Universal adversarial attacks on deep neural networks for medical image classification

Hokuto Hirano[1], Akinori Minagi[1], Kazuhiro Takemoto[1*]
1) Department of Bioscience and Bioinformatics, Kyushu Institute of Technology, Iizuka, Fukuoka 820-8502, Japan
**Corresponding author's e-mail: takemoto@bio.kyutech.ac.jp*

Abstract

Background. Deep neural networks (DNNs) are widely investigated in medical image classification to achieve automated support for clinical diagnosis. It is necessary to evaluate the robustness of medical DNN tasks against adversarial attacks, as high-stake decision making will be made based on the diagnosis. Several previous studies have considered simple adversarial attacks. However, the vulnerability of DNNs to more realistic and higher risk attacks have not been evaluated yet, i.e., universal adversarial perturbation (UAP), which is a single perturbation that can induce DNN failure in most classification tasks.

Methods. We focus on three representative DNN-based medical image classification tasks (i.e., skin cancer, referable diabetic retinopathy, and pneumonia classifications) and investigate their vulnerability of DNNs with various model architectures to UAPs.

Results. We demonstrate that the DNNs are vulnerable to both nontargeted UAPs, which cause a task failure resulting in an input being assigned an incorrect class, and to targeted UAPs, which cause the DNN to classify an input into a specific class. The almost imperceptible UAPs achieved > 80% success rates for nontargeted and targeted attacks. The vulnerability to UAPs barely depended on model architecture. Moreover, we discovered that adversarial retraining, which is known to be an effective method for adversarial defenses, increased the robustness of DNNs against UAPs in only limited cases.

Conclusion. Unlike previous assumptions, the results indicate that DNN-based clinical diagnosis is easier to deceive because of adversarial attacks. Adversaries can result in failed diagnoses at lower costs (e.g., without consideration of data distribution); moreover, they can affect the diagnosis. The effects of adversarial defenses may be not limited. Our findings emphasize that more careful consideration is required in developing DNNs for medical imaging and their practical applications.

Keywords: deep neural networks, medical imaging, adversarial attacks, security and privacy

Background

Deep neural networks (DNNs) are effective for image classification and are beginning to be applied to medical image diagnosis to empower physicians and accelerate decision making in clinical environments [1]. For example, DNNs have been used for classifying skin cancer from photographic images [2], referable diabetic retinopathy from optical coherence tomography (OCT) images of the retina [3], and pneumonia from chest X-ray images [3]; they have demonstrated high diagnostic performances. A meta-analysis [4] has indicated that the diagnostic performance of DNNs is equivalent to that of healthcare professionals.

Despite the high performance of DNNs, practical applications of DNNs to disease diagnosis are still debatable. High-stake decision-making will be made from disease diagnosis. Complex classifiers, including DNNs, can potentially cause catastrophic harm to society because they are often difficult to interpret [5]. More importantly, DNNs have a number of security concerns [6]; specifically, DNNs are known to be vulnerable to adversarial examples [7, 8], which are input images that cause misclassifications by DNNs and typically generated by adding specific, imperceptible perturbations to original input images that have been correctly classified using DNNs. The existence of adversarial examples questions the generalization ability of DNNs, reduces model interpretability, and limits the applications of deep learning in safety- and security-critical environments [9]; in particular, the adversarial examples cause not only misdiagnosis, but also various social disturbances [10]. The vulnerability of DNNs to adversarial attacks has been claimed in skin cancer classification [10] and pneumonia classification based on chest X-ray images [11].

Nevertheless, more focused investigations are required the vulnerability of DNNs to adversarial attacks. Previous studies have only considered input-dependent adversarial attacks (i.e., an individual adversarial perturbation is used such that each input image misclassifies). Such adversarial attacks are difficult tasks because they require high computational costs. More realistic adversarial attacks must be further considered. Notably, a single small perturbation that can induce DNN failure in most image classification tasks has been reported [12]. Such perturbations are called *universal adversarial perturbations (UAPs)*, as they are image agnostic. A previous study [12] has considered only UAPs for nontargeted attacks, which cause misclassification (i.e., a task failure resulting in an input image being assigned an incorrect class). However, we previously extended the algorithm for generating UAPs to enable targeted attacks [13], which caused the DNN to classify an input image into a specific class. UAPs are difficult to detect as such perturbations are extremely small and hence do not significantly affect data distributions. UAP-based adversarial attacks can be more straightforward to implement by adversaries in real-world environments. UAPs are vulnerable to security threats in medical image diagnosis; however, the vulnerability of UAPs is still poorly evaluated in DNN-based medical image diagnosis to date. In addition, defense strategies against UAPs are still poorly investigated although the vulnerability of DNNs to adversarial attacks indicates the need for strategies to address security concerns (i.e., adversarial defense [8]). Specifically, adversarial retraining is one of the few approaches that could not be defeated thus far [14].

The aim of this study to evaluate the vulnerability of DNNs to UAPs for medical image classification. We focused on representative medical image classifications: skin cancer classification from photographic images [2], referable diabetic retinopathy classification from OCT images [3], and pneumonia classification from chest X-ray images [3]. We obtained DNN models with various architectures for medical image diagnosis and investigated the vulnerability to nontargeted and targeted attacks based on UAPs. Moreover, adversarial defense was considered; in particular, we evaluated the increased robustness of DNNs to nontargeted and targeted UAPs using adversarial retraining [12, 14–16], a representative method for adversarial defenses.

Methods

Medical image datasets

We used three types of medical images: skin lesion images for skin cancer classification, OCT images for referable diabetic retinopathy classification, and chest X-ray images for pneumonia classification.

