



White-Box Target Attack for EEG-Based BCI Regression Problems

Lubin Meng¹, Chin-Teng Lin², Tzyy-Ping Jung³, and Dongrui Wu¹(✉)

¹ School of Artificial Intelligence and Automation,
Huazhong University of Science and Technology, Wuhan, Hubei, China
drwu@hust.edu.cn

² Centre of Artificial Intelligence, Faculty of Engineering and Information
Technology, University of Technology, Sydney, Australia

³ Swartz Center for Computational Neuroscience, Institute for Neural Computation,
University of California San Diego (UCSD), La Jolla, CA, USA

Abstract. Machine learning has achieved great success in many applications, including electroencephalogram (EEG) based brain-computer interfaces (BCIs). Unfortunately, many machine learning models are vulnerable to adversarial examples, which are crafted by adding deliberately designed perturbations to the original inputs. Many adversarial attack approaches for classification problems have been proposed, but few have considered target adversarial attacks for regression problems. This paper proposes two such approaches. More specifically, we consider white-box target attacks for regression problems, where we know all information about the regression model to be attacked, and want to design small perturbations to change the regression output by a pre-determined amount. Experiments on two BCI regression problems verified that both approaches are effective. Moreover, adversarial examples generated from both approaches are also transferable, which means that we can use adversarial examples generated from one known regression model to attack an unknown regression model, i.e., to perform black-box attacks. To our knowledge, this is the first study on adversarial attacks for EEG-based BCI regression problems, which calls for more attention on the security of BCI systems.

Keywords: Adversarial attack · Brain-computer interfaces · Regression · Target attack · White-box attack

1 Introduction

Machine learning has been widely used to solve many difficult tasks. One of them is brain-computer interfaces (BCIs). BCIs enable a user to directly communicate with a computer via brain signals [18], and have attracted lots of research interest recently [10, 11]. Electroencephalogram (EEG) is the most frequently used input signal in BCIs, because of its low-cost and non-invasive nature. Three commonly used BCI paradigms are motor imagery (MI) [15], event-related potentials

(ERP) [16, 20], and steady-state visual evoked potentials (SSVEP) [12]. Machine learning can be used to extract more generalizable features [23] and construct more accurate models [22], and hence makes BCIs more robust and user-friendly.

Recent research has shown that many machine learning models are vulnerable to adversarial examples. By adding deliberately designed perturbations to legitimate data, adversarial examples can cause large changes in the model outputs. The perturbations are usually so small that they are hardly noticeable by a human or a computer program, but can dramatically degrade the model performance. For example, in image recognition, adversarial examples can easily mislead a classifier to give a wrong output [6]. In speech recognition, adversarial examples can generate audio that sounds meaningless to a human, but be understood as a meaningful voice command by a smart phone [2]. Our recent work [24] also showed that adversarial examples can dramatically degrade the classification accuracy of EEG-based BCIs.

There are many different approaches for crafting adversarial examples. Szegedy *et al.* [17] first discovered the existence of adversarial examples in 2014, and proposed an optimization-based approach, L-BFGS, to find them. Goodfellow *et al.* [6] proposed a fast gradient sign method (FGSM) in 2014, which can rapidly find adversarial examples by searching for perturbations in the direction the loss has the fastest change. Carlini and Wagner [3] proposed the CW method in 2017, which can find adversarial examples with very small distortions.

All above approaches focused on classification problems, which find perturbations that can push the original examples cross the decision boundary. Jagielski *et al.* [7] conducted the first *non-target* adversarial attacks for linear regression models. This paper considers *target* adversarial attacks for regression problems, which change the model output by a pre-determined amount. Our contributions are:

1. We propose two approaches, based on optimization and gradient, respectively, to perform white-box target attack for regression problems.
2. We validate the effectiveness of our proposed approaches in two EEG-based BCI regression problems (drowsiness estimation and reaction time estimation). They can craft adversarial EEG trials that a human cannot distinguish from the original EEG trials, but can dramatically change the outputs of the BCI regression model.
3. We show that adversarial examples crafted by our approaches are transferable: adversarial examples crafted from a ridge regression model can also successfully attack a neural network model, and vice versa. This makes black-box attacks possible.

