



PAPER

Universal adversarial perturbations for CNN classifiers in EEG-based BCIs

RECEIVED
4 March 2021REVISED
12 June 2021ACCEPTED FOR PUBLICATION
28 June 2021PUBLISHED
15 July 2021Zihan Liu^{1,4}, Lubin Meng^{1,4}, Xiao Zhang¹, Weili Fang² and Dongrui Wu^{1,3,*} ¹ Key Laboratory of the Ministry of Education for Image Processing and Intelligent Control, School of Artificial Intelligence and Automation, Huazhong University of Science and Technology, Wuhan 430074, People's Republic of China² School of Design and Environment, National University of Singapore, 117566 Singapore³ Zhejiang Lab, Hangzhou 311121, People's Republic of China⁴ These authors contributed equally to this work.

* Author to whom any correspondence should be addressed.

E-mail: drwu@hust.edu.cn**Keywords:** brain-computer interface, convolutional neural network, electroencephalogram, universal adversarial perturbation**Abstract**

Objective. Multiple convolutional neural network (CNN) classifiers have been proposed for electroencephalogram (EEG) based brain-computer interfaces (BCIs). However, CNN models have been found vulnerable to universal adversarial perturbations (UAPs), which are small and example-independent, yet powerful enough to degrade the performance of a CNN model, when added to a benign example. *Approach.* This paper proposes a novel total loss minimization (TLM) approach to generate UAPs for EEG-based BCIs. *Main results.* Experimental results demonstrated the effectiveness of TLM on three popular CNN classifiers for both target and non-target attacks. We also verified the transferability of UAPs in EEG-based BCI systems. *Significance.* To our knowledge, this is the first study on UAPs of CNN classifiers in EEG-based BCIs. UAPs are easy to construct, and can attack BCIs in real-time, exposing a potentially critical security concern of BCIs.

1. Introduction

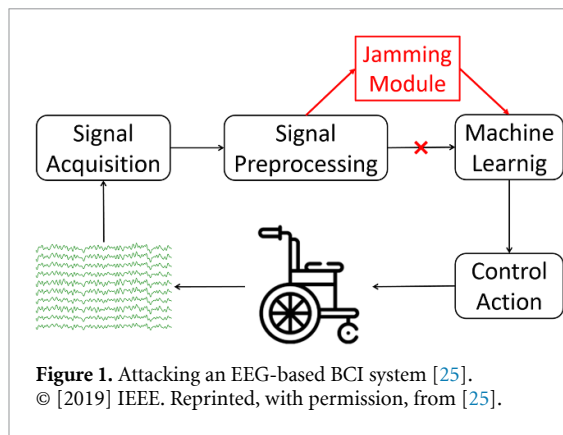
A brain-computer interface (BCI) enables people to interact directly with a computer using brain signals. Due to its low-cost and convenience, electroencephalogram (EEG), which records the brain's electrical activities from the scalp, has become the most widely used input signal in BCIs. There are several popular paradigms in EEG-based BCIs, e.g. P300 evoked potentials [1–4], motor imagery (MI) [5], steady-state visual evoked potentials (SSVEP) [6], etc.

Deep learning, which eliminates manual feature engineering, has become increasingly popular in decoding EEG signals in BCIs. Multiple convolutional neural network (CNN) classifiers have been proposed for EEG-based BCIs. Lawhern *et al* [7] proposed EEGNet, a compact CNN model demonstrating promising performance in several EEG-based BCI tasks. Schirrneister *et al* [8] proposed a deep CNN model (DeepCNN) and a shallow CNN model (ShallowCNN) for EEG classification. There were also studies that converted EEG signals to spectrograms or topoplots and then fed them into deep learning

classifiers [9–11]. This paper focuses on CNN classifiers which take raw EEG signals as the input, but our approach should also be extendable to other forms of inputs.

Albeit their promising performance, it was found that deep learning models are vulnerable to adversarial attacks [12, 13], in which deliberately designed tiny perturbations can significantly degrade the model performance. Many successful adversarial attacks have been reported in image classification [14–17], speech recognition [18], malware detection [19], etc.

According to the purpose of the attacker, adversarial attacks can be categorized into two types: non-target attacks and target attacks. In a non-target attack, the attacker wants the model output to an adversarial example (after adding the adversarial perturbation) to be wrong, but does not force it into a particular class. Two typical non-target attack approaches are DeepFool [20] and universal adversarial perturbation (UAP) [21]. In a target attack, the model output to an adversarial example should always be biased into a specific wrong class.



Some typical target attack approaches are the iterative least-likely class method [15], adversarial transformation networks [22], and projected gradient descent (PGD) [23]. There are also approaches that can be used in both non-target and target attacks, e.g. L-BFGS [12], the fast gradient sign method (FGSM) [14], the C&W method [24], the basic iterative method [15], etc.

An EEG-based BCI system usually consists of four parts: signal acquisition, signal preprocessing, machine learning, and control action. Zhang and Wu [25] explored the vulnerability of CNN classifiers under adversarial attacks in EEG-based BCIs, and discovered that adversarial examples do exist there. They injected a jamming module between signal preprocessing and machine learning to perform the adversarial attack, as shown in figure 1. They successfully attacked three CNN classifiers (EEGNet, DeepCNN, and ShallowCNN) in three different scenarios (white-box, gray-box, and black-box). Their results exposed a critical security problem in EEG-based BCIs, which had not been investigated before. As pointed out in [25], ‘*EEG-based BCIs could be used to control wheelchairs or exoskeleton for the disabled [26], where adversarial attacks could make the wheelchair or exoskeleton malfunction. The consequence could range from merely user confusion and frustration, to significantly reducing the user’s quality of life, and even to hurting the user by driving him/her into danger on purpose. In clinical applications of BCIs in awareness evaluation/detection for disorder of consciousness patients [26], adversarial attacks could lead to misdiagnosis.*

Albeit their success, Zhang and Wu’s approaches [25] had the following limitations:

- An adversarial perturbation needs to be computed specifically for each input EEG trial, which is inconvenient.
- To compute the adversarial perturbation, the attacker needs to wait for the complete EEG trial to be collected; however, by that time the EEG trial has gone, and the perturbation cannot be actually added to it. So, Zhang and Wu’s approaches are theoretically important, but may

not be easily implementable in practice. For better practicability, we should be able to perform the attack as soon as an EEG trial starts.

This paper introduces UAP for BCIs, which is a universal perturbation template computed offline and added to any EEG trial in real-time. Compared with Zhang and Wu’s approaches [25], it has two corresponding advantages:

- A UAP is computed once and applicable to any EEG trial, instead of being computed specifically for each input EEG trial.
- A UAP can be added as soon as an EEG trial starts, or anywhere during an EEG trial, thus the attacker does not need to know the number of EEG channels, the starting time, and the length of a trial.

So, it relieves the two limitations of [25] simultaneously, making the attack more practical⁴.

Studies on UAPs appeared in the literature very recently. Moosavi-Dezfooli *et al* [21] discovered the existence of UAPs, and verified that they can fool state-of-the-art machine learning models in image classification. Their method for crafting the UAPs, based on DeepFool [20], solves a complex optimization problem. The same idea was later used in attacking speech recognition systems [28]. Behjati *et al* [29] proposed a gradient projection based approach for generating UAPs in text classification. Mopuri *et al* [30] proposed a generalizable and data-free approach for crafting UAPs, which is independent of the underlying task. All these approaches were for non-target attacks. Target attacks using UAPs are more challenging, because the perturbation needs to be both universal and targeting at a particular class. To our knowledge, the only study on UAPs for target attacks was [31], where Hirano *et al* integrated a simple iterative method for generating non-target UAPs and FGSM for target attacks to generate UAPs for target attacks.

This paper investigates UAPs in EEG-based BCIs. We make the following three contributions:

- To our knowledge, this is the first study on UAPs for EEG-based BCIs, which make adversarial attacks in BCIs more convenient and more practical.
- We propose a novel total loss minimization (TLM) approach for generating a UAP for EEG

⁴ Our most recent research [27] also developed adversarial perturbation templates to accommodate causality, but it considered traditional P300 and SSVEP based BCI speller pipelines, which perform feature extraction and classification separately, whereas in this paper the UAPs are developed for end-to-end deep learning classifiers in P300 and MI tasks.

trials, which can achieve better attack performance with a smaller perturbation, compared with the traditional DeepFool based approach.

- (c) Our proposed TLM can perform both non-target attacks and target attacks. To our knowledge, no one has studied optimization based UAPs for target attacks before.

The remainder of this paper is organized as follows: section 2 introduces two approaches to generate UAPs for EEG trials. Section 3 describes our experimental setting. Sections 4 and 5 present the experimental results on non-target attacks and target attacks, respectively. Finally, section 6 draws conclusions and points out several future research directions.

2. Universal adversarial perturbations

This section first introduces an iterative algorithm for crafting a UAP for EEG trials, then presents the details of our proposed TLM approach. All source code can be downloaded at https://github.com/ZihanLiu95/UAP_EEG.

We distinguish between two types of attacks:

- *Target attacks*, in which the attacker wants all adversarial examples to be classified into a *specific* class. For example, for 3-class MI (left-hand, right-hand, and feet), the attacker may want all left-hand and right-hand adversarial trials to be misclassified as feet trials.
- *Non-target attacks*, in which the attacker wants the adversarial examples to be misclassified, but does not care which class they are classified into. In the above example, a left-hand adversarial trial could be misclassified as a right-hand trial, or a feet trial.

2.1. Problem setup

To attack a BCI system, the adversarial perturbations need to be added to benign EEG signals in real-time.

Let $X_i \in \mathbb{R}^{C \times T}$ be the i th raw EEG trial ($i = 1, \dots, n$), where C is the number of EEG channels and T the number of time domain samples. Let $\mathbf{x} \in \mathbb{R}^{C \times T \times 1}$ be the vector form of X_i , which concatenates all columns of X_i into a single column. Let $k(\mathbf{x}_i)$ be the estimated label from the target CNN model, $\mathbf{v} \in \mathbb{R}^{C \times T \times 1}$ be the UAP, and $\tilde{\mathbf{x}}_i = \mathbf{x}_i + \mathbf{v}$ be the adversarial EEG trial after adding the UAP. Then, \mathbf{v} needs to satisfy:

$$\left. \begin{aligned} \frac{1}{n} \sum_{i=1}^n I(k(\mathbf{x}_i + \mathbf{v}) \neq k(\mathbf{x}_i)) &\geq \delta \\ \|\mathbf{v}\|_p &\leq \xi \end{aligned} \right\}, \quad (1)$$

where $\|\cdot\|_p$ is the L_p norm, and $I(\cdot)$ is the indicator function which equals 1 if its argument is true, and 0 otherwise. The parameter $\delta \in (0, 1]$ determines the desired attack success rate (ASR), and ξ constrains the

magnitude of \mathbf{v} . Briefly speaking, the first constraint requires the UAP to achieve a desired ASR, and the second constraint ensures the UAP is small.