In a previous study [2], skin lesion images (Red-Green-Blue color) were obtained from the International Skin Imaging Collaboration (ISIC) 2018 dataset (challenge2018.isic-archive.com/task3/training/), in which the images were classified into seven classes: melanoma (MEL), melanocytic nevus (NV), basal cell carcinoma (BCC), actinic keratosis/Bowens disease (intraepithelial carcinoma; AKIEC), benign keratosis (solar lentigo/seborrheic keratosis/lichen planus-like keratosis; BKL), dermatofibroma (DF), and vascular lesion (VASC). The dataset comprised 10,015 images. We randomly divided these images into 7,000 training images (778 MEL images, 4689 NV images, 370 BCC images, 229 AKIEC images, 764 BKL images, 76 DF images, 94 VASC images) and 3,015 test images (335 MEL images, 2016 NV images, 144 BCC images, 98 AKIEC images, 335 BKL images, 39 DF images, 48 VASC images).

The OCT images and chest X-ray images (grayscale) were obtained from a previous study [3] (data.mendeley.com/datasets/rscbjbr9sj/3). The OCT images were classified into four classes: choroidal neovascularization with neovascular membrane and associated subretinal fluid (CNV), diabetic macular edema with retinal-thickening-associated intraretinal fluid (DME), multiple drusen present in early age-related macular degeneration (DRUSEN), and normal retina with preserved foveal contour and absence of any retinal fluid/edema (NM). The original dataset comprised 37,455 CNV images, 11,598 DME images, 8,866 DRUSEN images, and 51,390 NM images. We constructed a class-balanced training image set and test image set by randomly selecting 1,960 images and 840 images per class without duplicates, respectively. We finally obtained 7,840 training images and 3,360 test images.

The chest X-ray images were classified into binary classes: no pneumonia (NORMAL) and viral or bacterial pneumonia (PNEUMONIA). The original dataset comprised 1,583 NORMAL images and 4,273 PNEUMONIA images. We constructed a class-balanced training image set and test image set by randomly selecting 900 and 270 images per class without duplicates, respectively. We finally obtained 1,800 training images and 540 test images.

Transfer learning methods

Based on previous studies [2, 3], we obtained the DNN models using transfer learning; in particular, we fine-tuned DNN models pretrained using the ImageNet dataset [17] with a medical image dataset. We mainly considered the Inception V3 architecture [18], following previous studies. To investigate the effect of model architecture on adversarial robustness, we considered different model architectures: VGG16 [19], VGG19 [19], ResNet50 [20], Inception ResNet V2 [21], DenseNet 121 [22], and DenseNet 169 [22]. For each model architecture, we replaced the original last fully connected (FC) layer with a new FC layer in which the output size is the number of classes. The images were resized to 299 × 299 pixels. All layer parameters were fine-tuned using the training images in a medical image dataset. We used the stochastic gradient descent optimizer with a momentum of 0.9. The batch size and number of epochs were set to 32 and 50, respectively. The learning rates were scheduled based on the number of epochs: 0.001 for ≤ 40 epochs, 1e–4 for 41–45 epochs, and 1e–5 for > 45 epochs. To avoid overfitting, data augmentation was considered: random image rotations with the angle ranging between –5° and 5° and random image shifts with 5% of the height and 5% of width. For skin cancer classification, we adapted oversampling to account for imbalances in the dataset. The transfer learning procedures were performed using Keras (version 2.2.4; keras.io).

Universal adversarial perturbations

Simple iterative algorithms [12, 13] were used to generate UAPs for nontargeted and targeted attacks. The details of the algorithms are described in [12, 13]. We used the nontargeted UAP algorithm available in the Adversarial Robustness 360 Toolbox (ART) [23] (version 1.0; github.com/Trusted-AI/adversarial-robustness-toolbox). The targeted UAP algorithm was implemented by modifying the nontargeted UAP algorithm in the ART in our previous study [13] (github.com/hkthirano/targeted_UAP_CIFAR10).

The algorithms consider a classifier and generate nontargeted (targeted) UPAs $\boldsymbol{\rho}$ from an input image set $\boldsymbol{X}$, under the constraint that the L_p norm of the perturbation is equal to or less than a small ξ value (i.e., $\|\boldsymbol{\rho}\|_p \leq \xi$). The algorithms start with $\boldsymbol{\rho} = \mathbf{0}$ (no perturbation) and iteratively update $\boldsymbol{\rho}$ by additively obtaining an adversarial perturbation for an input image $\boldsymbol{x}$, which is randomly selected from $\boldsymbol{X}$ without replacement. These iterative updates continue until the number of iterations reaches the maximum $i_{\max}$.

The fast gradient sign method (FGSM) [7] is used to obtain an adversarial perturbation for the input image; meanwhile, the original UAP algorithm [12] uses the DeepFool method [24]. This is because the FGSM is used for both nontargeted and targeted attacks, and DeepFool requires a higher computational cost than the FGSM and only generates a nontargeted adversarial example for the input image. The FGSM generates the adversarial perturbation for $\boldsymbol{x}$ based on the loss gradient [7], with attack strength parameter ϵ.

Nontargeted and targeted UAPs were generated using the training images in the dataset. Parameter ϵ was set to 0.0024; cases where $p = 2$ and ∞ were considered. Parameter ξ was determined based on ratio ζ of the L_p norm of the UAP to the average L_p norm of an image in the dataset. For the ISIC 2018 (skin lesion image) dataset, the average L_∞

and L_2 norms were 237 and 85,662, respectively. For the OCT image dataset, the average L_∞ and L_2 norms were 253 and 15,077, respectively. For the chest X-ray image dataset, the average L_∞ and L_2 norms were 253 and 40,738, respectively. Parameter $i_{\max}$ was set to 15.

To compare the performances of the generated UAPs with those of the random controls, we generated random vectors (random UAPs) sampled uniformly from the sphere of a specified radius [12].

