medical images and shed light on why medical AEs can be more easily detected, compared with those of natural images.
- We propose a hierarchical feature constraint (HFC), a novel plug-in that can be applied to existing white-box attacks to avoid being detected.
- Extensive experiments validate the effectiveness of our HFC to help existing adversarial attacks bypass state-of-the-art adversarial detectors simultaneously with small perturbations.

## 2 Related Work

**Preliminaries.** Given a clean image  $x$  with its label  $y \in [1, 2, \dots, Y]$  and a DNN classifier  $h$ , the classifier predicts the class of the input example  $y'$  via:

$$y' = \arg \max_k p(k|x) \equiv \frac{\exp(l_k(x))}{\sum_{j=1}^K \exp(l_j(x))}. \quad (1)$$

where the logits  $l_k(x)$  (with respect to class  $k$ ) is given as  $l_k(x) = \sum_{n=1}^N w_{nk} * z_n(x) + b_k$ , in which  $z_n(x)$  is the  $n^{th}$  activation of the penultimate layer that has  $N$  dimensions;  $w_{nk}$  and  $b_k$  are the weights and the bias from the final dense layer, respectively.

**Adversarial Attack.** A common way is to manipulate the classifier’s prediction by minimizing the classification error between the prediction and target class  $c$ , while keeping the AE  $x_{adv}$  within a small  $\epsilon$ -ball of the  $L_p$ -norm [27] centered at the original sample  $x$ , i.e.,  $\|x_{adv} - x\|_p \leq \epsilon$ . In this paper, we focus on typical  $L_\infty$  adversarial attacks, which are most commonly used due to its consistency with respect to human perception [27]. The existing  $L_\infty$  approaches can be categorized into two parts. The first one is gradient-based approaches, such as the fast gradient sign method (FGSM) [10], basic iterative method (BIM) [16], the momentum

iterative method (MIM) [6] and projected gradient descent (PGD) [27]. BIM computes AE as follows:

$$x_0^* = x; \quad x_{t+1}^* = \Pi_\epsilon(x_t^* - \alpha \cdot \text{sign}(\nabla_x J(h(x_t^*), c))). \quad (2)$$

where  $J$  is often chosen as the cross-entropy loss;  $\epsilon$  is the  $L_\infty$  norm bound;  $\alpha$  is the step size; and  $\Pi(\cdot)$  is the project function. Differently from BIM, PGD use a random start  $x_0 = x + U^d(-\epsilon, \epsilon)$ , where  $U^d(-\epsilon, \epsilon)$  is the uniform noise between  $-\epsilon$  and  $\epsilon$ .

The second category is optimization-based methods, among which one of the representative approach is the Carlini and Wanger (CW) attack [3]. According to [27], the  $L_\infty$  version of CW attack can be solved by using the following objective function:

$$\hat{J} = \max(l_c(x_t^*) - l_{y_{\max} \neq c}(x_t^*), -\kappa). \quad (3)$$

where  $l_c$  is the logits with respect to the target class;  $l_{y_{\max} \neq c}$  is the maximum logits of the remaining classes; and  $\kappa$  is a parameter managing the confidence.

**Adversarial Defenses.** Various proactive defenses [12, 36] have been proposed to defend adversarial attacks, such as feature squeezing [39], distillation network [31], JPEG compression [7], gradient masking [30]. Per [5], the adversarial training [10, 27, 38] is the most robust defense, which augments the training set with AEs but consumes too much training time. However, these defenses can be bypassed either completely or partially by adaptive attacks [1, 2, 37]. Different from the challenging proactive defense, reactive defenses have been developed to detect AEs from clean examples with high accuracy [42]. For example, learning-based methods (e.g., RBF-SVM [23], DNN [28]) train a decision boundary between clean and adversarial features, while anomaly-detection based methods directly reject the outlier feature of the AEs. Specifically, kernel density estimation (KDE) [8] and multivariate Gaussian model (MGM) [22] modeled the normal feature distribution. Local intrinsic dimensionality (LID) [24] characterized the dimensional properties of the adversarial subspaces. The degree of the outlier was measured by Mahalanobis distance (MAHA) [19]. Especially, medical AEs were proven to be easy to detect with nearly 100% accuracy by these approaches [22, 24].