The attacks proposed in this paper may pose serious security and safety problems in real-world BCI applications. For example, an EEG-based BCI system may be used to monitor the driver’s drowsiness level and urge him/her to take breaks accordingly. An attack that deliberately changes the estimated drowsiness level from a high value to a low value may overload the driver, and hence cause accidents.

The remainder of this paper is organized as follows: Sect. 2 introduces several typical adversarial attack approaches for classification problems. Section 3 pro-

poses two white-box target attack approaches for regression problems. Section 4 evaluates the performances of our proposed approaches in two EEG-based BCI regression problems. Section 5 draws conclusion.

2 Adversarial Attacks for Classification Problems

This section introduces two typical adversarial attack approaches for classification problems, which are extended to regression problems in the next section.

2.1 Adversarial Attack Types

Assume a valid benign example $\mathbf{x} \in [0, 1]^k$ (k is the dimensionality of \mathbf{x}) is classified into Class y by a classifier $f(\mathbf{x})$. It is possible to find an adversarial example $\mathbf{x}' \in [0, 1]^k$, which is very similar to the original sample \mathbf{x} according to some distance metric d , but is misclassified to $f(\mathbf{x}') \neq y$. According to how $f(\mathbf{x}')$ is different from y , there can be two types of attacks:

1. *Target attack*, in which all adversarial examples are classified into a pre-determined class $y' \neq y$.
2. *Non-target attack*, whose goal is to construct adversarial examples that will be misclassified, but does not require them to be misclassified into a particular class.

According to how much knowledge the attacker can obtain about the target model (the model to be attacked), adversarial attacks can also be categorized into:

1. *White-box attack*, in which the attacker knows all information about the target model, such as its architecture and all parameter values.
2. *Black-box attack*, in which the attacker does not know the architecture and parameters of the target model; instead, he/she can feed some inputs to it and observe its outputs. In this way, he/she can obtain some training examples, and train a substitute model to craft adversarial examples to attack the target model. This approach makes use of the transferability of the adversarial examples [13].

2.2 White-Box Target Attack Approaches

This paper considers white-box target attacks only. Assume we know the architecture and all parameters of the classifier $f(\mathbf{x})$. We want to craft an adversarial example \mathbf{x}' from an input \mathbf{x} so that $f(\mathbf{x}') = y_t$, where y_t is a fixed class for all \mathbf{x}' .

Two representative target attack approaches for classification problems are:

1. *Carlini and Wagner (CW)* [3], which improves L-BFGS [17]. It introduces a new variable ω so that

$$\delta = \frac{1}{2}(\tanh(\omega) + 1) - \mathbf{x} \quad (1)$$

automatically satisfies the constraint $\mathbf{x}' = \mathbf{x} + \boldsymbol{\delta} = \frac{1}{2}(\tanh(\boldsymbol{\omega}) + 1) \in [0, 1]^k$. $\boldsymbol{\omega}$ in (1) is the variable to be optimized, which can assume any value in $(-\infty, \infty)$. Given \mathbf{x} , \mathbf{w} is found through:

$$\min \left\| \frac{1}{2}(\tanh(\boldsymbol{\omega}) + 1) - \mathbf{x} \right\|_2 + c \cdot \ell \left(\frac{1}{2}(\tanh(\boldsymbol{\omega}) + 1) \right) = \min \|\boldsymbol{\delta}\|_2 + c \cdot \ell(\mathbf{x}') \quad (2)$$

where c is a trade-off parameter, and

$$\ell(\mathbf{x}') = \max \left(\max_{i \neq y_t} Z(\mathbf{x}')_i - Z(\mathbf{x}')_{y_t}, -\lambda \right), \quad (3)$$

in which $Z(\mathbf{x}')_i$ is the logits of the target model in Class i , and λ controls the confidence of the adversarial example. A large λ forces the adversarial example to be classified into the target class y_t with high confidence.

2. *Iterative Target Class Method (ITCM)*¹ [9], which modifies FGSM [6], an efficient approach for non-target attacks:

$$\mathbf{x}' = \mathbf{x} + \epsilon \cdot \text{sign}(\nabla_{\mathbf{x}} J(\mathbf{x}, y_{true})), \quad (4)$$

where ϵ controls the amplitude of the perturbation, J is a loss function, and y_{true} is the true label of \mathbf{x} .