Vulnerability evaluation

The fooling rate R_f and targeted attack success rate R_s were computed to evaluate the vulnerability of the DNN models to nontargeted UAPs and targeted UAPs, respectively. R_f for an image set is defined as adversarial images for which predicted labels are inconsistent with the labels predicted from their associated clean images to all images in the set (i.e., the probability that the labels predicted from clean images are inconsistent with the label predicted from their adversarial images). R_s for an image set is the proportion of adversarial images classified into the target class to all images in the set. It is noteworthy that R_s has a baseline, defined as R_s observed without UAPs. Class (label) composition in image data and prediction performance of DNNs affect the baseline. In this study, for the OCT and chest X-ray image datasets, the baselines of R_s of targeted UAPs to a specified class were ~25% and ~50%, respectively. For the skin lesion dataset, the baselines of R_s of targeted UAPs to MEL and NV were ~10% and ~65%, respectively.

Additionally, we obtained the confusion matrices to evaluate the change in prediction due to UAPs for each class. The confusion matrices were normalized to account for an imbalanced dataset (ISIC 2018 dataset).

Adversarial retraining

Adversarial retraining was performed to increase the robustness of the DNN models to UAPs [12, 15]; in particular, the models were fine-tuned with adversarial images. The procedure was described in a previous study [12]. A brief description is as follows. i) Ten UAPs against a DNN model were generated using the algorithm (for generating a nontargeted or targeted UAP) with the (clean) training image set. ii) A modified training image set was obtained by randomly selecting half of the training images and combining them with the remaining, in which each image was perturbed by a UAP randomly selected from 10 UAPs. iii) The model was fine-tuned by performing five additional epochs of training on the modified training image set. iv) A new UAP against the fine-tuned model was generated using the algorithm with the training image set. v) The R_f and R_s of the UAP for the test images were computed. Steps i)–v) were repeated five times.

Results

Performance for medical image classification

The test and training accuracies of seven DNN models for three medical image datasets

are summarized in Table S1 in Addition file 1. The DNN models achieved good accuracies. For the skin lesion image, OCT image, chest X-ray image datasets, the test accuracies averaged over seven models were 87.3%, 95.8%, and 98.4%, respectively; specifically, the test accuracies of Inception V3 models, which were frequently used in previous studies on medical image diagnosis (e.g., [2, 3]), were 87.7%, 95.5%, and 97.6%, respectively. The normalized confusion matrices for the Inception V3 models on the test images are shown in Fig. S1 in Addition file 1.

Nontargeted universal adversarial attacks

Meanwhile, the DNN models showed vulnerability to nontargeted UAPs. The fooling rates R_f for both the training and test images increased rapidly with the perturbation magnitude ζ and reached a high R_f, despite a low ζ (2%–6%). Figure 1 shows the case of nontargeted UAPs with $p = 2$ against the Inception V3 models. The UAPs with $\zeta = 4\%$ achieved R_f of > 80% for the skin lesion (Fig. 1A) and chest X-ray image datasets (Fig. 1C), whereas slightly larger UAPs (with $\zeta = 6\%$) were required to achieve R_f of ~70% for the OCT image dataset (Fig. 1B). The R_f of the UAPs was significantly higher than those of random UAPs. The confusion matrices on test images show that the models classified most images into several specific classes (i.e., dominant classes) due to the UAPs for the skin lesion and OCT image datasets. Specifically, most skin lesion images tended to be classified into AKIEC or DF (Fig. 1D); moreover, most OCT images were classified into CNV (Fig. 1E). For the chest X-ray image dataset, the model wrongly predicted the true labels (Fig. 1F). A high R_f at a low ζ and dominant labels were observed in the case of UAP with $p = \infty$ against the Inception V3 models for all medical image datasets (Fig. S2 in Addition file 1). However, the skin lesion images tended to be classified into more broad classes: BCC, AKIEC, BKL, or DF (Fig. S2D in Addition file 1).

The vulnerability to nontargeted UAPs is typically observed in other models. Table 1 shows the R_f of the UAPs against the DNN models for the test images in the medical image datasets. Overall, the small UAPs ($\zeta = 4\%$ for the skin lesion and chest X-ray image datasets, and $\zeta = 6\%$ for the OCT image dataset) achieved a high R_f (70%–90%). The R_f of the UAPs were significantly higher than those of random UAPs. R_f may depend on model architectures; specifically, the R_f of the UAPs against the VGG16 and VGG19 models were ~50% for the chest X-ray image dataset, whereas those of the UAPs against the other models were almost between 70% and 80%. This indicates that the models classified images into either NORMAL or PNEUMONIA. In the case of UAPs with $p = 2$, the VGG16 and VGG19 models classified most test images into PNEUMONIA and NORMAL, respectively (Fig. S3 in Addition file 1). In the case of UAPs with $p = \infty$, both VGG16 and VGG19 models predicted most test images as NORMAL. This indicates that the confusion matrix patterns (dominant classes) might change according to the model architecture and p. Additionally, the change in confusion matrix patterns (on test images) was observed in the skin lesion and OCT image datasets. For example, the VGG16 model classified most skin lesion images into BKL owing to the UAP with $\zeta = 4\%$ and $p = 2$ (Figure S4A), whereas the Inception V3 models classified them into AKIEC or DF (Fig. 1D). The ResNet 50 model classified most OCT images into DME owing to the UAP with $\zeta = 6\%$ and $p = 2$ (Fig. S4B in Addition file 1), whereas Inception V3 models classified

them into CNV (Fig. 1E).