Next, we describe how a UAP can be crafted for EEG data. We first introduce DeepFool [20], a white-box attack (the attacker has access to all information of the victim model, including its architecture, parameters, and training data) approach for crafting an adversarial perturbation for a *single* input example, and then extend it to crafting a UAP for *multiple* examples. Finally, we propose a novel TLM approach to craft UAP, which can be applied to both non-target attacks and target attacks.

2.2. DeepFool-based UAP

DeepFool is an approach for crafting an adversarial perturbation for a single input example.

Consider a binary classification problem, where the labels are $\{-1, 1\}$. Let \mathbf{x} be an input example, and f an affine classification function $f(\mathbf{x}) = \mathbf{w}^T \mathbf{x} + b$. Then, the predicted label is $k(\mathbf{x}) = \text{sign}f(\mathbf{x})$. The minimal adversarial perturbation \mathbf{r}^* should move \mathbf{x} to the decision hyperplane $\mathcal{F} = \{\mathbf{x}^* : \mathbf{w}^T \mathbf{x}^* + b = 0\}$, i.e.

$$\mathbf{r}^* = -\frac{f(\mathbf{x})}{\|\mathbf{w}\|_2} \mathbf{w}. \quad (2)$$

CNN classifiers are nonlinear. So, an iterative procedure [20] is used to identify the adversarial perturbation, by approximately linearizing $f(\mathbf{x}_t) \approx f(\mathbf{x}_t) + \nabla f(\mathbf{x}_t)^T \mathbf{r}_t$ around \mathbf{x}_t at Iteration t , where ∇ is the gradient of $f(\mathbf{x}_t)$. Then, the minimal perturbation at Iteration t is computed as:

$$\min_{\mathbf{r}_t} \|\mathbf{r}_t\|, \quad \text{s.t. } f(\mathbf{x}_t) + \nabla f(\mathbf{x}_t)^T \mathbf{r}_t = 0. \quad (3)$$

The perturbation \mathbf{r}_t at Iteration t is computed using the closed-form solution in (2), and then $\mathbf{x}_{t+1} = \mathbf{x}_t + \mathbf{r}_t$ is used in the next iteration. The iteration stops when \mathbf{x}_{t+1} starts to change the classification label. The pseudocode is given in algorithm 1.

Algorithm 1 can be extended to multi-class classification by using the one-versus-all scheme to find the closest hyperplane. Experiments in [20] demonstrated that DeepFool can achieve comparable attack performance as FGSM [14], but the magnitude of the perturbation is smaller, which is more desirable.

UAPs were recently discovered in image classification by Moosavi-Dezfooli et al [21], who showed that a fixed adversarial perturbation can fool multiple state-of-the-art CNN classifiers on multiple images. They developed a DeepFool-based iterative algorithm to craft the UAP, which satisfies (1). A UAP is designed by proceeding iteratively over all examples in the dataset $\mathbf{X} = \{\mathbf{x}_i\}_{i=1}^n$. In each iteration, DeepFool is used to compute a minimum perturbation $\Delta \mathbf{v}_i$ for the current perturbed point $\mathbf{x}_i + \mathbf{v}$, and then $\Delta \mathbf{v}_i$ is aggregated into \mathbf{v} .

Algorithm 1. DeepFool [20] for generating an adversarial perturbation for a *single* input example.

Input: \mathbf{x} , an input example
 f , the classification function.
Output: \mathbf{r}^* , the adversarial perturbation.
 $\mathbf{x}_0 = \mathbf{x}$
 $t = 0$
while $\text{sign}f(\mathbf{x}_t) = \text{sign}f(\mathbf{x}_0)$ **do**
 $\mathbf{r}_t = -\frac{f(\mathbf{x}_t)}{\|\nabla f(\mathbf{x}_t)\|_2^2} \nabla f(\mathbf{x}_t)$
 $\mathbf{x}_{t+1} = \mathbf{x}_t + \mathbf{r}_t$
 $t = t + 1$
end
 $\mathbf{r}^* = \sum_{i=0}^t \mathbf{r}_i$.

More specifically, if the current universal perturbation \mathbf{v} cannot fool the classifier on \mathbf{x}_i , then a minimum extra perturbation $\Delta \mathbf{v}_i$ that can fool \mathbf{x}_i is computed by solving the following optimization problem:

$$\min_{\Delta \mathbf{v}_i} \|\Delta \mathbf{v}_i\|_2, \quad \text{s.t. } k(\mathbf{x}_i + \mathbf{v} + \Delta \mathbf{v}_i) \neq k(\mathbf{x}_i). \quad (4)$$

To ensure the constraint $\|\mathbf{v}\|_p \leq \xi$ is satisfied, the updated universal perturbation \mathbf{v} is further projected onto the ℓ_p ball of radius ξ centered at 0. The projection operator $\mathcal{P}_{p,\xi}$ is defined as:

$$\mathcal{P}_{p,\xi}(\mathbf{v}) = \arg \min_{\|\mathbf{v}'\|_p \leq \xi} \|\mathbf{v} - \mathbf{v}'\|_2. \quad (5)$$

Then, the UAP can be updated by $\mathbf{v} = \mathcal{P}_{p,\xi}(\mathbf{v} + \Delta \mathbf{v}_i)$ in each iteration. This process is repeated on the entire dataset until the maximum number of iterations is reached, or the ASR on the perturbed dataset $\mathbf{X}_v = \{\mathbf{x}_i + \mathbf{v}\}_{i=1}^n$ exceeds the target ASR threshold $\delta \in (0, 1]$, i.e.

$$\text{ASR}(\mathbf{X}_v, \mathbf{X}) = \frac{1}{n} \sum_{i=1}^n I(k(\mathbf{x}_i + \mathbf{v}) \neq k(\mathbf{x}_i)) \geq \delta. \quad (6)$$

The pseudo-code of the DeepFool-based algorithm is given in algorithm 2.

2.3. Our proposed TLM-based UAP

Different from the DeepFool-based algorithm, TLM directly optimizes an objective function w.r.t. the UAP by batch gradient descent. In white-box attacks, the parameters of the victim model are known and fixed, and hence we can view the UAP as a variable to minimize an objective function on the entire training set.

Specifically, we solve the following optimization problem by gradient descent:

$$\min_{\mathbf{v}} E_{\mathbf{x} \sim D} l(\mathbf{x} + \mathbf{v}, y) + \alpha \cdot C(\mathbf{x}, \mathbf{v}), \quad \text{s.t. } \|\mathbf{v}\|_p \leq \xi, \quad (7)$$

where $l(\mathbf{x} + \mathbf{v}, y)$ is a loss function, which evaluates the effect of the UAP on the target model. $C(\mathbf{x}, \mathbf{v})$

Algorithm 2. DeepFool-based algorithm for generating a UAP [21].

Input: $\mathbf{X} = \{\mathbf{x}_i\}_{i=1}^n$, n input examples
 k , the classifier
 ξ , the maximum ℓ_p norm of the UAP
 δ , the desired ASR
 M , the maximum number of iterations.
Output: \mathbf{v} , a UAP.
 $\mathbf{v} = \mathbf{0}$
 $\mathbf{X}_v = \mathbf{X}$
for $m = 1, \dots, M$ **do**
 if $\text{ASR}(\mathbf{X}_v, \mathbf{X}) < \delta$ **then**
 for Each $\mathbf{x}_i \in \mathbf{X}$ **do**
 if $k(\mathbf{x}_i + \mathbf{v}) == k(\mathbf{x}_i)$ **then**
 Use DeepFool to compute the minimal perturbation $\Delta \mathbf{v}_i$ in (4);
 Update the perturbation by (5):
 $\mathbf{v} \leftarrow \mathcal{P}_{p,\xi}(\mathbf{v} + \Delta \mathbf{v}_i)$;
 end
 end
 $\mathbf{X}_v = \{\mathbf{x}_i + \mathbf{v}\}_{i=1}^n$;
 else
 Break;
 end
end
Return \mathbf{v} .

is the constraint on the perturbation \mathbf{v} , and α the regularization coefficient. Our proposed approach is highly flexible, as the attacker can choose different optimizers, loss functions, and/or constraints, according to the specific task.

For non-target attacks, the loss function l can be defined as $l(\mathbf{x}, y) = \log(p_y(\mathbf{x}))$, in which y is the true label of \mathbf{x} . This loss function forces the UAP to make the model to have minimum confidence on the true label y . TLM optimizes the perturbation by increasing the expected loss of the model on the training data \mathbf{D} as much as possible. When UAP can affect enough samples in the dataset, and the distribution difference between the test data and \mathbf{D} is small, UAP will have a high probability of affecting the test data. In practice, we could also use $\arg \max_j p_j(\mathbf{x})$, i.e. the predicted label, to replace y if the true label is not available.

For target attacks, the loss function should force the UAP to maximize the model's confidence on the target label y_t specified by attacker; hence, l can be defined as $l(\mathbf{x}, y_t) = -\log(p_{y_t}(\mathbf{x}))$. In fact, both $l(\mathbf{x}, y)$ and $l(\mathbf{x}, y_t)$ are essentially the negative cross-entropy.

There are also various options for the constraint function $C(\mathbf{x}, \mathbf{v})$. In most cases, we can simply set $C(\mathbf{x}, \mathbf{v})$ as L1 or L2 regularization on the UAP \mathbf{v} ; however, it can also be a more sophisticated function, e.g. a metric function to detect whether the input is an adversarial example or not. When a new metric function for detecting adversarial examples is proposed, our approach can also be utilized to test its reliability: we set C as the metric function to check whether we can still find an adversarial example. Given the diversity of metric functions, we only consider L1 or

Algorithm 3. The proposed TLM approach for generating a UAP.