### 3 Why Are Medical AEs Easy to Detect?


**Vulnerability of Representations.** A stress test is performed to evaluate the vulnerability of the features of medical and natural images. Specifically, we aim to manipulate the features as much as possible by adversarial attacks. In the test, we decrease ( $\downarrow$ ) and increase ( $\uparrow$ ) the features by replacing the loss function  $J$  in BIM (Eq. 2) by  $J_\downarrow^* = \text{mean}(f^l(x))$  and  $J_\uparrow^* = -\text{mean}(f^l(x))$ , respectively, where  $f^l(x)$  is the feature from the  $l^{\text{th}}$  activation layer. We execute the stress test on the medical image dataset (Fundoscopy [13]) and natural image dataset (CIFAR-10 [15]). The results in Table 1 demonstrate that the features of the medical images can be altered more drastically, i.e., **the medical image representations are much more vulnerable than natural images.**

**Table 1.** Comparison of the robustness between medical images and natural images. We calculate the mean values of the activation layers from ResNet-50 before and after the adversarial attacks. The stress test uses BIM with perturbations under the constraint  $L_\infty = 8/256$ .

Dataset	Fundoscopy			CIFAR-10		
	36	45	48	36	45	48
Normal	.0475	.1910	.3750	.0366	.1660	.1900
Adversarial ( $\downarrow$ )	.0322	.0839	.0980	.0312	.1360	.1480
Adversarial ( $\uparrow$ )	.0842	.7306	2.0480	.0432	.2030	.2640
Difference ( $\downarrow$ )	.0153	.1071	.2770	.0054	.0300	.0420
Difference ( $\uparrow$ )	.0367	.5396	1.6730	.0066	.0370	.0740

**Consistency of Gradient Direction.** To investigate how the feature vulnerability can be exploited by adversarial attacks, we then focus on the loss function and the corresponding gradient on the final logits output  $l_k(x)$ . In each iteration of the  $L_\infty$  attack introduced above,  $J$  and  $\hat{J}$  increase the logits of the target class and decrease the others simultaneously. As a result, gradients point to a similar direction across iteration, which will be back-propagated according to the chain rule.

**Theorem 1.** Consider a binary disease diagnosis network and its representations from the penultimate layer, the directions of the corresponding gradients are fixed during each iteration under adversarial attack.<sup>1</sup>



**Fig. 2.** (a) The similarity between the change values (under adversarial attack) from the penultimate layer and  $(w_{i0} - w_{i1})$ . (b) The similarity of the changes of features from the penultimate layer between different attacks and the different iterations from basic iterative method (BIM).

**Implication.** The partial derivative of cross-entropy loss  $J$  with respect to the activation value  $z_i$  of  $i$ -th node in the penultimate layer is computed as:

$$\nabla_{z_i} J(h(x), y_1) = (1 - p_1)(w_{i0} - w_{i1}), \quad (4)$$

<sup>1</sup> We provide the proof in the supplementary material.

where  $p_1$  denotes the probability of class 1 and  $w_{ij}$  denotes the weight between  $i$ -th node in penultimate layer and  $j$ -th node in the last layer. Accordingly, the activation  $z_i$  with larger  $w_{i0} - w_{i1}$  will increase more (guided by the gradient) under attack. We plot the similarity between the value changes and  $w_{i0} - w_{i1}$  in Fig. 2(a). Similar conclusions can be derived for other attacks. As shown in Fig. 2(b), **the features are pushed toward a similar direction during different iterations of different attacks.**

## 4 Adversarial Attack with a Hierarchical Feature Constraint

We investigate the strategies that the attackers can leverage to bypass detection. Our key intuition is to exploit the vulnerability of medical features in the opposite way, i.e., pushing and hiding it in the distribution of normal medical features.