ITCM performs target attack by replacing y_{true} in (4) by the target class y_t . It also improves the attack performance by taking multiple small steps of α in the gradient direction and clipping the maximum perturbation to ϵ , instead of taking a single large step of ϵ in (4):

$$\begin{aligned} \mathbf{x}'_0 &= \mathbf{x}, \\ \mathbf{x}'_{m+1} &= \text{Clip}_{\mathbf{x}, \epsilon} \{ \mathbf{x}'_m - \alpha \cdot \text{sign}(\nabla_{\mathbf{x}'_m} J(\mathbf{x}'_m, y_t)) \}, \end{aligned} \quad (5)$$

where $\text{Clip}_{\mathbf{x}, \epsilon}(\mathbf{x}')$ ensures the difference between each dimension of \mathbf{x}' and the corresponding dimension of \mathbf{x} does not exceed ϵ .

3 White-Box Target Attack for Regression Problems

The section proposes two white-box target attack approaches for regression problems.

Let \mathbf{x} be an input, y the groundtruth output, and $g(\mathbf{x})$ the regression model. Target attack aims to generate a small perturbation $\boldsymbol{\delta}$ such that the adversarial example $\mathbf{x}' = \mathbf{x} + \boldsymbol{\delta}$ can change the regression output to $g(\mathbf{x}') \geq y + t$, where $t > 0$ is a predefined target²:

$$\min_{\mathbf{x}'} \|\mathbf{x}' - \mathbf{x}\|_2, \quad \text{s.t.} \quad g(\mathbf{x}') - y \geq t \quad (7)$$

¹ It is also called the iterative least-likely class method in [9].

² The regression output can also be changed to $g(\mathbf{x}') \leq y - t$. Without loss of generality, $g(\mathbf{x}') \geq y + t$ is considered in this paper.

3.1 CW for Regression (CW-R)

To extend the CW target attack approach from classification to regression, we optimize the following loss function:

$$\min_{\omega} \left\| \frac{1}{2}(\tanh(\omega) + 1) - \mathbf{x} \right\|_2 + c \cdot \ell(\mathbf{x}, \omega, t), \quad (8)$$

where

$$\ell(\mathbf{x}, \omega, t) = \max \left\{ g(\mathbf{x}) + t - g \left(\mathbf{x} + \frac{1}{2}(\tanh(\omega) + 1) \right), 0 \right\} \quad (9)$$

$$= \max \{ g(\mathbf{x}) + t - g(\mathbf{x}'), 0 \}. \quad (10)$$

The constructed adversarial example is then $\mathbf{x}' = \mathbf{x} + \frac{1}{2}(\tanh(\omega) + 1)$.

The pseudocode of the proposed CW method for regression (CW-R) is shown in Algorithm 1. It uses iterative binary search to find the optimal trade-off parameter c .

Algorithm 1. CW for regression (CW-R).

Input: \mathbf{x} , the original example;
 $g(\mathbf{x})$, the target regression model;
 t , the minimum change of the output;
 M , the number of the iterations;
 c_0 , initialization of the trade-off parameter;
 N , the number of binary search steps for the optimal trade-off parameter.

Output: \mathbf{x}' , the adversarial example.

Initialize $c \leftarrow c_0$, $\omega_1 \leftarrow \text{random}$, $d_{\min} \leftarrow \infty$, $\bar{c} \leftarrow 1e4$, $\underline{c} \leftarrow 0$

for $n = 1 : N$ **do**

for $m = 1 : M$ **do**

$\mathbf{x}'_m \leftarrow \frac{1}{2}(\tanh(\omega_m) + 1)$;

$\ell \leftarrow \|\mathbf{x}'_m - \mathbf{x}\|_2 + c \cdot \max\{g(\mathbf{x}) + t - g(\mathbf{x}'_m), 0\}$;

$\omega_{m+1} \leftarrow \omega_m - \alpha \cdot \frac{\partial \ell}{\partial \omega_m}$;

if $g(\mathbf{x}'_m) \geq g(\mathbf{x}) + t$ **and** $\|\mathbf{x}'_m - \mathbf{x}\|_2 \leq d_{\min}$ **then**