Input: $\mathbf{X}_{train} = \{\mathbf{x}_{train,i}\}_{i=1}^n$, n training examples
 $\mathbf{X}_{val} = \{\mathbf{x}_{val,i}\}_{i=1}^m$, m validation examples
 k , the classifier
 ξ , the maximum ℓ_p norm of the UAP
 α , the regularization coefficient
 δ , the desired ASR
 M , the maximum number of epochs

Output: \mathbf{v}_{best} , a UAP.

$\mathbf{v} = \mathbf{0}$;
 $r = 0$;
for $m = 1, \dots, M$ **do**
 for Each mini-batch $D \in \mathbf{X}_{train}$ **do**
 Update \mathbf{v} in (7) for D with an optimizer
 Constrain \mathbf{v} by (5): $\mathbf{v} \leftarrow \mathcal{P}_{p,\xi}(\mathbf{v})$, or directly clip
 \mathbf{v} into $[-\xi, \xi]$;
 end
 $\mathbf{X}_{val,\mathbf{v}} = \{\mathbf{x}_{val,i} + \mathbf{v}\}_{i=1}^m$;
 if $ASR(\mathbf{X}_{val,\mathbf{v}}, \mathbf{X}_{val}) > r$ **then**
 $r = ASR(\mathbf{X}_{val,\mathbf{v}}, \mathbf{X}_{val})$;
 $\mathbf{v}_{best} = \mathbf{v}$;
 end
 if $r > \delta$ **then**
 Break;
 end
end
Return \mathbf{v}_{best} .

L2 regularization in this paper. Other metric functions and defense strategies for TLM-UAP will be considered in our future research.

To ensure the constraint $\|\mathbf{v}\|_p \leq \xi$, we can project the UAP \mathbf{v} into the l_p ball of radius ξ centered at 0 by the projection function in (5) after each optimization iteration, or simply clip its amplitude into $[-\xi, \xi]$, which is equivalent to using the l_∞ ball. The latter was used in this paper for its simplicity.

Due to the transferability of adversarial examples, i.e. adversarial examples generated by one model may also be used to attack another one, we can perform TLM-UAP attacks in a gray-box attack scenario. In this case, the attacker only has access to the training set of the victim model, instead of its architecture and parameters. The attacker can train a substitute model on the known training set to generate a UAP, which can then be used to attack the victim model.

The TLM-UAP can also be simplified, e.g. the same perturbation is designed for all EEG channels, or a mini TLM-UAP is added to an arbitrary location of an EEG trail. The corresponding experimental results are shown in sections 4.2 and 5. The pseudo-code of our proposed TLM approach is given in algorithm 3.

3. Experimental settings

This section introduces the experimental settings for validating the performance of our proposed TLM approach.

3.1. The three BCI datasets

The following three BCI datasets were used in our experiments, as in [25, 32]:

P300 evoked potentials (P300): The P300 dataset, first introduced in [33], contained eight subjects. In the experiment, each subject faced a laptop on which six images were flashed randomly to elicit P300 responses. The goal was to classify whether the image was a target or non-target. The 32-channel EEG data was downsampled to 256 Hz, bandpass filtered to [1, 40] Hz, and epoched to [0, 1]s after each image onset. Then, we normalized the data using $\frac{x - \text{mean}(x)}{10}$, and clipped the resulting values to $[-5, 5]$. Each subject had about 3300 trials.

Feedback error-related negativity (ERN): The ERN dataset [34] was used in a Kaggle challenge⁵. The EEG signals were collected from 26 subjects and consisted of two classes (bad-feedback and good-feedback). The entire dataset was partitioned into a training set (16 subjects) and a test set (10 subjects). We used all 26 subjects in the experiments. The 56-channel EEG signals were downsampled to 200 Hz, bandpass filtered to [1, 40] Hz, epoched to [0, 1.3] s after each stimulus, and z-normalized. Each subject had 340 trials.

Motor imagery (MI): The MI dataset was Dataset 2A⁶ in BCI Competition IV [35]. The EEG signals were collected from nine subjects and consisted of four classes: the imagined movements of the left hand, right hand, both feet, and tongue. The 22-channel EEG signals were downsampled to 128 Hz, bandpass filtered to [4, 40] Hz, epoched to [0, 2] s after each imagination prompt, and standardized using an exponential moving average window with a decay factor of 0.999, as in [7]. Each subject had 576 trials, with 144 in each class.

3.2. The three CNN models

The following three CNN models were used in our experiments, as in [25, 32]:

EEGNet: EEGNet [7] is a compact CNN architecture for EEG-based BCIs. It consists of two convolutional blocks and a classification block. To reduce the number of model parameters, EEGNet uses depth-wise and separable convolutions [36] instead of traditional convolutions.

DeepCNN: DeepCNN [8] consists of four convolutional blocks and a softmax layer for classification, which is deeper than EEGNet. Its first convolutional block is specially designed to handle EEG inputs, and the other three are standard convolutional blocks.

ShallowCNN: Inspired by filter bank common spatial patterns [37], ShallowCNN [8] is specifically tailored to decode band power features. Compared with DeepCNN, ShallowCNN uses a larger kernel in

⁵ www.kaggle.com/c/inria-bci-challenge.

⁶ www.bbc.de/competition/iv/.

Table 1. Parameters for generating DF-UAP and TLM-UAP. $\|v\|_\infty$ was used in computing the norm of the UAPs.

	ξ	δ	M	α	Constraint
DF-UAP	0.2	0.8	10	—	—
TLM-UAP	0.2	1.0	500	0	No

temporal convolution, and then a spatial filter, squaring nonlinearity, a mean pooling layer and logarithmic activation function.

3.3. The two experimental settings

We considered two experimental settings:

Within-subject experiments: Within-subject five-fold cross-validation was used in the experiments. For each individual subject, all EEG trials were divided into five non-overlapping blocks. Three blocks were selected for training, one for validation, and the remaining one for test. We made sure each block was used in test once, and reported the average results.

Cross-subject experiments: For each dataset, leave-one-subject-out cross-validation was performed. Assume a dataset had N subjects, and the N th subject was selected as the test subject. In training, trials from the first $N - 1$ subjects were mixed, and divided into 75% for training and 25% for validation in early stopping.

When training the victim models on the first two datasets, we applied weights to different classes to accommodate the class imbalance, according to the inverse of its proportion in the training set. We used the cross entropy loss function and Adam optimizer [38]. Early stopping was used to reduce overfitting.

The parameters for generating DF-UAP and TLM-UAP are shown in table 1. It should be noted that TLM-UAP was trained with no constraint, and we set δ to 1.0 and used early stopping (patience = 10) to decide whether to stop the iteration or not. We replaced the true labels y in (7) with the predicted ones, as in real-world applications we do not have the true labels.

3.4. The two performance measures

Both raw classification accuracy (RCA) and balanced classification accuracy (BCA) [3] were used as the performance measures. The RCA is the ratio of the number of correctly classified samples to the number of total samples, and the BCA is the average of the individual RCAs of different classes.

The BCA is necessary, because some BCI paradigms (e.g. P300) have intrinsic significant class imbalance, and hence using RCA alone may be misleading sometimes.

4. Non-target attack results

This section presents the experimental results in non-target attacks on the three BCI datasets. Recall that a non-target attack forces a model to misclassify an adversarial example to any class, instead of a specific class.

For notation convenience, we denote the UAP generated by the DeepFool-based algorithm (algorithm 2) as *DF-UAP*, and the UAP generated by the proposed TLM (algorithm 3) as *TLM-UAP*.

4.1. Baseline performances

We compared the UAP attack performance with two baselines:

4.1.1. Clean baseline

We evaluated the baseline performances of the three CNN models on the clean (unperturbed) EEG data, as shown in the first part of table 2. For all three datasets and all three classifiers, generally RCAs and BCAs of the within-subject experiments were higher than their counterparts in the cross-subject experiments, which is reasonable, since individual differences cause inconsistency among EEG trials from different subjects.

4.1.2. Noisy baseline

We added clipped Gaussian random noise $\xi \cdot \max(-1, \min(1, \mathcal{N}(0, 1)))$ to the original EEG data, which had the same maximum amplitude ξ as the UAP and satisfied the constraint $\|v\|_p \leq \xi$ after clipping. If the random noise under the same magnitude constraint can significantly degrade the classification performance, then there is no need to compute a UAP. The results are shown in table 2. Random noise with the same amplitude as the UAP did not degrade the classification performance in most cases, except sometimes on the MI dataset. This suggests that the three CNN classifiers are generally robust to random noise, therefore we should deliberately design the adversarial perturbations.

4.2. White-box non-target attack performances

First consider white-box attacks, where we have access to all information of the victim model, including its architecture, parameters, and training data. The performances of DF-UAP and TLM-UAP in white-box non-target attacks are shown in the second part of table 2. We also performed non-parametric Mann-Whitney U tests [39] on the RCAs (BCAs) of DF-UAP and TLM-UAP to check if there were statistically significant differences between them (marked with ‘*’; $p < 0.01$). Observe that:

- (a) After adding DF-UAP or TLM-UAP, both the RCAs and the BCAs were significantly reduced,

Table 2. The ratio of the number of correctly classified samples to the number of total samples (RCAs), and the mean RCAs of different classes (BCAs), of the three CNN classifiers in different non-target attack scenarios on the three datasets ($\xi = 0.2$). For each attack type on each CNN model and each dataset, the best performances are marked in bold. Statistically significantly different RCAs/BCAs between DF-UAP and TLM-UAP were marked with ^{**} (non-parametric Mann-Whitney U test; $p < 0.01$).