**Modeling the Distribution of Normal Features.** We model the distribution of normal features using a Gaussian mixture model (GMM).

$$p(f^l(x)) = \sum_{k=1}^K \pi_k \mathcal{N}(f^l(x) | \mu_k, \Sigma_k), \quad (5)$$

where  $p$  is the probability density of sample  $x$  in the target class  $c$ ;  $f^l(\cdot)$  denotes the feature of the  $l^{\text{th}}$  activation layer;  $\pi_k$  is the mixture coefficient subject to  $\sum_{k=1}^K \pi_k = 1$ ;  $\mu_k$  and  $\Sigma_k$  are the mean and covariance matrix of the  $k$ -th Gaussian component in the GMM. These parameters are trained by the expectation-maximization (EM) algorithm [4] on the normal features belonging to the target class  $c$ .

**Hierarchical Feature Constraint (HFC).** We propose a hierarchical feature constraint as a simple-yet-effective add-on that can be applied to existing attacks to avoid being detected. First, we compute the log-likelihood of the AE feature under the distribution of normal features. Specifically, for a given input  $x$ , we separately compute the log-likelihood of an AE feature relative to each component and find the most likely Gaussian component:

$$k' = \arg \max_k \ln(\pi_k \mathcal{N}(f^l(x) | \mu_k, \Sigma_k)). \quad (6)$$

To hide the adversarial representation, we maximize the log-likelihood of this chosen component. Conceptually, it encourages adversarial features in the high-density regions of the distribution of normal features. Then we encourage hiding in all DNN layers. The hierarchical feature constraint induces a loss ( $J_{\text{HFC}}$ ) that is formulated as Eq. 7. It can be used as an add-on to existing attacks by directly adding  $J_{\text{HFC}}$  to any original attack object function  $J_{\text{original}}$ , i.e.,  $J = J_{\text{original}} + J_{\text{HFC}}$ .

$$J_{\text{HFC}} = \sum_{l=1}^L \frac{\lambda^l}{2} (f^l(x) - \mu_{k'}^l)^\top (\Sigma_{k'}^l)^{-1} (f^l(x) - \mu_{k'}^l), \quad (7)$$

where  $\lambda^l$  weights the contribution of constraint in layer  $l$ . The pseudo-code of the adversarial attack boosted with HFC can be found in supplemental material.

## 5 Experiments

### 5.1 Setup

**Datasets.** Following the literature [9], we use two public datasets on typical medical disease diagnosis tasks. The first one is the Kaggle Fundoscopy dataset [13] on the diabetic retinopathy (DR) classification task, which consists of 3,663 fundus images labeled to one of the five levels from “no DR” to “mid/moderate/severe/proliferate DR”. Following [9], we consider all fundoscopies with DR as the same class. The other one is the Kaggle Chest X-Ray [14] dataset on the pneumonia classification task, which consists of 5,863 X-Ray images labeled with “Pneumonia” or “Normal”. We split both datasets into three subsets: *Train*, *AdvTrain* and *AdvTest*. For each dataset, we randomly select 80% of the samples as *Train set* to train the DNN classifier, and treat the left samples as the *Test set*. The incorrectly classified (by the diagnosis network) test samples are discarded. Then we use 70% of the samples (*AdvTrain*) in the *Test set* to train the adversarial detectors and evaluate their effectiveness with the remaining ones (*AdvTest*).

**DNN Models.** We choose the ResNet-50 [11] and VGG-16 [34] models pre-trained on ImageNet. Both models achieve high Area Under Curve (AUC) scores on Fundoscopy and Chest X-Ray datasets: ResNet-50 achieves 99.5% and 97.0%, while VGG-16 achieves 99.3% and 96.5%, respectively.

**Adversarial Attacks and Detectors.** Following [25], we choose three representative attacks, i.e., BIM, PGD and CW, against our models. For the adversarial detectors, we use kernel density (KD) [8], bayesian uncertainty (BU) [8], local intrinsic dimensionality (LID) [24], Mahalanobis distance (MAHA) [19], RBF-SVM [23], and deep neural network (DNN) [28]. The parameters for KD, LID, BU and MAHA are set per the original papers. For KD, BU, and RBF-SVM, we extract features from the penultimate layers. For DNN, we train a classifier for each layer and ensemble by summing up their logits.