The nontargeted UAPs (with $\zeta = 4\%$ for the skin lesion and chest X-ray image datasets, and with $\zeta = 6\%$ for the OCT image dataset) were almost imperceptible. Figure 2 shows the nontargeted UAPs with $p = 2$ against the Inception V3 models and examples of adversarial images for the medical image datasets. The models predicted the original images as their actual classes; however, they classified the adversarial images into any incorrect class due to the nontargeted UAPs. The UAPs with $p = \infty$ and against the other DNN models were also almost imperceptible for the skin lesion (Fig. S5 in Addition file 1), OCT (Fig. S6 in Addition file 1), and chest X-ray image datasets (Fig. S7 in Addition file 1). The UAP patterns tended to depend on the model architecture for each medical image dataset (Figs. S5–S7 in Addition file 1). The transferability of UAPs, which indicates that UAPs generated based on DNNs with one model architecture can be used to deceive DNNs with another model architecture, was not confirmed for the OCT (Table S3 in Addition file 1) and chest X-ray image datasets (Table S4 in Addition file 1); however, a weak transferability of UAPs was observed in the skin lesion image dataset (Table S5 in Addition file 1). Specifically, the nontargeted UAPs with $p = 2$ generated based on the Inception V3 models achieved R_f of ~45%, ~2%, and ~10% on average against the DNNs with another model architecture for the skin lesion, OCT, and chest X-ray image datasets, respectively.

Targeted universal adversarial attacks

Vulnerability to targeted UAPs was confirmed. Table 2 shows the targeted attack success rates R_s of the UAPs with $p = 2$ against the DNN models for the test images in the medical image datasets. As representative examples, we considered targeted attacks to the most significant case and the control in each medical image dataset. For the skin lesion image dataset, targeted attacks to MEL and NV were considered for the skin lesion image dataset. For the OCT image dataset, targeted attacks to CNV and NM were considered. For the chest X-ray image dataset, targeted attacks to PNEUMONIA and NORMAL were considered. Overall, a high (> 85%) R_s was observed regardless of the model architecture despite small UAPs (with $\zeta = 2\%$ for the skin lesion and chest X-ray image datasets, and with $\zeta = 6\%$ for the OCT image dataset). Furthermore, the confusion matrices (Fig.3) indicate that the UAP-based targeted attacks were successful: most ($R_s\%$ of) test images were classified into the targeted class owing to the UAPs. However, a smaller R_s was partly observed according to the model architectures and datasets. For the skin lesion image dataset, the R_s of the UAPs against the VGG16 (~40%) and VGG19 (~65%) models were lower than those (~90%) of the UAPs against the other models. For the targeted attacks to NM in the OCT image dataset, the R_s (30%–40%) of the UAPs against the VGG and DensNet models were lower than those (~85%) of the UAPs against the other models. The R_s of random UAPs were almost equivalent to those of the baselines. The R_s of the UAPs were significantly higher than those of the random UAPs. Furthermore, a high R_s for a small ζ were observed for the targeted UAPs with $p = \infty$ (Table S2 in Addition file 1). However, the R_s for targeted attacks to MEL were lower overall, compared with the R_s of the UAPs with $p = 2$. For example, the R_s of the UAPs with $p = 2$ and $p = \infty$ against the Inception V3 model were ~95% and ~75%, respectively.

The targeted UAPs (with $\zeta = 2\%$ for the skin lesion and chest X-ray image datasets, and with $\zeta = 6\%$ for the OCT image dataset) were also almost imperceptible. Figure 3 shows the targeted UAPs with $p = 2$ against the Inception V3 models and examples of adversarial images for the medical image datasets. The models predicted the original images as their actual classes; however, they classified the adversarial images into the targeted classed owing to the UAPs. The UAPs with $p = \infty$ and against the other DNN models were also almost imperceptible. For the skin lesion image dataset, Figures S8 and S9 show the targeted attacks to NV and MEL, respectively. For the OCT image dataset, Figures S10 and S11 in Addition file 1 show the targeted attacks to NM and CNV, respectively. For the chest X-ray image dataset, Figures S12 and S13 in Addition file 1 show the targeted attacks to NORMAL and PNEUMONIA, respectively. The UAP patterns tended to depend on the model architecture for each medical image dataset (Figs. S8–S13 in Addition file 1). The transferability of UAPs was not confirmed for the skin lesion (Table S6), OCT (Table S7 in Addition file 1), and chest X-ray image datasets (Table S8 in Addition file 1); specifically, the R_s observed when the targeted UAPs with $p = 2$ generated based on the Inception V3 model that attacked the DNN models with another architecture was almost equivalent to their baselines of R_s (~10%, ~25%, and ~50%, for the skin lesion, OCT, and chest X-ray image datasets, respectively).

Adversarial retraining

We analyzed the usefulness of adversarial retraining against universal adversarial attacks. Figure 5 shows the effect of adversarial retraining on the R_f of nontargeted UAPs with $p = 2$ against the Inception V3 models for the skin lesion, OCT, and chest X-ray image datasets. $\zeta = 4\%$ for the skin lesion and chest X-ray image datasets; $\zeta = 6\%$ for the OCT image dataset. Adversarial retraining did not affect the test accuracy. For the OCT image dataset, R_f decreased with the iterations of adversarial retraining; specifically, R_f decreased from 70.2% to 13.1 % after five iterations (Fig. 5B); however, ~40% of the NM images were still classified into an incorrect class (DME; Fig. 5E). The effect of adversarial retraining on R_f was limited for the skin lesion (Fig. 5A) and chest X-ray image datasets (Fig. 5B). For the chest X-ray image dataset, R_f decreased from 81.7% to 46.7%. A R_f of ~50% indicates that the model classified most images into either one of two classes; specifically, most images were classified into NORMAL at the fifth iteration (Fig. 5F). For the skin lesion image dataset, any remarkable decrease in R_f due to adversarial retraining was not confirmed; specifically, R_f decreased from 92.2% to 82.1% (Fig. 5A). Most images were classified into MEL at the fifth iteration (Fig. 5C). However, the dominant classes changed for each iteration. For example, the dominant classes were AKIEC and BKL at the third and fourth iterations, respectively (Fig. S14 in Addition file 1).