Experiment	Dataset	Victim model	Baseline			White-Box Attack			Gray-Box Attack					
			Clean	Noisy	DF-UAP	TLM-UAP	Substitute model (DF-UAP)			Substitute model (TLM-UAP)				
							EEGNet	DeepCNN	ShallowCNN	EEGNet	DeepCNN	ShallowCNN	EEGNet	DeepCNN
Within-subject	P300	EEGNet	.81/.79	.80/.79	.18/.51	.17*/.50*	.21/.52	.40/.62	.17*/.20*	.19*/.50*	.40/.62	.17*/.20*	.19*/.50*	.25*/.54*
		DeepCNN	.82/.78	.82/.78	.20/.51	.19*/.50*	.33/.58	.49/.65	.20*/.51*	.20*/.51*	.49/.65	.20*/.51*	.20*/.51*	.30*/.57*
		ShallowCNN	.80/.75	.80/.74	.46/.62	.20*/.51*	.62/.69	.58/.66	.44*/.62*	.43*/.61*	.58/.66	.44*/.62*	.43*/.61*	.29*/.54*
	ERN	EEGNet	.76/.73	.69/.64	.32/.51	.30*/.50	.58/.64	.69/.67	.45*/.57*	.53*/.61	.69/.67	.45*/.57*	.53*/.61	.62*/.66*
		DeepCNN	.69/.65	.69/.65	.44/.55	.40*/.52	.65/.64	.66/.62	.60*/.62	.59*/.61	.66/.62	.60*/.62	.59*/.61	.64/.62
		ShallowCNN	.70/.68	.69/.67	.52/.57	.40*/.53*	.66/.65	.62/.62	.65/.65	.64/.65	.62/.62	.66/.66	.64/.65	.57*/.61*
MI	EEGNet	.61/.61	.47/.48	.29/.29	.25*/.25*	.37/.38	.38/.38	.34/.34	.31*/.31*	.38/.38	.34/.34	.31*/.31*	.34/.34	
	DeepCNN	.50/.50	.47/.47	.35/.35	.26*/.26*	.47/.47	.44/.44	.41/.41	.32*/.32*	.44/.44	.49/.49	.32*/.32*	.39/.39	
	ShallowCNN	.74/.74	.68/.68	.29/.29	.25*/.25*	.55*/.55*	.32/.32	.40/.40	.35*/.35*	.32/.32	.64/.63	.35*/.35*	.27*/.27*	
Cross-subject	P300	EEGNet	.68/.63	.69/.63	.19/.51	.17*/.50*	.22/.52	.27/.54	.17*/.50*	.17*/.50*	.27/.54	.17*/.50*	.18*/.50*	
		DeepCNN	.69/.64	.70/.64	.20/.51	.18*/.50*	.24/.52	.28/.54	.18*/.50*	.18*/.50*	.28/.54	.18*/.50*	.18*/.50*	
		ShallowCNN	.67/.62	.66/.62	.27/.54	.19*/.50*	.32/.55	.32/.55	.20*/.51*	.20*/.51*	.32/.55	.20*/.51*	.18*/.50*	
	ERN	EEGNet	.67/.68	.67/.65	.31/.51	.29*/.50*	.53/.59	.43/.59	.29*/.50*	.31*/.50*	.43/.59	.29*/.50*	.29*/.50*	
		DeepCNN	.69/.69	.71/.70	.34/.52	.31*/.50*	.49/.54	.36/.53	.29*/.50*	.31*/.50*	.36/.53	.29*/.50*	.29*/.50*	
		ShallowCNN	.69/.69	.68/.68	.38/.55	.29*/.50*	.60/.63	.42/.56	.35*/.54*	.38*/.55*	.42/.56	.35*/.54*	.38*/.55*	
MI	EEGNet	.44/.44	.38/.38	.30/.30	.25*/.25*	.30/.30	.30/.30	.29/.29	.29*/.29*	.30/.30	.29/.29	.29*/.29*		
	DeepCNN	.47/.47	.44/.44	.33/.33	.25*/.25*	.38/.38	.28/.28	.32/.32	.30/.30	.28/.28	.40/.40	.26/.26		
	ShallowCNN	.47/.47	.43/.43	.27/.27	.25*/.25*	.31*/.31*	.30/.30	.26/.26	.35/.35	.30/.30	.26/.26	.26*/.26*		

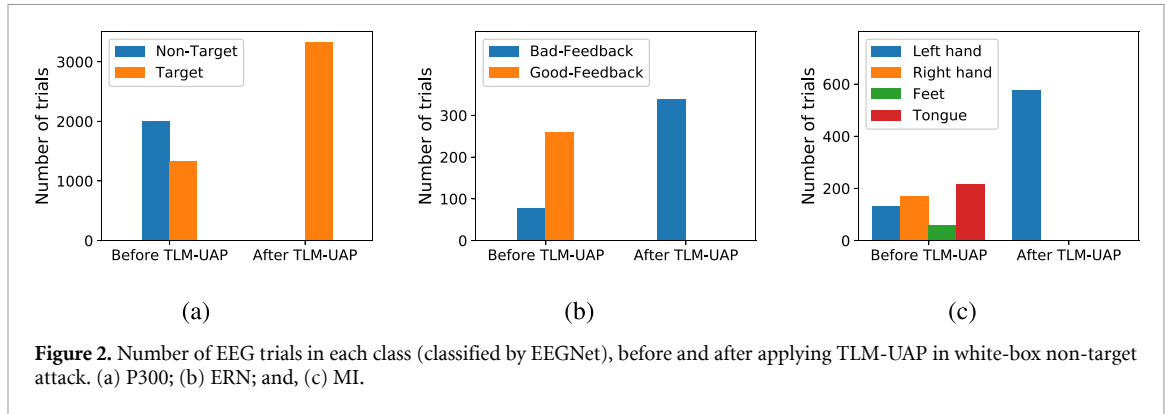


Figure 2. Number of EEG trials in each class (classified by EEGNet), before and after applying TLM-UAP in white-box non-target attack. (a) P300; (b) ERN; and, (c) MI.

suggesting the effectiveness of DF-UAP and TLM-UAP attacks.

- (b) In most cases, TLM-UAP significantly outperformed DF-UAP. This may be due to the fact that, as shown in (7), TLM-UAP optimizes the ASR directly on the entire dataset, whereas DF-UAP considers each sample individually, which may be easily trapped into a local minimum.
- (c) The BCAs of the P300 and ERN datasets were close to 0.5 after DF-UAP or TLM-UAP attacks, whereas the RCAs were lower than 0.5, implying that most test EEG trials were classified into the minority class to achieve the best attack performance.

Figure 2 shows the number of EEG trials in each class on the three datasets, classified by EEGNet before and after applying TLM-UAP. Generally, the trials originally classified into the majority class were misclassified into the minority class in binary classification after applying TLM-UAP. This is reasonable. Assume the minority class contains $p\%$ ($p < 50$) of the trials. Then, misclassifying all minority class trials into the majority class gives an ASR of $p\%$, whereas misclassifying all majority class trials into the minority class gives an ASR of $(100 - p)\%$. Clearly, the latter is larger. Similarly, the trials were misclassified into a minority class (but not necessarily the smallest class) in multi-class classification.

An example of the EEG trial before and after applying Gaussian random noise and TLM-UAP in the time domain is shown in figure 3(a). The signal distortion caused by the Gaussian random noise was greater than that caused by TLM-UAP under the same maximum amplitude, indicating the effectiveness of the constraint on the TLM-UAP amplitude during optimization. Figure 3(b) shows the spectrogram of the EEG trial before and after applying TLM-UAP. The TLM-UAP was so small in both the time domain and the frequency domain that it is barely visible, and hence difficult to detect. We further show the topoplots of an EEG trial before and after white-box non-target TLM-UAP attack in figure 3(c). The

difference is again very small, which may not be detectable by human eyes or a computer program.

We also investigated an easily implementable channel-invariant TLM-UAP attack, which added the same perturbation to all EEG channels. The white-box attack results are shown in table 3. Compared with the clean and noisy baselines in table 2, the RCAs and BCAs in table 3 are smaller, i.e. the attacks were effective; however, the attack performances were worse than those of DF-UAP and TLM-UAP in table 2, which is intuitive, as the channel-invariant TLM-UAP had more constraints.

4.3. Generalization of TLM-UAP on traditional classification models

It's also interesting to evaluate the generalization performance of the proposed TLM-UAP approach on traditional BCI classification models. We used the TLM-UAP generated from CNN models to attack some traditional non-CNN models, i.e. xDAWN [40] spatial filtering and logistic regression (LR) for P300 and ERN, and common spatial pattern [41] filtering and LR classifier for MI.

TLM-UAPs with three different amplitudes were generated by three CNN models, and then used to attack the traditional models. The results are shown in table 4. Observe that:

- (a) Generally, TLM-UAPs generated by CNN models were more effective to degrade the performance of the traditional models than random Gaussian noise, i.e. TLM-UAP can generalize to traditional non-CNN models.
- (b) The amplitude of TLM-UAP heavily affected the attack performance. Intuitively, a larger amplitude resulted in a greater model performance reduction.
- (c) Different models may have different resistance to TLM-UAPs. Compared with the attack performances on CNN models in table 2, the traditional model was more robust than CNN models on the ERN dataset, but less robust on the MI dataset.