**Metrics.** We choose three metrics to evaluate the effectiveness of the adversarial detector and the proposed HFC bypassing method: 1) *True positive rate at 90% true negative rate (TPR@90)*: The detector will drop 10% of the normal samples to reject more adversarial attacks; 2) *Area Under Curve (AUC) score*; 3) *Adversarial accuracy (Adv. Acc)*: The success rate of the targeted adversarial attack against diagnosis network.

**Hyperparameters.** We set<sup>2</sup>  $K = 64$  and 1 for Fundoscopy and Chest X-Ray datasets, respectively. For the  $l^{th}$  layer, we compute the mean value of each

---

<sup>2</sup> The hyperparameter analysis can be found in the supplementary material.

channel separately, and set  $\lambda_l$  to  $\frac{1}{C_l}$ , where  $C_l$  is the number of channels. As tiny perturbations in medical images can cause drastic increase in the loss [25], we set  $\alpha = 0.02/256$  and  $T = 2\epsilon/\alpha$ .

**Table 2.** The point-wise performances of HFC. The metrics scores on the *left* and *right* of the slash are the performances (%) of the adversarial detectors under the attack *without* and *with* HFC, respectively. All of the attacks are evaluated on ResNet-50 with constraint  $L_\infty = 1/256$ .


Fundoscopy	BIM (Adv. Acc = 99.5)		PGD (Adv. Acc = 99.5)		CW (Adv. Acc = 99.5)	
	AUC	TPR@90	AUC	TPR@90	AUC	TPR@90
KD	99.0/74.2	96.8/20.5	99.4/73.4	98.6/13.2	99.5/74.7	99.1/19.6
MAHA	99.6/6.4	99.5/0.0	100/4.2	100/0.0	99.8/33.0	99.5/0.0
LID	99.8/78.3	100/40.6	99.6/73.2	98.6/35.5	98.8/73.4	97.7/33.3
SVM	99.5/28.6	99.1/0.0	99.8/23.1	99.5/0.0	99.8/27.0	99.5/0.0
DNN	100/60.0	100/12.8	100/58.6	100/8.2	100/62.6	100/15.1
BU	58.9/37.4	9.1/0.0	61.9/35.9	9.1/0.0	93.0/32.8	73.1/5.0
Chest X-Ray	BIM (Adv. Acc = 90.9)		PGD (Adv. Acc = 90.9)		CW (Adv. Acc = 98.9)	
	AUC	TPR@90	AUC	TPR@90	AUC	TPR@90
KD	100/73.1	100/6.8	100/82.3	100/50.5	99.2/71.5	98.4/15.7
MAHA	100/0.0	100/0.0	100/0.0	100/0.0	100/22.4	100/0.0
LID	100/48.6	100/1.8	100/49.1	100/1.5	99.2/64.5	98.4/14.4
SVM	100/16.7	100/6.9	100/5.8	100/0.0	100/21.2	100/0.0
DNN	100/31.8	100/0.7	100/33.7	100/0.0	100/61.6	100/5.2
BU	49.9/26.1	19.2/0.0	49.2/26.2	22.7/0.0	98.3/26.2	94.8/0.0

## 5.2 Experimental Results

**Bypassing Adversarial Detectors.** We compare the performances of the existing adversarial attacks *with* and *without* HFC on different DNN classifiers, datasets, and perturbation constraints. The visualization in Fig. 3(a) shows that HFC successfully moves the AE feature (orange) from outlier regions to the regions (cyan) where normal features (purple) reside. Quantitative results in Table 2 show that the HFC boosts all evaluated adversarial attacks to bypass all evaluated detectors simultaneously, with high accuracy and a small perturbation constraint  $L_\infty = 1/256$ . Furthermore, as shown in Fig. 3(b), a larger perturbation budget provides HFC more room to manipulate the representations, compromising the detectors more drastically (as the dotted lines decrease).