Figure 6 shows the effect of adversarial retraining on the R_s of targeted UAPs with $p = 2$ against the Inception V3 models for the skin lesion, OCT, and chest X-ray image datasets. As representative examples, we considered targeted attacks to the most significant cases: MEL, CNV, and PNEUMONIA for the skin lesion, OCT, and chest X-ray image datasets, respectively. $\zeta = 2\%$ for the skin lesion and chest X-ray image datasets; $\zeta = 6\%$ for the OCT image dataset. Adversarial retraining did not affect the test accuracy and reduced R_s for all medical image datasets (Figs. 6A–6C). For the OCT and chest X-ray image dataset,

R_s decreased from ~95% to the baseline R_s (~25% and ~50%, respectively) after five iterations. For the skin lesion image dataset, R_s decreased from ~95% to ~30%; however, R_s at the fifth iteration was higher than the baseline (~10%). The confusion matrices (Figs. 6D–6F) indicated that adversarial retraining was useful against UAP-based targeted attacks: most images were correctly classified into the original classes despite the adversarial attacks. However, the effect of adversarial retraining was partially limited for the skin lesion image dataset. For example, 30% of the NV images were still classified into the target class (MEL) despite five iterations of adversarial retraining (Fig. 6C). Furthermore, ~20% of BKL and VASC images were still classified into the target class.

Discussion

We showed the vulnerability of the DNN models for medical image classification to small UAPs. Furthermore, previous studies [10, 11] have indicated the vulnerability to adversarial attacks toward medical DNNs; however, they were limited to input image-dependent perturbations. In this study, we demonstrated that almost imperceptible UAPs caused DNN misclassifications. Unlike previous assumptions, the results indicate that a DNN-based medical image diagnosis is easier to deceive. Adversaries can result in failed DNN-based medical image diagnoses at lower costs (i.e., using a single perturbation); specifically, they do not need to consider the distribution and diversity of input images when attacking DNNs using UAPs, as UPAs are image agnostic. The vulnerability to UAPs was confirmed in various model architectures. The vulnerability to UAPs may be a universal feature in DNNs.

We demonstrated that nontargeted attacks based on UAPs were possible (Figs. 1 and 2; Table 1). Most images were classified into a few specific classes for the skin lesion and OCT image (multiclass) datasets. This result is consistent with the existence of dominant classes in UAP-based nontargeted attacks [12]. For the skin lesion image dataset, the dominant classes of AKIEC and DF observed in this study may be due to the imbalanced dataset. The numbers of AKIEC and DF images are relatively fewer than those of other class images. As the algorithm considers maximizing the R_f, a relatively large R_f is achieved when all inputs are classified into AKIEC and DF due to UAPs. The use of imbalanced datasets may be one of the causes of vulnerability to UAPs. For the OCT image (binary-class) dataset, the DNN models wrongly predicted the actual labels because of R_f maximization; however, the existence of dominant classes was partly confirmed according to the model architecture. These misclassifications result in false positives and false negatives in medical diagnosis. False positives may cause unwanted mental stress to patients, whereas false negatives may result in significant misdiagnoses involving human lives; specifically, they fail to perform early detection and render therapeutic strategies difficult. Moreover, they can cause the social credibility of medical doctors and medical organizations to be undermined.

The transferability of nontargeted UAPs across model architectures was limited (Tables S3–S5 in Addition file 1). This indicates that UAPs are architecture specific, inconsistent with a previous study [12]. This discrepancy might be due to differences in the image datasets. Specifically, the number of classes (2–7) in the medical image datasets was lower than that (1,000) in the dataset in the previous study. This study partly considered grayscale

images, whereas the previous study used colored images. Transferability may be observed in datasets comprising colored images with more classes. In fact, a weak transferability was observed for the skin lesion image dataset (Table S5 in Addition file 1).

Furthermore, we showed that targeted attacks based on UAPs were possible in medical image diagnosis (Figs. 3 and 4; Table 2), although the UAPs were not transferable across model architectures (Tables S6–S8 in Addition file 1). The results imply that adversaries can control DNN-based medical image diagnoses. As targeted attacks are more realistic, they may result in more significant security concerns compared with nontargeted attacks. In particular, adversaries can obtain any diagnosis; specifically, they can intentionally cause not only problems resulting from misdiagnosis, but also various social disturbances. As mentioned in a previous study [10], adversarial attacks can be used for insurance fraud as well as drug and device approval adjustments, thereby fraudulently providing and obtaining high-quality care when DNNs are used for decision making.

We considered adversarial retraining, which is known to be an effective method for adversarial defenses [14], to reduce the vulnerability to UAPs. However, the effect of adversarial retraining was limited for nontargeted UAPs (Fig. 5). For targeted attacks, adversarial retraining reduced the vulnerability to UAPs significantly but did not completely avoid it (particularly for the skin lesion image dataset; Fig. 6). In addition, adversarial retraining requires high computational costs, as it is an iterative fine-tuning method. Alternative simple methods such as dimensionality reduction (e.g., principle component analysis), distributional detection (e.g., maximum mean discrepancy), and normalization detection (e.g., dropout randomization) are available; however, they are known to be easily detected as adversarial examples [15]. Despite the recent development in adversarial defenses such as regularized surrogate loss optimization [25] and the use of a discontinuous activation function [26], many promising defense methods have failed [27]. Defending against adversarial attacks is a cat-and-mouse game [10]. Furthermore, properties inherent to image processing may cause misclassifications. For instance, DNN-based image reconstructions are often performed for purifying adversarial examples [28]; however, it causes image artifacts, resulting in misclassifications by DNNs [29]. It may be difficult to completely avoid security concerns caused by adversarial attacks.