$\mathbf{x}' \leftarrow \mathbf{x}'_m$;

$d_{\min} \leftarrow \|\mathbf{x}'_m - \mathbf{x}\|_2$;

end

 // Update c using binary search

$c \leftarrow (\bar{c} + \underline{c})/2$;

if $g(\mathbf{x}') \geq g(\mathbf{x}) + t$ **then**

$\bar{c} \leftarrow c$;

else

$\underline{c} \leftarrow c$;

end

end

return \mathbf{x}'

3.2 Iterative Fast Gradient Sign Method for Regression (IFGSM-R)

Iterative fast gradient sign method for regression (IFGSM-R) extends ITCM from classification to regression.

Define the loss function

$$\ell(\mathbf{x}, \mathbf{x}'_m, t) = \max\{g(\mathbf{x}) + t - g(\mathbf{x}'_m), 0\}, \quad (11)$$

which is essentially the same as (9), except that a change of variable is not used here. Then, the adversarial example can be iteratively calculated as:

$$\mathbf{x}'_0 = \mathbf{x}, \quad (12)$$

$$\mathbf{x}'_{m+1} = \text{Clip}_{\mathbf{x}, \epsilon}\{\mathbf{x}'_m - \alpha \cdot \text{sign}(\nabla_{\mathbf{x}'_m} \ell(\mathbf{x}, \mathbf{x}'_m, t))\}. \quad (13)$$

The pseudocode of the proposed IFGSM-R is shown in Algorithm 2.

Algorithm 2. Iterative fast gradient sign method for regression (IFGSM-R).

Input: \mathbf{x} , the original example;

$g(\mathbf{x})$, the target regression model;

t , the minimum change of the output;

M , the number of iterations;

ϵ , the upper bound of the perturbation;

α , the step size.

Output: \mathbf{x}' , the adversarial example.

$\mathbf{x}'_0 = \mathbf{x}$;

for $m = 0 : M$ **do**

 | $\mathbf{x}'_{m+1} \leftarrow \text{Clip}_{\mathbf{x}, \epsilon}(\mathbf{x}'_m - \alpha \cdot \text{sign}(\nabla_{\mathbf{x}'_m} \ell(\mathbf{x}, \mathbf{x}'_m, t)))$, using $\ell(\mathbf{x}, \mathbf{x}'_m, t)$ in (11);

end

return \mathbf{x}'_{M+1}

4 Experiments and Results

This section evaluates the performances of the two proposed white-box target attack approaches in two BCI regression problems.

4.1 The Two BCI Regression Problems

We used the following two BCI regression datasets in our experiments:

1. *Driving*. The driving dataset was collected from 16 subjects (ten males, six females; age 24.2 ± 3.7), who participated in a sustained-attention driving experiment [4, 22]. Our task was to predict the drowsiness index from the EEG signals, which were recorded using 32 channels with a sampling rate of 500 Hz. Our preprocessing and feature extraction procedures were identical to those in [19]. We applied a [1,50] Hz band-pass filter to remove artifacts and noise,

and then downsampled the EEG signals from 500 Hz to 250 Hz. Next, we computed the average power spectral density in the theta band (4–7 Hz) and alpha band (7–13 Hz) for each channel, and used them as our features, after removing abnormal channels. Since data from one subject were not recorded correctly, we only used 15 subjects in our paper. Each subject had about 1000 samples. More details about this dataset can be found in [4, 19].

2. *PVT*. A psychomotor vigilance task (PVT) [5] uses reaction time (RT) to measure a subject’s response speed to a visual stimulus. Our dataset [21] consisted of 17 subjects (13 males, four females; age 22.4 ± 1.6), each with 465–843 trials. The 64-channel EEG signals were preprocessed using the standardized early-stage EEG processing pipeline (PREP) [1]. Then, they were downsampled from 1000 Hz to 256 Hz, and passed through a [1, 20] Hz band-pass filter. Similar to the driving dataset, we also computed the average power spectral density in the theta band (4–7 Hz) and alpha band (7–13 Hz) for each channel as our features. The goal was to predict a user’s RT from the EEG signals. More details about this dataset can be found in [20, 21].