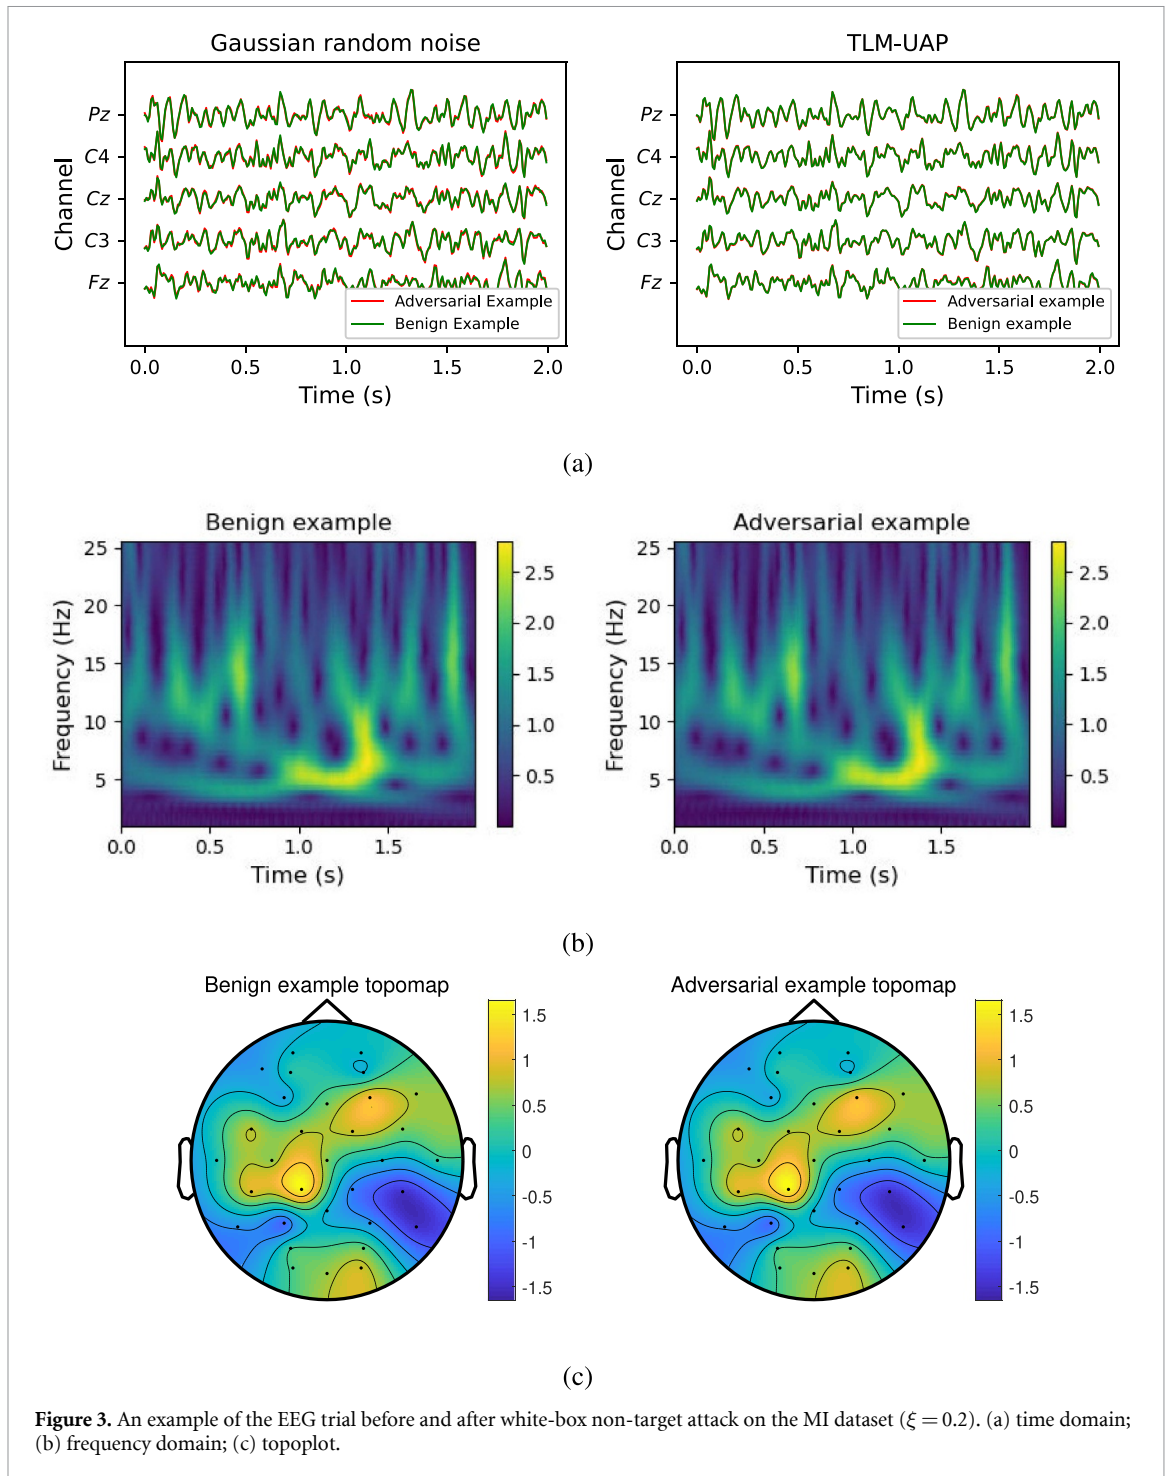


Figure 3. An example of the EEG trial before and after white-box non-target attack on the MI dataset ($\xi = 0.2$). (a) time domain; (b) frequency domain; (c) topoplot.

4.4. Transferability of UAP in gray-box attacks

Transferability is one of the most threatening properties of adversarial examples, which means that adversarial examples generated by one model may also be able to attack another one. This section explores the transferability of DF-UAP and TLM-UAP.

A gray-box attack scenario was considered: the attacker only has access to the training set of the victim model, instead of its architecture and parameters. In this situation, the attacker can train a substitute model on the same training set to generate a UAP,

which was then used to attack the victim model. The results are shown in the last part of table 2. We can observe that:

- (a) The classification performances degraded after gray-box attacks, verifying the transferability of both DF-UAP and TLM-UAP.
- (b) In most cases, TLM-UAP led to larger classification performance degradation of the RCA and BCA than DF-UAP, suggesting that our proposed TLM-UAP was more effective than DF-UAP.

Table 3. The ratio of the number of correctly classified samples to the number of total samples (RCAs) and the average of the individual RCAs of different classes (BCAs) of the three CNN classifiers after channel-invariant TLM-UAP attacks on the three datasets ($\xi = 0.2$).

Experiment	Dataset	Victim model	RCA/BCA			
			Baseline		TLM-UAP	Channel-invariant TLM-UAP
			Clean	Noise		
Within-subject	P300	EEGNet	.81/.79	.80/.79	.17/.50	.47/.66
		DeepCNN	.82/.78	.82/.78	.19/.50	.60/.72
		ShallowCNN	.80/.75	.80/.74	.20/.51	.64/.72
	ERN	EEGNet	.76/.73	.69/.64	.30/.50	.48/.58
		DeepCNN	.69/.65	.69/.65	.40/.52	.53/.59
		ShallowCNN	.70/.68	.69/.67	.40/.53	.56/.61
		EEGNet	.61/.61	.47/.48	.25/.25	.36/.37
		DeepCNN	.50/.50	.47/.47	.26/.26	.52/.51
		ShallowCNN	.74/.74	.68/.68	.25/.25	.74/.74
Cross-subject	P300	EEGNet	.61/.61	.47/.48	.25/.25	.36/.57
		DeepCNN	.69/.64	.70/.64	.18/.50	.40/.58
		ShallowCNN	.67/.62	.66/.62	.19/.50	.49/.61
	ERN	EEGNet	.67/.68	.67/.65	.29/.50	.53/.62
		DeepCNN	.69/.69	.71/.70	.31/.50	.46/.56
		ShallowCNN	.69/.69	.68/.68	.29/.50	.57/.61
	MI	EEGNet	.44/.44	.38/.38	.25/.25	.33/.33
		DeepCNN	.47/.47	.44/.44	.25/.25	.33/.33
		ShallowCNN	.47/.47	.43/.43	.25/.25	.38/.38

Table 4. The ratio of the number of correctly classified samples to the number of total samples (RCAs), and the mean RCAs of different classes (BCAs), of the traditional classifiers after TLM-UAP attacks with different ξ on the three datasets.

Dataset	Victim model	ξ	RCA/BCA				
			Baseline		Generation model		
			Clean	Noise	EEGNet	DeepCNN	ShallowCNN
P300	xDAWN + LR	0.05	.73/.73	.75/.73	.66/.71	.61/.70	.71/.73
		0.1	.73/.73	.75/.73	.56/.68	.48/.65	.66/.71
		0.2	.73/.73	.74/.73	.38/.60	.30/.56	.57/.68
ERN	xDAWN + LR	0.05	.67/.65	.67/.65	.66/.64	.65/.64	.66/.65
		0.1	.67/.65	.66/.65	.64/.65	.64/.64	.65/.65
		0.2	.67/.65	.67/.65	.64/.64	.62/.64	.63/.64
MI	CSP + LR	0.05	.58/.58	.57/.57	.47/.47	.52/.52	.42/.42
		0.1	.58/.58	.42/.42	.31/.31	.32/.32	.28/.28
		0.2	.58/.58	.27/.27	.30/.30	.26/.26	.27/.27

4.5. Characteristics of UAP

Additional experiments were performed in this subsection to analyze the characteristics of TLM-UAP.

4.5.1. Signal-to-perturbation ratio (SPR)

We computed SPRs of the perturbed EEG trials, including applying random noise, DF-UAP and TLM-UAP in white-box attacks. We treated the original EEG trials as clean signals, and computed the SPRs in cross-subject experiments. The results are shown in table 5. We performed non-parametric Mann–Whitney U tests on the SPRs of DF-UAP and TLM-UAP to check if there were statistically significant differences between them (marked in bold; $p < 0.001$). In most cases, the SPRs of the adversarial examples perturbed by TLM-UAP were higher than those perturbed by DF-UAP, i.e. the UAP crafted by TLM had a smaller magnitude, and hence may be more difficult to detect. This

is because in addition to $\|\mathbf{v}\|_p \leq \xi$, TLM-UAP is also bounded by the constraint function $C(\mathbf{x}, \mathbf{v})$ in (7).