A simple solution for avoiding adversarial attacks is to render DNNs closed source and publicly unavailable; however, this hinders the accelerated development of medical DNNs and practical applications of DNNs to automated support for clinical diagnosis. Because the amount of medical image data is limited, collaboration among multiple institutions is required to achieve high diagnosis performance [30]. For similar reasons, medical DNNs are often developed by fine-tuning existing DNNs such as VGG, ResNet, and Inception pretrained using the ImageNet dataset (i.e., via transfer learning), although a previous study [30] debated the effect of transfer learning on the improvement in prediction performance for medical imaging; consequently, model architectures and model weighs may be important. Furthermore, DNNs are aimed for real-world usage (e.g., automated support for clinical diagnosis). The assumption that DNNs are closed source and publicly unavailable may be unrealistic. Even if DNNs are black-box (e.g., model architectures and weights are unknown and loss gradient is not accessible), adversarial attacks on DNNs may be possible. Several methods for adversarial attacks on black-box DNNs, which estimate adversarial

perturbations using only model outputs (e.g., confidence scores), have been proposed [31–33]. The development and operation of secure, privacy-preserving, and federated DNNs are required in medical imaging [6].

Conclusion

We demonstrated the vulnerability of DNNs for medical image classification to nontargeted and targeted UAPs. Our findings emphasized that careful consideration is required in developing DNNs for medical imaging and their practical applications. Our study enhances our understanding of the vulnerabilities of DNNs to adversarial attacks and may help increase the security of DNNs. UAPs are useful for reliability evaluation and for designing the operation strategy of medical DNNs.

List of abbreviations

AKIEC: actinic keratosis/Bowens disease (intraepithelial carcinoma)

BCC: basal cell carcinoma

BKL: benign keratosis (solar lentigo/seborrheic keratosis/lichen planus-like keratosis)

CNV: neovascular membrane and associated subretinal fluid

DF: dermatofibroma

DME: diabetic macular edema with retinal-thickening-associated intraretinal fluid

DNN: deep neural network

DRUSEN: multiple drusen present in early age-related macular degeneration

DensNet: dense convolutional network

FC: fully connected

FGSM: fast gradient sign method

ISIC: International Skin Imaging Collaboration

MEL: melanoma

NM: normal retina with preserved foveal contour and absence of any retinal fluid/edema

NV: melanocytic nevus

OCT: optical coherence tomography

ResNet: residual network

UAP: universal adversarial perturbation

VASC: vascular lesion

VGG: visual geometry group

Declarations

Ethics approval and consent to participate

Not applicable.

Consent for publication

Not applicable.

Availability of data and material

All data generated and analyzed during this study are included in this published article and its supplementary information files. The code and data used in this study are available from our GitHub repository: github.com/hkthirano/MedicalAI-UAP.

Competing interests

The authors declare that they have no competing interests.

Funding

No funding was received.

Authors' contributions

KT conceived and designed the study. HH and AM prepared the data and models. HH coded and performed experimental evaluation. HH and KT interpreted the results. HH and KT wrote the manuscript. All authors provided the final approval for publication.

Acknowledgments

We would like to thank Editage (www.editage.jp) for English language editing.

References

1. Litjens G, Kooi T, Bejnordi BE, Setio AAA, Ciompi F, Ghafoorian M, et al. A survey on deep learning in medical image analysis. Med Image Anal. 2017;42 December 2012:60–88. doi:10.1016/j.media.2017.07.005.

2. Esteva A, Kuprel B, Novoa RA, Ko J, Swetter SM, Blau HM, et al. Dermatologist-level classification of skin cancer with deep neural networks. Nature. 2017;542:115–8. doi:10.1038/nature21056.

3. Kermany DS, Goldbaum M, Cai W, Valentim CCS, Liang H, Baxter SL, et al. Identifying Medical Diagnoses and Treatable Diseases by Image-Based Deep Learning. Cell. 2018;172:1122-1131.e9. doi:10.1016/j.cell.2018.02.010.

4. Liu X, Faes L, Kale AU, Wagner SK, Fu DJ, Bruynseels A, et al. A comparison of deep learning performance against health-care professionals in detecting diseases from medical imaging: a systematic review and meta-analysis. Lancet Digit Heal. 2019;1:e271–97. doi:10.1016/S2589-7500(19)30123-2.

5. Rudin C. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. Nat Mach Intell. 2019;1:206–15. doi:10.1038/s42256-019-0048-x.

6. Kaissis GA, Makowski MR, Rückert D, Braren RF. Secure, privacy-preserving and federated machine learning in medical imaging. Nat Mach Intell. 2020;2:305–11. doi:10.1038/s42256-020-0186-1.

7. Goodfellow IJ, Shlens J, Szegedy C. Explaining and harnessing adversarial examples. 2014. http://arxiv.org/abs/1412.6572.

8. Yuan X, He P, Zhu Q, Li X. Adversarial examples: attacks and defenses for deep learning. IEEE Trans Neural Networks Learn Syst. 2019;30:2805–24. doi:10.1109/TNNLS.2018.2886017.

9. Matyasko A, Chau L-P. Improved network robustness with adversary critic. 2018. http://arxiv.org/abs/1810.12576.

10. Finlayson SG, Bowers JD, Ito J, Zittrain JL, Beam AL, Kohane IS. Adversarial attacks on medical machine learning. Science (80-). 2019;363:1287–9. doi:10.1126/science.aaw4399.

11. Asgari Taghanaki S, Das A, Hamarneh G. Vulnerability Analysis of Chest X-Ray Image Classification Against Adversarial Attacks. In: Understanding and Interpreting Machine Learning in Medical Image Computing Applications. 2018. p. 87–94. doi:10.1007/978-3-030-02628-8_10.