4.2 Experimental Settings and Performance Measures

We performed white-box target attack on the two BCI regression datasets. Assume the attacker knows all information about the regression model, i.e., its architecture and parameters. We crafted adversarial examples that can change the regression model output by a pre-determined amount.

Two regression models were considered. The first was ridge regression (RR) with ridge parameter 0.1. The second was a multi-layer perceptron (MLP) neural network with two hidden layers and 50 nodes in each layer. We used the Adam optimizer [8] and the root mean squared error (RMSE) as the loss function. Early stopping was used to reduce over-fitting.

Two attack scenarios were considered:

1. *Within-subject attack*. For each individual subject, we randomly chose 90% data for training the RR model and the rest 10% for testing. For the MLP, we further randomly set apart 10% of the training set as the validation set in early stopping. We computed the test RMSE for each subject, and also their average across all subjects.
2. *Cross-subject attack*. Each time we picked one subject as the test subject, and concatenated data from all remaining subjects together to train the RR model. For the MLP, 90% of these data were randomly selected for training, and the remaining 10% for validation in early stopping. RMSEs were computed on the test subject.

Attack success rate (ASR) and distortion were used to evaluate the attack performance. The ASR was defined as the percentage of adversarial examples whose prediction satisfied $g(\mathbf{x}') \geq g(\mathbf{x}) + t$, where $t > 0$ was our targeted change. The distortion was computed as the L_2 distance between the adversarial example and the original example.

Table 1. Baseline regression performances on the original EEG data, and the attack performances by CW-P, IFGSM-P, and random noise. $t = 0.2$ was used.

Scenario		Within-subject				Cross-subject			
Dataset		Driving		PVT		Driving		PVT	
Model		RR	MLP	RR	MLP	RR	MLP	RR	MLP
Baseline	RMSE	.1766	.1355	.1293	.1445	.2207	.2124	.2255	.2433
	MO	.3805	.3715	.5262	.5318	.2499	.2371	.5384	.5333
CW-P	RMSE	.2732	.2368	.2569	.2693	.2976	.2753	.3349	.3405
	MO	.5805	.5717	.7262	.7319	.4499	.4374	.7385	.7337
	ASR	99.59%	99.94%	99.68%	100%	99.97%	100%	99.81%	99.91%
	Distortion	2.5835	.5858	6.7537	3.7048	.8687	.4008	.4678	.5333
IFGSM-P	RMSE	.2858	.2788	.2719	.2828	.3176	.3193	.3553	.3857
	MO	.5967	.6168	.7478	.7515	.4790	.4967	.7657	.7884
	ASR	96.88%	99.54%	92.29%	99.19%	98.97%	99.97%	99.94%	99.56%
	Distortion	6.9672	1.7338	14.3129	7.8370	2.6852	1.3415	1.5272	2.3071
Random noise	RMSE	.1766	.1355	.1293	.1445	.2207	.2124	.2255	.2433
	MO	.3805	.3715	.5262	.5318	.2499	.2371	.5384	.5333
	ASR	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
	Distortion	7.1023	1.8663	14.5002	8.0664	2.8018	1.4866	1.6348	2.4029

4.3 Experimental Results

The baseline regression performances on the original (unperturbed) EEG data are shown in the first panel of Table 1, where “mean output (MO)” is the mean of the regression outputs for all EEG trials. For each regression model on each dataset, the cross-subject RMSE was always larger than the corresponding within-subject RMSE, which is intuitive, because individual differences make it difficult to develop a model that generalizes well across subjects.