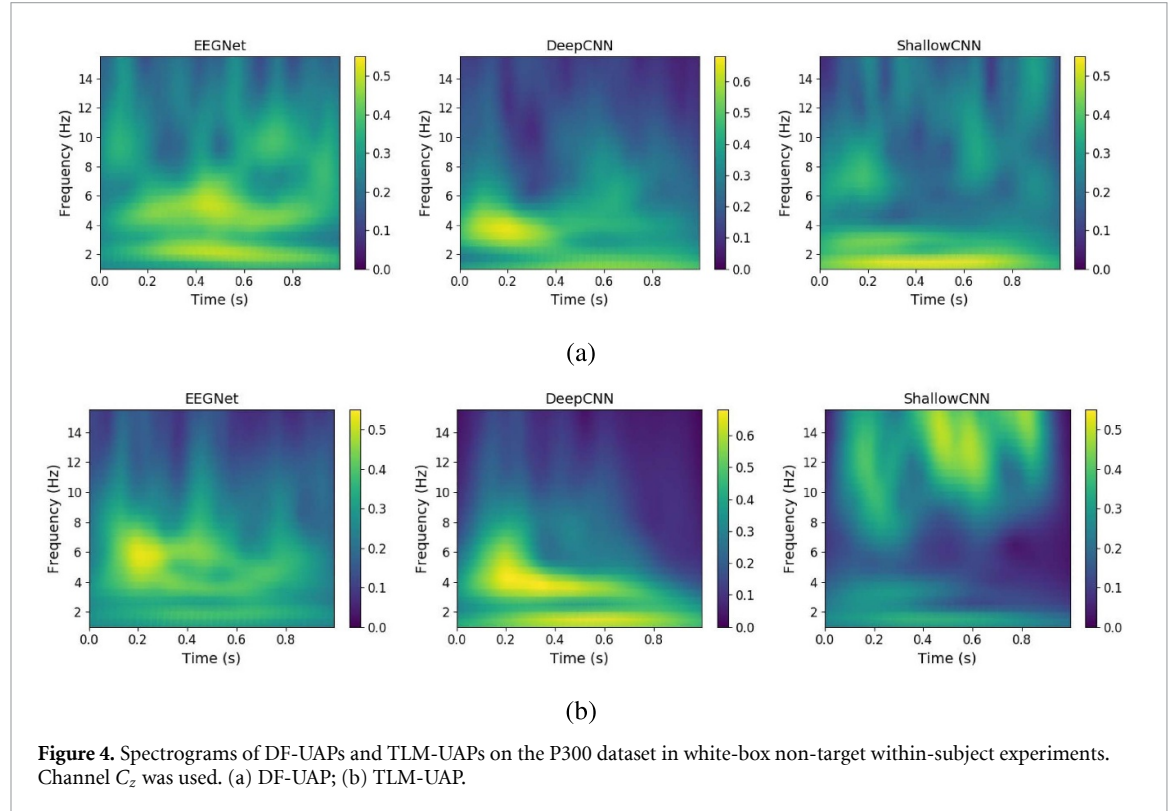
4.5.2. Spectrogram

In order to analyze the time-frequency characteristics of UAP, we compared the spectrograms of DF-UAP and TLM-UAP for the three classifiers in white-box attacks. The results are shown in figure 4.

DF-UAPs and TLM-UAPs share similar spectrogram patterns: for EEGNet and DeepCNN, the energy was mainly concentrated in the low-frequency areas, whereas it was more scattered for ShallowCNN. There were also some significant differences. For EEGNet, the energy of DF-UAP was concentrated in [0.1, 0.9]s and [0, 7]Hz, whereas the energy of TLM-UAP was concentrated in [0.1, 0.8]s and [3, 8] Hz. For DeepCNN, TLM-UAP seemed to affect a longer period of signals, i.e. [0.4, 0.8]s. For ShallowCNN, TLM-UAP

Table 5. SPRs (dB) of EEG trials perturbed by DF-UAP and TLM-UAP in white-box non-target attacks ($\xi = 0.2$). Statistically significantly different SPRs between DF-UAP and TLM-UAP were marked in bold (non-parametric Mann-Whitney U test; $p < 0.001$).

	Dataset	EEGNet	DeepCNN	ShallowCNN
DF-UAP	P300	16.99	17.00	17.85
	ERN	16.22	16.73	17.73
	MI	21.71	13.08	14.57
TLM-UAP	P300	21.17	19.92	20.58
	ERN	21.03	21.67	17.72
	MI	23.48	17.85	17.80



mainly perturbed the high-frequency areas, which were less uniform than DF-UAP.

These results also explained the cross-subject transferability results of TLM-UAP on the P300 dataset in table 2: TLM-UAPs generated from EEGNet and DeepCNN were more similar than those from ShallowCNN, so TLM-UAP generated from EEGNet (DeepCNN) was more effective in attacking DeepCNN (EEGNet), and their RCAs and BCAs were close.

4.5.3. Hidden-layer feature map

Section 4.2 shows that there is no big difference between the EEG trial before and after TLM-UAP attack in the time domain and the frequency domain, and on the topoplot. Figure 5 visualizes the feature map from the last convolution layer of EEGNet before and after adding the TLM-UAP to a clean MI EEG trial. The small perturbation was amplified by the complex nonlinear transformation of EEGNet, and hence the hidden layer feature maps were significantly different. This is intuitive, as otherwise the output of EEGNet would not change much.

4.6. Hyper-parameter sensitivity

This subsection analyzes the sensitivity of TLM-UAP to its hyper-parameters.

4.6.1. The magnitude of TLM-UAP

ξ is an important parameter in algorithm 3, which directly bounds the magnitude of the perturbation. We evaluated the TLM-UAP attack performance with respect to different ξ . As shown in figure 6, the RCA decreased rapidly as ξ increased and converged at $\xi = 0.2$ in most cases, suggesting that a small UAP is powerful enough to attack the victim model.

4.6.2. Training set size

It's interesting to study if the training set size affects the performance of TLM-UAP. Figure 7 shows the white-box non-target attack performance of TLM-UAP, which were trained with different numbers of EEG trials in cross-subject experiments on the MI dataset. It seems that we do not need a large training set to obtain an effective TLM-UAP. The same phenomenon was also observed in [21].

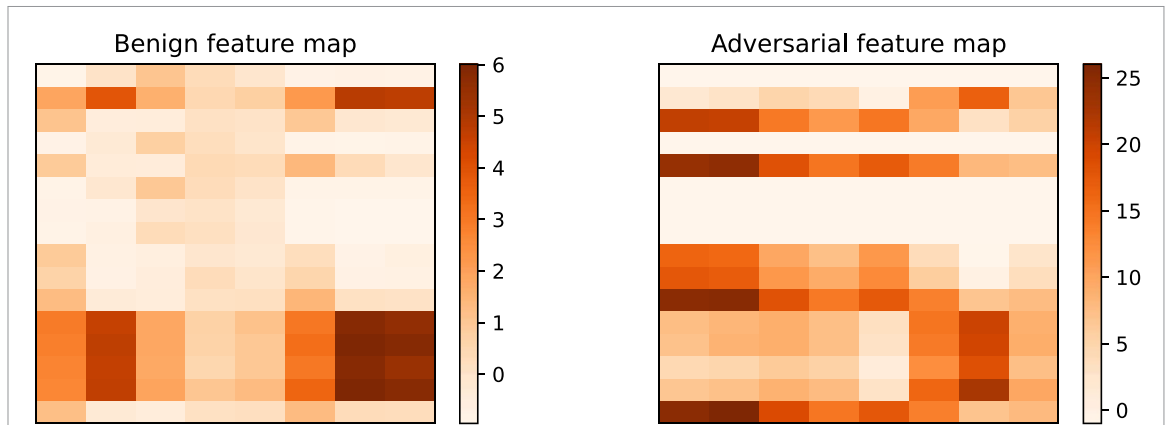


Figure 5. Feature map of EEGNet before and after white-box non-target within-subject TLM-UAP attack on the MI dataset ($\xi = 0.2$).

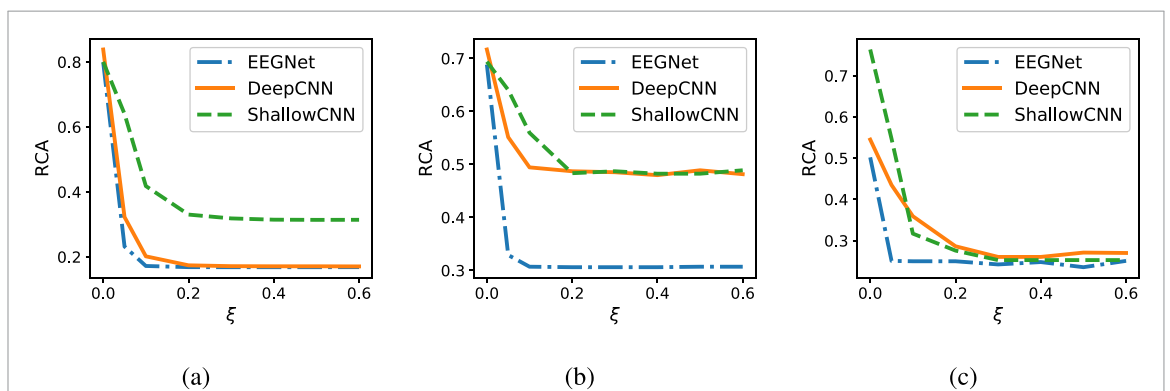


Figure 6. RCAs of the victim model after white-box non-target within-subject TLM-UAP attack, with respect to different ξ . (a) P300 dataset; (b) ERN dataset; and, (c) MI dataset.

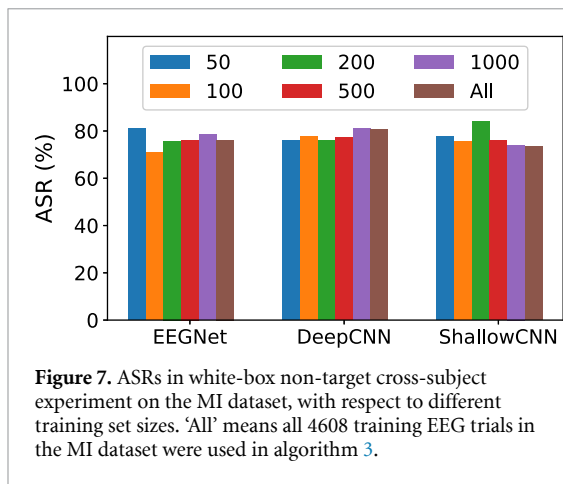


Figure 7. ASRs in white-box non-target cross-subject experiment on the MI dataset, with respect to different training set sizes. ‘All’ means all 4608 training EEG trials in the MI dataset were used in algorithm 3.

4.6.3. Constraint

We also compared different constraint C in (7): No constraint, $L1$ regularization ($\alpha = 10/10/5$ for EEGNet/DeepCNN/ShallowCNN), and $L2$ regularization ($\alpha = 100$). The SPRs on the three datasets are shown in table 6. Albeit similar attack performances, TLM-UAP trained with constraints led to a larger SPR (the SPRs in the ‘L1’ and ‘L2’ rows are larger than those in the corresponding ‘No’ row).

Figure 8 shows that adding different constraints significantly changed the waveforms of TLM-UAP. $L1$ regularization introduced sparsity, whereas $L2$ regularization reduced the perturbation magnitude.

We may also generate a TLM-UAP which satisfies other requirements by changing the constraint function C , such as perturbing certain EEG channels, or even against a metric function which is used to detect adversarial examples. We will leave these to our future research.