12. Moosavi-Dezfooli SM, Fawzi A, Fawzi O, Frossard P. Universal adversarial perturbations. Proc - 30th IEEE Conf Comput Vis Pattern Recognition, CVPR 2017. 2017;2017-Janua:86–94.

13. Hirano H, Takemoto K. Simple iterative method for generating targeted universal adversarial perturbations. In: Proceedings of 25th International Symposium on Artificial Life and Robotics. 2020. p. 426–30. http://arxiv.org/abs/1911.06502.

14. Madry A, Makelov A, Schmidt L, Tsipras D, Vladu A. Towards Deep Learning Models Resistant to Adversarial Attacks. In: International Conference on Learning Representations. 2018. https://openreview.net/forum?id=rJzIBfZAb.

15. Carlini N, Wagner D. Adversarial examples are not easily detected. In: Proceedings

of the 10th ACM Workshop on Artificial Intelligence and Security - AISec '17. New York, New York, USA: ACM Press; 2017. p. 3–14. doi:10.1145/3128572.3140444.

16. Wong E, Rice L, Kolter JZ. Fast is better than free: Revisiting adversarial training. In: International Conference on Learning Representations. 2020. https://openreview.net/forum?id=BJx040EFvH.

17. Russakovsky O, Deng J, Su H, Krause J, Satheesh S, Ma S, et al. ImageNet Large Scale Visual Recognition Challenge. Int J Comput Vis. 2015.

18. Szegedy C, Vanhoucke V, Ioffe S, Shlens J, Wojna Z. Rethinking the Inception Architecture for Computer Vision. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). IEEE; 2016. p. 2818–26. doi:10.1109/CVPR.2016.308.

19. Simonyan K, Zisserman A. Very deep convolutional networks for large-scale image recognition. In: 3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings. 2015.

20. He K, Zhang X, Ren S, Sun J. Deep Residual Learning for Image Recognition. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). IEEE; 2016. p. 770–8. doi:10.1109/CVPR.2016.90.

21. Szegedy C, Ioffe S, Vanhoucke V, Alemi AA. Inception-v4, inception-ResNet and the impact of residual connections on learning. In: 31st AAAI Conference on Artificial Intelligence, AAAI 2017. 2017.

22. Huang G, Liu Z, Van Der Maaten L, Weinberger KQ. Densely connected convolutional networks. In: Proceedings - 30th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017. 2017.

23. Nicolae M-I, Sinn M, Tran MN, Buesser B, Rawat A, Wistuba M, et al. Adversarial Robustness Toolbox v1.0.0. 2018. http://arxiv.org/abs/1807.01069.

24. Moosavi-Dezfooli S-M, Fawzi A, Frossard P. DeepFool: a simple and accurate method to fool deep neural networks. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). IEEE; 2016. p. 2574–82. doi:10.1109/CVPR.2016.282.

25. Zhang H, Yu Y, Jiao J, Xing E, Ghaoui L El, Jordan M. Theoretically Principled Trade-off between Robustness and Accuracy. In: Chaudhuri K, Salakhutdinov R, editors. Proceedings of the 36th International Conference on Machine Learning. Long Beach, California, USA: PMLR; 2019. p. 7472–82. http://proceedings.mlr.press/v97/zhang19p.html.

26. Xiao C, Zhong P, Zheng C. Enhancing adversarial defense by k-winners-take-all. Proc 8th Int Conf Learn Represent. 2020. http://arxiv.org/abs/1905.10510.

27. Croce F, Hein M. Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks. Proc 37th Int Conf Mach Learn. 2020. doi:2003.01690.

28. Hwang U, Park J, Jang H, Yoon S, Cho NI. PuVAE: a variational autoencoder to purify adversarial examples. IEEE Access. 2019;7:126582–93. doi:10.1109/ACCESS.2019.2939352.

29. Antun V, Renna F, Poon C, Adcock B, Hansen AC. On instabilities of deep learning in image reconstruction and the potential costs of AI. Proc Natl Acad Sci. 2020;:201907377. doi:10.1073/pnas.1907377117.

30. Chang K, Balachandar N, Lam C, Yi D, Brown J, Beers A, et al. Distributed deep learning networks among institutions for medical imaging. J Am Med Informatics Assoc. 2018;25:945–54.

31. Chen J, Su M, Shen S, Xiong H, Zheng H. POBA-GA: Perturbation optimized black-box adversarial attacks via genetic algorithm. Comput Secur. 2019;85:89–106. doi:10.1016/j.cose.2019.04.014.

32. Guo C, Gardner JR, You Y, Wilson AG, Weinberger KQ. Simple black-box adversarial attacks. Proc 36th Int Conf Mach Learn. 2019;:2484–93. http://arxiv.org/abs/1905.07121.

33. Co KT, Muñoz-González L, de Maupeou S, Lupu EC. Procedural noise adversarial examples for black-box attacks on deep convolutional networks. In: Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security. New York, NY, USA: ACM; 2019. p. 275–89. doi:10.1145/3319535.3345660.

Tables

Table 1: Fooling rates R_f (%) of nontargeted UAPs against various DNN models for test images of skin lesion, OCT, and chest X-ray image datasets. $\zeta = 4\%$ for skin lesion and chest X-ray image datasets. $\zeta = 6\%$ for OCT image dataset. Values in brackets are R_f of random UAPs (random controls).