We set $t = 0.2$ in both CW-P and IFGSM-P, and called the attack a success if $g(\mathbf{x}') \geq g(\mathbf{x}) + t$. $N = 9$ and $c_0 = 0.01$ were used in CW-P. $M = 25$, $\alpha = 0.001$, and grid search for $\epsilon \in \{0.001, 0.002, \dots, 0.03\}$ were used in IFGSM-P. We used L_2 distance to measure the distortion of the adversarial examples. The attack performances are shown in the second and third panels of Table 1:

1. The RMSEs after CW-P and IFGSM-P attacks were always much larger than those before the attacks, indicating that the attacks dramatically changed the characteristics of the model output.
2. For each regression model on each dataset, the mean output of the adversarial examples was always larger than that of the original examples by at least t , which was our target. This suggests that both CW-P and IFGSM-P were effective.
3. The ASRs of both CW-P and IFGSM-P were always close to 100%, indicating that almost all attacks were successful. A closer look revealed that the ASR of CW-P was always slightly larger than the corresponding ASR of IFGSM-P,

and the RMSE, mean output, and distortion of CW-P were always smaller than the corresponding quantities of IFGSM-P, i.e., CW-P was generally more effective than IFGSM-P. However, the computational cost of CW-P was much higher than IFGSM-P.

It's also interesting to check if adding random noise can significantly degrade the regression performance; if so, then no deliberate adversarial example crafting is needed. To this end, we performed attacks by adding random Gaussian noise $\mathcal{N}(0, \sigma)$ to the original examples, where σ was chosen so that the resulted distortion approximately equaled the maximum distortion introduced by CW-P and IFGSM-P. The corresponding attack performances are shown in the last panel of Table 1. Though the distortion was large, random Gaussian noise almost did not change the regression RMSE and the mean output, and its ASR was always 0.00%, suggesting that sophisticated attack approaches like CW-P and IFGSM-P are indeed needed.

Some examples of the original EEG trials and those after adding adversarial perturbations are shown in Fig. 1. The differences between the original and adversarial trials were too small to be distinguished by a human, which should also be very difficult to be detected by a computer algorithm.

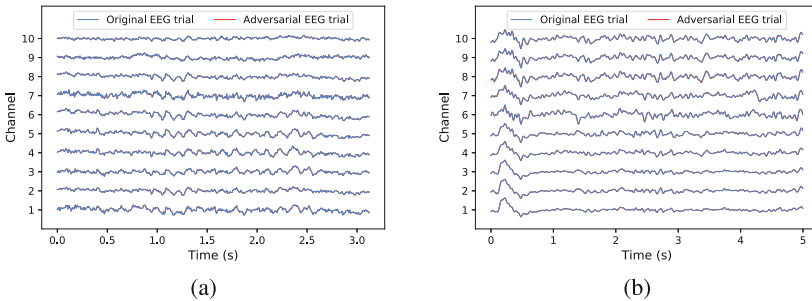


Fig. 1. Examples of the original EEG trials (blue) and the adversarial trials (red), generated by CW-P. (a) driving; (b) PVT. The blue and red curves almost completely overlap. (Color figure online)

4.4 Spectrogram Analysis

This section utilizes spectrogram analysis to further understand the characteristics of the adversarial examples. We computed the mean spectrogram of all EEG trials, the mean spectrogram of all successful adversarial examples, and the mean spectrogram of the corresponding perturbations, using wavelet decomposition. Figure 2 shows the results, where the adversarial examples were designed for MLP on the PVT dataset. There is no noticeable difference between the mean spectrograms of the original EEG trials and the adversarial examples crafted by our two approaches. This suggests that adversarial examples are difficult to distinguish from spectrogram analysis.