4.7. Influence of the Batch size

Mini-batch gradient descent was used to optimize TLM-UAP. Table 7 shows the change of RCAs and BCAs w.r.t. different batch sizes. Generally, the proposed TLM-UAP is insensitive to batch size within [16, 64].

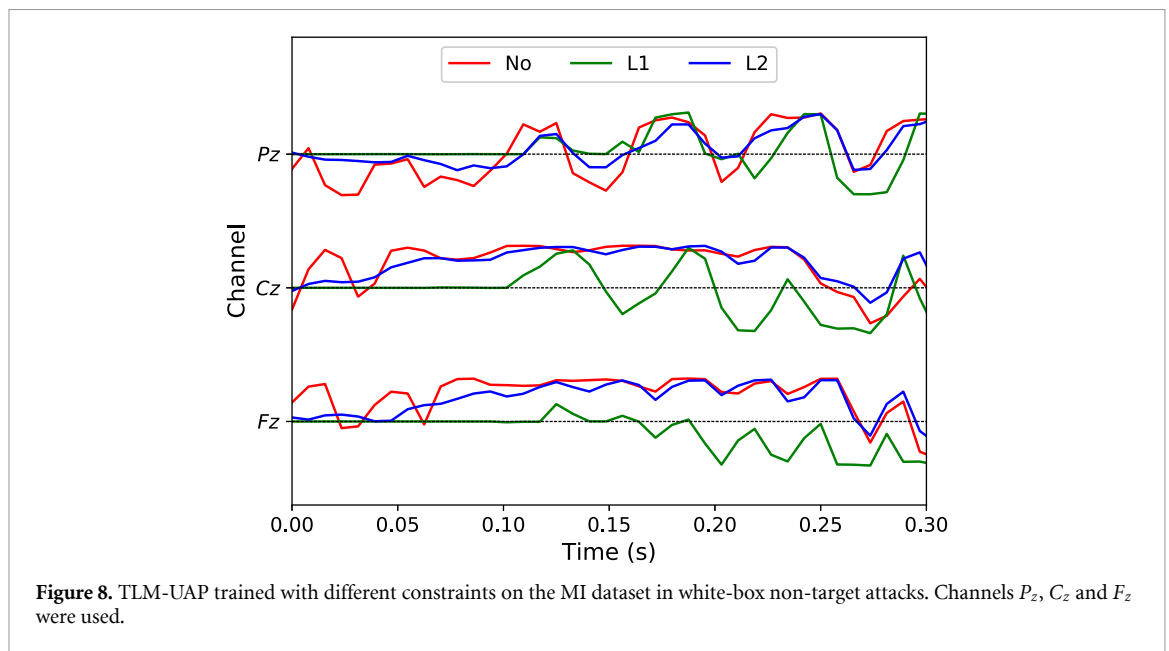
5. Target attack results

Our TLM approach is also capable of performing target attacks, which can be easily achieved by changing the loss function l in (7).

We performed white-box target attacks in cross-subject experiments on the three datasets and

Table 6. Mean RCAs (%) and SPRs (dB) on the three datasets using different constraints in white-box non-target attacks ($\xi = 0.2$).

Dataset	Constraint	Mean RCA	SPR of TLM-UAP		
			EEGNet	DeepCNN	ShallowCNN
P300	No	17.18	14.89	14.71	14.45
	L1	17.36	18.39	17.82	17.16
	L2	17.85	21.17	19.92	20.58
ERN	No	30.96	19.91	20.70	17.02
	L1	29.24	21.45	22.05	17.11
	L2	30.66	21.03	21.67	17.72
MI	No	25.05	22.88	15.46	16.11
	L1	25.08	23.35	53.76	16.88
	L2	25.06	23.48	17.85	17.80

**Table 7.** The ratio of the number of correctly classified samples to the number of total samples (RCAs), and the mean RCAs of different classes (BCAs), of the three CNN classifiers after TLM-UAP attacks with different batch sizes on the three datasets.

Dataset	Victim model	RCA/BCA		
		Batch size 16	Batch size 32	Batch size 64
P300	EEGNet	.17/.50	.17/.50	.17/.50
	DeepCNN	.17/.50	.17/.50	.19/.50
	ShallowCNN	.21/.51	.20/.51	.20/.51
ERN	EEGNet	.31/.50	.31/.50	.30/.50
	DeepCNN	.41/.52	.40/.52	.40/.52
	ShallowCNN	.38/.51	.38/.51	.40/.53
MI	EEGNet	.25/.25	.25/.25	.25/.25
	DeepCNN	.29/.29	.29/.29	.26/.26
	ShallowCNN	.26/.26	.26/.26	.25/.25

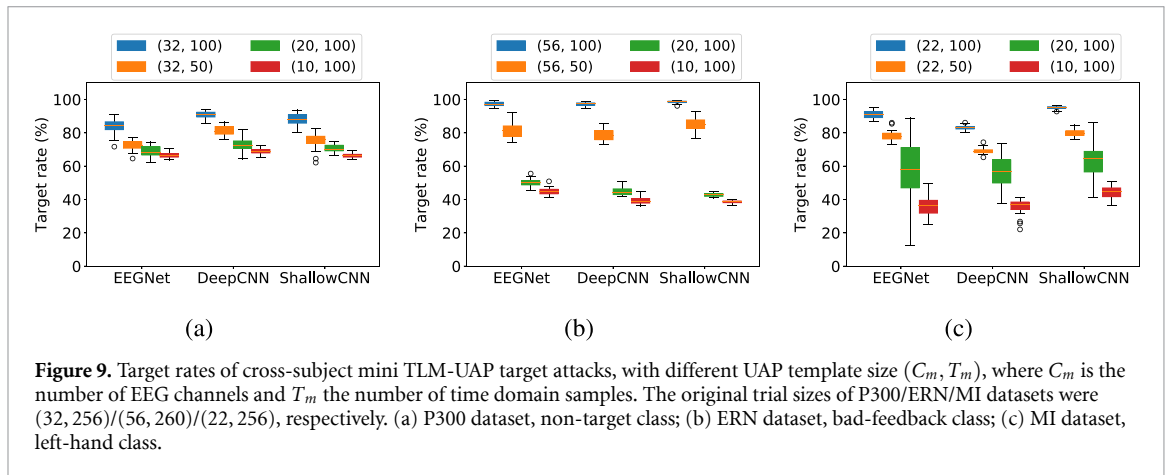
evaluated the *target rate*, which is the number of samples classified to the target class divided by the number of total samples. The results are shown in table 8. TLM-UAPs had close to 100% target rates in white-box target attacks, indicating that our approach can manipulate the BCI systems to output whatever the attacker wants, which may be more dangerous than non-target attacks. For example, in a BCI-driven wheelchair, a target TLM-UAP attack may force all

commands to be interpreted as a specific command (e.g. going forward), and hence run the user into danger.

To further simplify the implementation of TLM-UAP, we also considered smaller template size, i.e. mini TLM-UAP with a small number of channels and time domain samples, which can be added anywhere to an EEG trail. Mini TLM-UAPs are more practical and flexible, because they do not require the

Table 8. Target rates of the three CNN classifiers in cross-subject white-box target TLM-UAP attacks on the three datasets ($\xi = 0.2$).

Dataset	Victim model	Target class	Baseline		TLM-UAP
			Clean	Noisy	Attack
P300	EEGNet	Non-target	.6627	.6463	.9629
		Target	.3373	.3510	.9572
	DeepCNN	Non-target	.6755	.6637	.9416
		Target	.3245	.3116	.9373
	ShallowCNN	Non-target	.6505	.6597	.8904
		Target	.3495	.3499	.8306
ERN	EEGNet	Bad	.3537	.3741	.9980
		Good	.6463	.6300	.9971
	DeepCNN	Bad	.3770	.3309	.9912
		Good	.6230	.6739	.9976
	ShallowCNN	Bad	.3033	.2910	.9741
		Good	.6967	.7160	.9888
MI	EEGNet	Left	.3152	.1350	.9821
		Right	.2830	.2056	.9850
		Feet	.1545	.1954	.9994
		Tongue	.2473	.5380	1.000
	DeepCNN	Left	.2535	.1765	.8839
		Right	.3491	.2207	.9238
		Feet	.2282	.3155	.9659
		Tongue	.1692	.2544	.9938
	ShallowCNN	Left	.2872	.1952	.9151
		Right	.2537	.1746	.9443
		Feet	.2647	.2838	.9819
		Tongue	.2124	.3673	.9983



attacker to know the exact number of EEG channels and the exact length and starting time of an EEG trial. During optimization, we randomly placed the mini TLM-UAP at different locations (both channel-wise and time-wise) of EEG trials and tried to make the attacks successful. During test, the mini TLM-UAP was randomly added to 30 different locations of each EEG trail. The results are shown in figure 9. Generally, all mini TLM-UAPs were effective. However, their effectiveness decreased when the number of used channels (C_m) and/or the template length (T_m) decreased, which is intuitive. These results suggest that a mini TLM-UAP may be used to achieve a better compromise between the attack performance and the implementation difficulty.

6. Conclusions and future research

Multiple CNN classifiers have been proposed for EEG-based BCIs. However, CNN models are vulnerable to UAPs, which are small and example-independent perturbations, yet powerful enough to significantly degrade the performance of a CNN model when added to a benign example. This paper has proposed a novel TLM approach to generate UAP for EEG-based BCI systems. Experimental results demonstrated its effectiveness in attacking three popular CNN classifiers for both non-target and target attacks. We also verified the transferability of the UAPs in EEG-based BCI systems. To our knowledge, this is the first study on UAPs of CNN

classifiers in EEG-based BCIs. It exposes a potentially critical security problem in BCIs, and hopefully will lead to the design of safer BCIs.

Our future research will enhance the transferability of TLM-UAP in deep learning, and also consider how to attack traditional machine learning models in EEG-based BCIs. More importantly, we will design effective strategies to defend against UAP attacks. Multiple approaches, e.g. adversarial training [14], defensive distillation [42], ensemble adversarial training [43], and so on [23, 44–46], have been proposed to defend against adversarial examples in other application domains. As TLM is a perturbation based first-order gradient optimization approach, PGD [23] training may be used to defend against it.

Data availability statement

No new data were created or analysed in this study.