Model architecture	Skin lesion		OCT		Chest X-ray	
	$p = 2$	$p = \infty$	$p = 2$	$p = \infty$	$p = 2$	$p = \infty$
Inception V3	92.2 (14.1)	90.0 (11.8)	70.2 (1.0)	73.9 (3.4)	81.7 (2.4)	79.8 (3.0)
VGG16	87.6 (4.9)	86.4 (3.5)	72.4 (0.2)	74.9 (1.8)	49.8 (2.2)	50.0 (2.2)
VGG19	89.2 (5.2)	87.0 (3.7)	72.8 (0.4)	74.7 (2.1)	49.3 (3.9)	49.3 (4.4)
ResNet50	91.9 (11.6)	87.9 (10.1)	71.2 (1.1)	74.8 (5.4)	72.6 (7.2)	73.0 (7.4)
Inception ResNet V2	94.5 (16.7)	90.3 (15.2)	69.6 (1.4)	74.0 (3.2)	78.0 (2.6)	77.0 (3.3)
DenseNet 121	93.8 (12.0)	82.9 (10.2)	68.8 (1.3)	73.0 (3.6)	69.8 (3.9)	71.7 (4.1)
DenseNet 169	93.8 (11.7)	84.2 (9.1)	50.3 (1.3)	72.3 (4.0)	67.6 (2.8)	71.3 (3.7)

Table 2: Targeted attack success rates R_s (%) of targeted UAPs with $p = 2$ against various DNN models to each target class. R_s were for test images. $\zeta = 2\%$ for skin lesion and chest X-ray image datasets. $\zeta = 6\%$ for OCT image dataset. Values in brackets are R_s of random UAPs (random controls).

Model architecture / Target class	Skin lesion		OCT		Chest X-ray	
	NV	MEL	NM	CNV	NORMAL	PNEUMONIA
Inception V3	93.3 (65.6)	94.4 (12.2)	84.1 (25.7)	95.9 (24.8)	96.1 (52.8)	93.3 (47.2)
VGG16	89.6 (71.7)	40.4 (8.3)	32.4 (25.4)	97.7 (24.9)	95.6 (50.2)	95.0 (49.8)
VGG19	91.6 (72.1)	64.6 (8.7)	41.2 (25.9)	97.5 (24.9)	97.6 (51.7)	95.2 (48.3)
ResNet50	97.9 (66.5)	92.4 (11.8)	84.9 (25.8)	98.5 (24.5)	95.7 (53.5)	95.2 (46.5)
Inception ResNet V2	92.4 (61.0)	97.3 (16.1)	84.5 (25.6)	96.2 (24.7)	98.3 (53.1)	93.9 (46.9)
DenseNet 121	92.1 (65.2)	90.5 (13.4)	41.8 (25.3)	88.1 (24.7)	94.8 (51.9)	92.0 (48.1)
DenseNet 169	92.9 (65.8)	92.9 (12.2)	41.7 (25.0)	92.7 (24.2)	95.7 (52.0)	93.1 (48.0)

Figure captions

Fig. 1: Vulnerability to nontargeted UAPs with $p = 2$. Line plots of the fooling rate R_f against Inception V3 model versus perturbation magnitude ζ for skin lesion (A), OCT (B), and chest X-ray (C) image datasets. Legend label indicates image set used for computing R_f. Additional argument "(random)" indicates that random UAPs were used instead of UAPs. Normalized confusion matrices for Inception V3 models attacked using UAPs on test images of skin lesion (D), OCT (E), and chest X-ray (F) image datasets are also shown. $\zeta = 4\%$ in (D) and (F). $\zeta = 6\%$ in (E).

Fig. 2: Nontargeted UAPs with $p = 2$ against Inception V3 models and their adversarial images for skin lesion (A), OCT (B), and chest X-ray image datasets (C). $\zeta = 4\%$ in (A) and (C). $\zeta = 6\%$ in (B). Labels in brackets beside the images are the predicted classes. The original (clean) images are correctly classified into their actual labels. UAPs are emphatically displayed for clarity; in particular, each UAP is scaled by a maximum of 1 and minimum of 0.

Fig. 3: Normalized confusion matrices for Inception V3 models attacked with targeted UPAs with $p = 2$ on test images in skin lesion (left panels), OCT (middle panels), and chest X-ray image datasets (right panels). $\zeta = 2\%$ for skin lesion and chest X-ray image datasets. $\zeta = 6\%$ for OCT image dataset.

Fig. 4: Targeted UAPs with $p = 2$ against Inception V3 models and their adversarial images for skin lesion (A), OCT (B), and chest X-ray image datasets. $\zeta = 2\%$ in (A) and (C). $\zeta = 6\%$ in (B). Labels in brackets beside the images are predicted classes. Original (clean) images were correctly classified into their actual labels. Adversarial images were classified into the target classes. UAPs are emphatically displayed for clarity; in particular, each UAP is scaled by a maximum of 1 and minimum of 0.

Fig. 5: Effect of adversarial retraining on robustness of nontargeted UAPs with $p = 2$ against Inception V3 models for skin lesion, OCT, and chest X-ray image datasets. $\zeta = 4\%$ for the skin lesion and chest X-ray image datasets. $\zeta = 6\%$ for OCT image dataset. Top panels indicate scatter plots of fooling rate R_f (%) of UAPs versus number of iterations for adversarial retraining. Bottom panels indicate normalized confusion matrices for fine-tuned models obtained after five iterations of adversarial retraining. These confusion matrices are on adversarial test images.

Fig. 6: Effect of adversarial retraining on robustness of targeted UAPs with $p = 2$ against Inception V3 models for skin lesion, OCT, and chest X-ray image dataset. $\zeta = 2\%$ for skin lesion and chest X-ray image datasets. $\zeta = 6\%$ for OCT image dataset. Top panels indicate scatter plots of targeted attack success rate R_s (%) of UAPs versus number of iterations for adversarial retraining. Bottom panels indicate normalized confusion matrices for fine-tuned models obtained after five iterations of adversarial retraining. These confusion matrices are on adversarial test images.

Additional files

Additional file 1: Supplementary tables and figures. (PDF)