**Comparison to Other Adaptive Attacks.** We compare the proposed HFC to other adaptive attacks designed to manipulate the deep representations and bypass the detectors: 1) Generate AEs with internal representation similar to a random guide image [33]; 2) Choose a guide image with its representation closest to the input [33]; 3) Minimize the loss terms of KDE and cross-entropy





**Fig. 3.** (a) Visualization of 2D t-SNE of clean and adversarial features generated from BIM and HFC, extracted from ResNet-50 on Chest X-Ray. (b) The AUC scores in the solid lines and dotted lines are the performances of the adversarial detectors against BIM *with* and *without* HFC.

simultaneously [2]; 4) Minimize the loss term of LID and cross-entropy at the same time [1]. As shown in Table 3, our proposed HFC bypasses the five detectors and greatly outperforms other competitive adaptive attacks.

**Table 3.** AUC scores (%) of the proposed HFC and other adaptive attacks. The AEs are generated on ResNet-50 under constraint  $L_\infty = 1/256$ .

Fundoscopy	KD	MAHA	LID	SVM	DNN	Chest X-Ray	KD	MAHA	LID	SVM	DNN
Random	75.1	86.1	91.7	48.2	93.7	Random	77.0	64.0	91.0	13.0	79.3
Closest	77.0	64.0	91.0	13.0	79.3	Closest	80.1	38.3	71.3	9.9	87.7
KDE	51.6	86.5	90.9	45.3	95.0	KDE	58.2	66.9	71.7	15.3	95.6
LID	87.6	85.4	93.4	61.2	96.2	LID	84.0	66.6	77.1	28.6	96.6
HFC	74.2	6.4	78.3	28.6	60.0	HFC	70.8	0.0	53.6	16.7	32.6

**Table 4.** The performance of BIM *w/* and *w/o* HFC under the semi-white-box setting. The AEs are generated under constraint  $L_\infty = 4/256$ . AUC scores (%) are used as metrics.

DNN model	Fundoscopy	KD	MAHA	LID	SVM	DNN	Adv. Acc
ResNet-50	BIM	98.7	100.0	99.5	92.1	100.0	83.2
	BIM_HFC	78.0	9.1	68.0	16.9	43.8	68.2
VGG-16	BIM	72.3	89.6	81.5	48.1	95.2	88.6
	BIM_HFC	50.9	18.2	64.6	28.8	16.7	73.2

**Semi-white-box Attack.** We evaluate the proposed method in a more difficult scenario: only the architecture information of DNN models is available. The attacker tries to confuse the victim model and bypass its adversarial detectors simultaneously, without knowing model parameters. The results in Table 4 show that our HFC can help BIM compromise most of the detectors and manipulate the DNN models, which poses more disturbing concerns to the safety of DNN-based diagnosis networks.

## 6 Conclusion

In this paper, we investigate the characteristics of medical AEs in feature space. A stress test is performed to reveal the greater vulnerability of medical image features, compared to the natural images. Then we theoretically prove that existing adversarial attacks tend to alter the vulnerable features in a fixed direction. As a result, the adversarial features become outliers and easy to detect. However, an attacker can exploit this vulnerability to hide the adversarial features. We propose a novel hierarchical feature constraint (HFC), a simple-yet-effective add-on that can be applied to existing attacks, to avoid AEs from being detected. Extensive experiments validate the effectiveness of HFC, which also significantly outperforms other adaptive attacks. It reveals the limitation of the current methods for detecting medical AEs in the feature space. We hope it can inspire more defenses in future work.