Acknowledgment

This research was supported by the Technology Innovation Project of Hubei Province of China (2019AEA171), Zhejiang Lab (2021KE0AB04), the Hubei Province Funds for Distinguished Young Scholars (2020CFA050), the National Natural Science Foundation of China (61873321 and U1913207), the National Social Science Foundation of China (19ZDA104 and 20AZD089), and the International Science and Technology Cooperation Program of China (2017YFE0128300).

ORCID iD

Dongrui Wu  <https://orcid.org/0000-0002-7153-9703>

References

- [1] Sutton S, Braren M, Zubin J and John E R 1965 Evoked-potential correlates of stimulus uncertainty *Science* **150** 1187–8
- [2] Farwell L A and Donchin E 1988 Talking off the top of your head: toward a mental prosthesis utilizing event-related brain potentials *Electroencephalogr. Clin. Neurophysiol.* **70** 510–23
- [3] Dongrui W 2017 Online and offline domain adaptation for reducing BCI calibration effort *IEEE Trans. Human-Machine Syst.* **47** 550–63
- [4] Dongrui W, Lawhern V J, Hairston W D and Lance B J 2016 Switching EEG headsets made easy: reducing offline calibration effort using active weighted adaptation regularization *IEEE Trans. Neural Syst. Rehabil. Eng.* **24** 1125–37
- [5] Pfurtscheller G and Neuper C 2001 Motor imagery and direct brain-computer communication *Proc. IEEE* **89** 1123–34
- [6] Zhu D, Bieger J, Molina G G and Aarts R M 2010 A survey of stimulation methods used in SSVEP-based BCIs *Computat. Intell. Neurosci.* **702357**
- [7] Lawhern V J, Solon A J, Waytowich N R, Gordon S M, Hung C P and Lance B J 2018 EEGNet: a compact convolutional neural network for EEG-based brain-computer interfaces *J. Neural Eng.* **15** 056013
- [8] Schirrmeister R T, Springenberg J T, Fiederer L D J, Glasstetter M, Eggensperger K, Tangermann M, Hutter F, Burgard W and Ball T 2017 Deep learning with convolutional neural networks for EEG decoding and visualization *Human Brain Mapp.* **38** 5391–420
- [9] Bashivan P, Rish I, Yeasin M and Codella N 2016 Learning representations from EEG with deep recurrent-convolutional neural networks *Proc. Int' Conf. on Learning Representations (San Juan, Puerto Rico)*
- [10] Tabar Y R and Halici U 2017 A novel deep learning approach for classification of EEG motor imagery signals *J. Neural Eng.* **14** 016003
- [11] Tayeb Z, Fedjaev J, Ghaboosi N, Richter C, Everding L, Xingwei Q, Yingyu W, Cheng G and Conradt Jorg 2019 Validating deep neural networks for online decoding of motor imagery movements from EEG signals *Sensors* **19** 210
- [12] Szegedy C, Zaremba W, Sutskever I, Bruna J, Erhan D, Goodfellow I J and Fergus R 2014 Intriguing properties of neural networks *Proc. Int' Conf. on Learning Representations (Banff, Canada)*
- [13] Biggio B, Corona I, Maiorca D, Nelson B, Šrđić N, Laskov P, Giacinto G and Roli F 2013 Evasion attacks against machine learning at test time *Proc. Joint Conf. on Machine Learning and Knowledge Discovery in Databases (Berlin, Germany)* pp 387–402
- [14] Goodfellow I J, Shlens J and Szegedy C 2015 Explaining and harnessing adversarial examples *Proc. Int' Conf. on Learning Representations (San Diego, CA)*
- [15] Kurakin A, Goodfellow I J and Bengio S 2017 Adversarial examples in the physical world *Proc. Int' Conf. on Learning Representations (Toulon, France)*
- [16] Brown T B, Mané D, Roy A, Abadi M and Gilmer J 2017 Adversarial patch *CoRR* [abs/1712.09665](https://arxiv.org/abs/1712.09665)
- [17] Athalye A, Engstrom L, Ilyas A and Kwok K 2018 Synthesizing robust adversarial examples *Proc. 35th Int' Conf. on Machine Learning (Stockholm, Sweden)* pp 284–93
- [18] Carlini N and Wagner D A 2018 Audio adversarial examples: targeted attacks on speech-to-text *Proc. Symp. on Security and Privacy (San Francisco, CA)* pp 1–7
- [19] Grosse K, Papernot N, Manoharan P, Backes M and Patrick M 2016 Adversarial perturbations against deep neural networks for malware classification *CoRR* [abs/1606.04435](https://arxiv.org/abs/1606.04435)
- [20] Moosavi-Dezfooli S-M, Fawzi A and Frossard P 2016 DeepFool: a simple and accurate method to fool deep neural networks *Proc. Conf. on Computer Vision and Pattern Recognition (Las Vegas, NV)* pp 2574–82
- [21] Moosavi-Dezfooli S-M, Fawzi A, Fawzi O and Frossard P 2017 Universal adversarial perturbations *Proc. Conf. on Computer Vision and Pattern Recognition (Honolulu, HI)* pp 1765–73
- [22] Baluja S and Fischer I 2017 Adversarial transformation networks: learning to generate adversarial examples *CoRR* [abs/1703.09387](https://arxiv.org/abs/1703.09387)
- [23] Madry A, Makelov A, Schmidt L, Tsipras D and Vladu A 2018 Towards deep learning models resistant to adversarial attacks *Proc. Int'l. Conf. on Learning Representations (Vancouver, Canada)*
- [24] Carlini N and Wagner D 2017 Towards evaluating the robustness of neural networks *Proc. Symp. on Security and Privacy (San Jose, CA: IEEE)* pp 39–57
- [25] Zhang X and Wu D 2019 On the vulnerability of CNN classifiers in EEG-based BCIs *IEEE Trans. Neural Syst. Rehabil. Eng.* **27** 814–25
- [26] Jundong Li, Cheng K, Wang S, Fred Morstatter R P Trevino J T and Liu H 2016 Feature selection: a data perspective *CoRR* [abs/1601.07996](https://arxiv.org/abs/1601.07996)
- [27] Zhang X, Dongrui W, Ding L, Luo H, Lin C-T, Jung T-P and Chavarriaga R 2021 Tiny noise, big mistakes: adversarial perturbations induce errors in brain-computer interface spellers *Natl Sci. Rev.* **8**

- [28] Neekhara P, Hussain S, Pandey P, Dubnov S, McAuley J J and Koushanfar F 2019 Universal adversarial perturbations for speech recognition systems *CoRR* [abs/1905.03828](https://arxiv.org/abs/1905.03828)
- [29] Behjati M, Moosavi-Dezfooli S-M, Baghshah M S and Frossard P 2019 Universal adversarial attacks on text classifiers *Proc. IEEE Int'l Conf. on Acoustics, Speech and Signal Processing (Brighton, United Kingdom)* pp 7345–9
- [30] Mopuri K R, Ganeshan A and Radhakrishnan V B 2019 Generalizable data-free objective for crafting universal adversarial perturbations *IEEE Trans. Pattern Anal. Mach. Intell.* **41** 2452–65
- [31] Hirano H and Takemoto K 2020 Simple iterative method for generating targeted universal adversarial perturbations *Proc. of 25th Int'l. Symp. on Artificial Life and Robotics (Beppu, Japan)* pp 426–30
- [32] Jiang X, Zhang X and Dongrui W 2019 Active learning for black-box adversarial attacks in EEG-based brain-computer interfaces *Proc. Symp. Series on Computational Intelligence (Xiamen, China)*
- [33] Hoffmann U, Vesin J-M, Ebrahimi T and Diserens K 2008 An efficient P300-based brain-computer interface for disabled subjects *J. Neurosci. Methods* **167** 115–25
- [34] Margaux P, Emmanuel M, Sébastien D, Olivier B and Mattout Jémie 2012 Objective and subjective evaluation of online error correction during P300-based spelling *Adv. Human-Computer Interact.* **2012** 13
- [35] Tangermann M et al 2012 Review of the BCI competition IV *Front. Neurosci.* **6** 55
- [36] Chollet F 2017 Xception: deep learning with depthwise separable convolutions *Proc. Conf. on Computer Vision and Pattern Recognition (Honolulu, HI)* pp 1800–7
- [37] Ang K K, Chin Z Y, Zhang H and Guan C 2008 Filter bank common spatial pattern (FBCSP) in brain-computer interface *Proc. IEEE Int'l Conf. on Neural Networks (Hong Kong, China)*
- [38] Kingma D P and Jimmy B 2014 Adam: a method for stochastic optimization *CoRR* [abs/1412.6980](https://arxiv.org/abs/1412.6980)
- [39] Mann H B and Whitney D R 1947 On a test of whether one of two random variables is stochastically larger than the other *Ann. Math. Stat.* **18** 50–60
- [40] Rivet B, Souloumiac A, Attina V and Gibert G 2009 xDAWN algorithm to enhance evoked potentials: application to brain-computer interface *IEEE Trans. Biomed. Eng.* **56** 2035–43
- [41] Ramoser H, Muller-Gerking J and Pfurtscheller G 2000 Optimal spatial filtering of single trial EEG during imagined hand movement *IEEE Trans. Rehabil. Eng.* **8** 441–6
- [42] Papernot N, McDaniel P, Wu Xi, Jha S and Swami A 2016 Distillation as a defense to adversarial perturbations against deep neural networks *Proc. Symp. on Security and Privacy (San Jose, CA: IEEE)* pp 582–97
- [43] Tramèr F, Kurakin A, Papernot N, Goodfellow I, Boneh D and Patrick M 2018 Ensemble adversarial training: attacks and defenses *Proc. Int'l Conf. on Learning Representations (Vancouver, Canada)*
- [44] Xie C, Wang J, Zhang Z, Ren Z and Yuille A 2018 Mitigating adversarial effects through randomization *Proc. Int'l Conf. on Learning Representations (Vancouver, Canada)*
- [45] Guo C, Rana M, Cissé M and van der Maaten L 2017 Countering adversarial images using input transformations *CoRR* [abs/1711.00117](https://arxiv.org/abs/1711.00117)
- [46] Akhtar N, Liu J and Mian A 2018 Defense against universal adversarial perturbations *Proc. Conf. on Computer Vision and Pattern Recognition (Salt Lake City, UT)* pp 3389–98