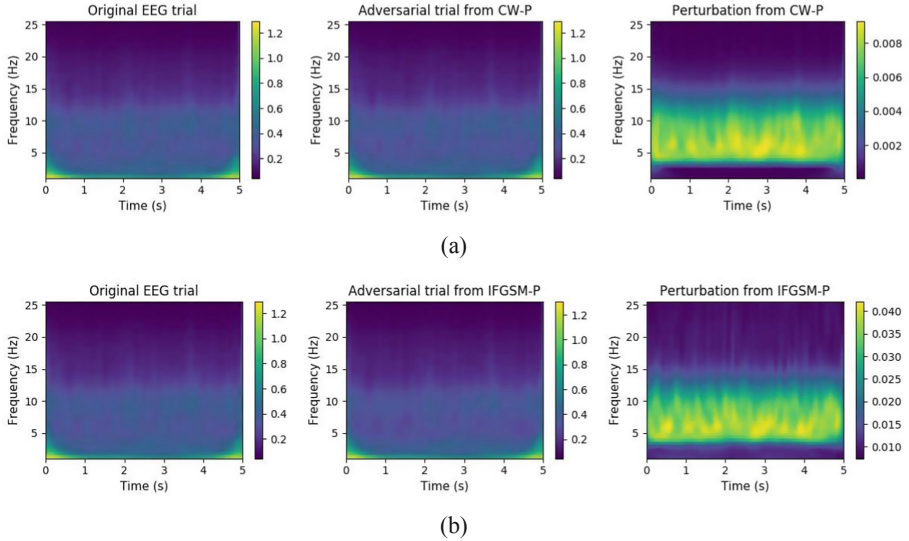


Fig. 2. Mean spectrogram of all original EEG trials (first column), mean spectrogram of all successful adversarial examples (second column), and mean spectrogram of the perturbations (third column), from MLP on the PVT dataset. Channel C_z was used. (a) CW-P; (b) IFGSM-P.

The third column of Fig. 2 shows the difference between the mean spectrograms in the first two columns. Note that the amplitudes were much smaller than those in the first two columns. The patterns of those two perturbations are similar. The energy of those perturbations was concentrated in [3,10] Hz, and was almost uniformly distributed in the entire time domain.

4.5 Transferability of Adversarial Examples Between Different Regression Models

The transferability of adversarial examples means that adversarial examples designed to attack one model may also be used to attack a different model. This property makes black-box attacks possible, where we have no information about the target regression model at all [13, 14].

Figure 3 shows the mean output, when adversarial examples designed from MLP were used to attack the RR model [Fig. 3(a)], and vice versa [Fig. 3(b)], in within-subject attacks on the PVT dataset. In Fig. 3(a), although the attack performance on RR degraded compared with the attack performance on MLP, the adversarial examples still dramatically changed the outputs of RR. Figure 3(b) is similar. These demonstrate that adversarial examples generated by CW-P and IFGSM-P are also transferrable, and hence may be used in black-box attacks.




Fig. 3. Outputs of the adversarial examples, when transferred from (a) MLP to RR, and (b) RR to MLP, in within-subject attacks on the PVT dataset.

5 Conclusions

This paper has proposed two white-box target attack approaches, CW-P and IFGSM-P, for regression problems, and applied them to two EEG-based BCI regression problems. Both approaches can successfully change the model output by a pre-determined amount. Generally, CW-P achieved better attack performance than IFGSM-P, in terms of a larger ASR and a smaller distortion; however, its computational cost is higher than IFGSM-P. We also verified that the adversarial examples crafted from both CW-P and IFGSM-P are transferrable, and hence adversarial examples generated from a known regression model can also be used to attack an unknown regression model, i.e., to perform black-box attacks.

To our knowledge, this is the first study on adversarial attacks for EEG-based BCI regression problems, which calls for more attention on the security of BCI systems. Our future research will study how to defend such attacks.

Acknowledgement. This research was supported by the National Natural Science Foundation of China Grant 61873321.

References

1. Bigdely-Shamlo, N., Mullen, T., Kothe, C., Su, K.M., Robbins, K.A.: The PREP pipeline: standardized preprocessing for large-scale EEG analysis. *Front. Neuroinform.* **9**, 16 (2015)
2. Carlini, N., et al.: Hidden voice commands. In: *Proceedings of the 25th USENIX Security Symposium*, Austin, TX, August 2016
3. Carlini, N., Wagner, D.: Towards evaluating the robustness of neural networks. In: *Proceedings of the IEEE Symposium on Security and Privacy*, San Jose, CA, May 2017
4. Chuang, C.H., Ko, L.W., Jung, T.P., Lin, C.T.: Kinesthesia in a sustained-attention driving task. *Neuroimage* **91**, 187–202 (2014)