## References

1. Athalye, A., Carlini, N., Wagner, D.: Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. In: ICLR (2018)
2. Carlini, N., Wagner, D.: Adversarial examples are not easily detected: bypassing ten detection methods. In: Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security, pp. 3–14 (2017)
3. Carlini, N., Wagner, D.: Towards evaluating the robustness of neural networks. In: IEEE Symposium on Security and Privacy, pp. 39–57 (2017)
4. Dempster, A.P., Laird, N.M., Rubin, D.B.: Maximum likelihood from incomplete data via the EM algorithm. *J. R. Stat. Soc.: Ser. B (Methodol.)* **39**(1), 1–22 (1977)
5. Dong, Y., et al.: Benchmarking adversarial robustness. In: CVPR (2020)
6. Dong, Y., et al.: Boosting adversarial attacks with momentum. In: CVPR, pp. 9185–9193 (2018)
7. Dziugaite, G.K., Ghahramani, Z., Roy, D.M.: A study of the effect of JPG compression on adversarial images. arXiv preprint [arXiv:1608.00853](https://arxiv.org/abs/1608.00853) (2016)
8. Feinman, R., Curtin, R.R., Shintre, S., Gardner, A.B.: Detecting adversarial samples from artifacts. arXiv preprint [arXiv:1703.00410](https://arxiv.org/abs/1703.00410) (2017)
9. Finlayson, S.G., Chung, H.W., Kohane, I.S., Beam, A.L.: Adversarial attacks against medical deep learning systems. *Science* **363**(6433), 1287–1289 (2018)
10. Goodfellow, I.J., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples. In: ICLR (2015)
11. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR, pp. 770–778 (2016)
12. He, X., Yang, S., Li, G., Li, H., Chang, H., Yu, Y.: Non-local context encoder: Robust biomedical image segmentation against adversarial attacks. In: AAAI, vol. 33, pp. 8417–8424 (2019)
13. Kaggle: APTOS 2019 Blindness Detection (2019). <https://www.kaggle.com/c/aptos2019-blindness-detection>
14. Kaggle: Chest X-Ray Images (Pneumonia) (2019). <https://www.kaggle.com/paultimothymooney/chest-xray-pneumonia>
15. Krizhevsky, A.: Learning multiple layers of features from tiny images. University of Toronto, May 2012

16. Kurakin, A., Goodfellow, I., Bengio, S.: Adversarial machine learning at scale. In: ICLR (2017)
17. Ji, W., et al.: Uncertainty quantification for medical image segmentation using dynamic label factor allocation among multiple raters. In: MICCAI on QUBIQ Workshop (2020)
18. Ji, W., et al.: Learning calibrated medical image segmentation via multi-rater agreement modeling. In: CVPR, pp. 12341–12351, June 2021
19. Lee, K., Lee, K., Lee, H., Shin, J.: A simple unified framework for detecting out-of-distribution samples and adversarial attacks. In: ICLR, pp. 7167–7177 (2018)
20. Li, H., et al.: High-resolution chest x-ray bone suppression using unpaired CT structural priors. *IEEE Trans. Med. Imaging* **39**, 3053–3063 (2020)
21. Li, H., Han, H., Zhou, S.K.: Bounding maps for universal lesion detection. In: Martel, A.L., et al. (eds.) MICCAI 2020. LNCS, vol. 12264, pp. 417–428. Springer, Cham (2020). [https://doi.org/10.1007/978-3-030-59719-1\\_41](https://doi.org/10.1007/978-3-030-59719-1_41)
22. Li, X., Zhu, D.: Robust detection of adversarial attacks on medical images. In: IEEE International Symposium on Biomedical Imaging, pp. 1154–1158. IEEE (2020)
23. Lu, J., Issaranon, T., Forsyth, D.: SafetyNet: detecting and rejecting adversarial examples robustly. In: ICCV, October 2017
24. Ma, X., et al.: Characterizing adversarial subspaces using local intrinsic dimensionality. In: ICLR (2018)
25. Ma, X., et al.: Understanding adversarial attacks on deep learning based medical image analysis systems. *Pattern Recogn.* **110**, 107332 (2020)
26. Maaten, L.V.D., Hinton, G.: Visualizing data using t-SNE. *J. Mach. Learn. Res.* **9**(Nov), 2579–2605 (2008)
27. Madry, A., Makelov, A., Schmidt, L., Tsipras, D., Vladu, A.: Towards deep learning models resistant to adversarial attacks. In: ICLR (2018)
28. Metzen, J.H., Genewein, T., Fischer, V., Bischoff, B.: On detecting adversarial perturbations. In: ICLR (2017)
29. Ozbulak, U., Van Messem, A., De Neve, W.: Impact of adversarial examples on deep learning models for biomedical image segmentation. In: Shen, D., et al. (eds.) MICCAI 2019. LNCS, vol. 11765, pp. 300–308. Springer, Cham (2019). [https://doi.org/10.1007/978-3-030-32245-8\\_34](https://doi.org/10.1007/978-3-030-32245-8_34)
30. Papernot, N., McDaniel, P., Goodfellow, I., Jha, S., Celik, Z.B., Swami, A.: Practical black-box attacks against machine learning. In: ASIA Computer and Communications Security, pp. 506–519 (2017)
31. Papernot, N., McDaniel, P., Wu, X., Jha, S., Swami, A.: Distillation as a defense to adversarial perturbations against deep neural networks. In: IEEE Symposium on Security and Privacy, pp. 582–597. IEEE (2016)
32. Paschali, M., Conjeti, S., Navarro, F., Navab, N.: Generalizability vs. robustness: investigating medical imaging networks using adversarial examples. In: Frangi, A.F., Schnabel, J.A., Davatzikos, C., Alberola-López, C., Fichtinger, G. (eds.) MICCAI 2018. LNCS, vol. 11070, pp. 493–501. Springer, Cham (2018). [https://doi.org/10.1007/978-3-030-00928-1\\_56](https://doi.org/10.1007/978-3-030-00928-1_56)
33. Sabour, S., Cao, Y., Faghri, F., Fleet, D.J.: Adversarial manipulation of deep representations. In: IEEE Symposium on Security and Privacy (2016)
34. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. In: ICLR (2015)
35. Szegedy, C., et al.: Intriguing properties of neural networks. In: ICLR (2014)
36. Taghanaki, S.A., Abhishek, K., Azizi, S., Hamarneh, G.: A kernelized manifold mapping to diminish the effect of adversarial perturbations. In: CVPR, pp. 11340–11349 (2019)

37. Tramer, F., Carlini, N., Brendel, W., Madry, A.: On adaptive attacks to adversarial example defenses. In: ICLR (2020)
38. Tramèr, F., Kurakin, A., Papernot, N., Goodfellow, I., Boneh, D., McDaniel, P.: Ensemble adversarial training: attacks and defenses. In: ICLR (2018)
39. Xu, W., Evans, D., Qi, Y.: Feature squeezing: detecting adversarial examples in deep neural networks. In: Network and Distributed System Security Symposium (2017)
40. Yao, Q., He, Z., Han, H., Zhou, S.K.: Miss the point: targeted adversarial attack on multiple landmark detection. In: Martel, A.L., et al. (eds.) MICCAI 2020. LNCS, vol. 12264, pp. 692–702. Springer, Cham (2020). [https://doi.org/10.1007/978-3-030-59719-1\\_67](https://doi.org/10.1007/978-3-030-59719-1_67)
41. Yao, Q., Xiao, L., Liu, P., Zhou, S.K.: Label-free segmentation of COVID-19 lesions in lung CT. *IEEE Trans. Med. Imaging* (2020)
42. Zheng, Z., Hong, P.: Robust detection of adversarial attacks by modeling the intrinsic properties of deep neural networks. In: Advances in Neural Information Processing Systems, pp. 7913–7922 (2018)
43. Zhou, S.K., et al.: A review of deep learning in medical imaging: imaging traits, technology trends, case studies with progress highlights, and future promises. *Proc. IEEE* **109**(5), 820–838 (2021)
44. Zhou, S.K., Rueckert, D., Fichtinger, G.: Handbook of Medical Image Computing and Computer Assisted Intervention. Academic Press, Cambridge (2019)