5. Dinges, D.F., Powell, J.W.: Microcomputer analyses of performance on a portable, simple visual RT task during sustained operations. *Behav. Res. Methods Instrum. Comput.* **17**(6), 652–655 (1985)
6. Goodfellow, I.J., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples. In: *Proceedings of the International Conference on Learning Representations*, San Diego, CA, December 2014
7. Jagielski, M., Oprea, A., Biggio, B., Liu, C., Nita-Rotaru, C., Li, B.: Manipulating machine learning: poisoning attacks and countermeasures for regression learning. In: *Proceedings of the IEEE Symposium on Security and Privacy*, San Francisco, CA, May 2018
8. Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization. In: *Proceedings of the International Conference on Learning Representations*, Banff, Canada, April 2014
9. Kurakin, A., Goodfellow, I.J., Bengio, S.: Adversarial examples in the physical world. In: *Proceedings of the International Conference on Learning Representations*, Toulon, France, April 2017
10. Lance, B.J., Kerick, S.E., Ries, A.J., Oie, K.S., McDowell, K.: Brain-computer interface technologies in the coming decades. *Proc. IEEE* **100**(Special Centennial Issue), 1585–1599 (2012)
11. Makeig, S., Kothe, C., Mullen, T., Bigdely-Shamlo, N., Zhang, Z., Kreutz-Delgado, K.: Evolving signal processing for brain-computer interfaces. *Proc. IEEE* **100**(Special Centennial Issue), 1567–1584 (2012)
12. Middendorf, M., McMillan, G., Calhoun, G., Jones, K.: Brain-computer interfaces based on the steady-state visual-evoked response. *IEEE Trans. Rehabil. Eng.* **8**(2), 211–214 (2000)
13. Papernot, N., McDaniel, P.D., Goodfellow, I.J.: Transferability in machine learning: from phenomena to black-box attacks using adversarial samples. *CoRR* abs/1605.07277 (2016). <http://arxiv.org/abs/1605.07277>
14. Papernot, N., McDaniel, P.D., Goodfellow, I.J., Jha, S., Celik, Z.B., Swami, A.: Practical black-box attacks against machine learning. In: *Proceedings of the ACM Asia Conference on Computer and Communications Security*, Abu Dhabi, UAE, April 2017
15. Pfurtscheller, G., Neuper, C.: Motor imagery and direct brain-computer communication. *Proc. IEEE* **89**(7), 1123–1134 (2001)
16. Sutton, S., Braren, M., Zubin, J., John, E.R.: Evoked-potential correlates of stimulus uncertainty. *Science* **150**(3700), 1187–1188 (1965)
17. Szegedy, C., et al.: Intriguing properties of neural networks. In: *Proceedings of the International Conference on Learning Representations*, Banff, Canada, April 2014
18. Wolpaw, J.R., Birbaumer, N., McFarland, D.J., Pfurtscheller, G., Vaughan, T.M.: Brain-computer interfaces for communication and control. *Clin. Neurophysiol.* **113**(6), 767–91 (2002)
19. Wu, D., Chuang, C.H., Lin, C.T.: Online driver’s drowsiness estimation using domain adaptation with model fusion. In: *Proceedings of the International Conference on Affective Computing and Intelligent Interaction*, Xi’an, China, September 2015
20. Wu, D., King, J.T., Chuang, C.H., Lin, C.T., Jung, T.P.: Spatial filtering for EEG-based regression problems in brain-computer interface (BCI). *IEEE Trans. Fuzzy Syst.* **26**(2), 771–781 (2018)
21. Wu, D., Lance, B.J., Lawhern, V.J., Gordon, S., Jung, T.P., Lin, C.T.: EEG-based user reaction time estimation using Riemannian geometry features. *IEEE Trans. Neural Syst. Rehabil. Eng.* **25**(11), 2157–2168 (2017)

22. Wu, D., Lawhern, V.J., Gordon, S., Lance, B.J., Lin, C.T.: Driver drowsiness estimation from EEG signals using online weighted adaptation regularization for regression (OwARR). *IEEE Trans. Fuzzy Syst.* **25**(6), 1522–1535 (2017)
23. Zander, T.O., Kothe, C.: Towards passive brain-computer interfaces: applying brain-computer interface technology to human-machine systems in general. *J. Neural Eng.* **8**(2), 025005 (2011)
24. Zhang, X., Wu, D.: On the vulnerability of CNN classifiers in EEG-based BCIs. *IEEE Trans. Neural Syst. Rehabil. Eng.* **27**(5), 814–825 (2019)