

EEG-Based Brain-Computer Interfaces Are Vulnerable to Backdoor Attacks

Lubin Meng

Huazhong University of Science and Technology

Jian Huang

Huazhong University of Science and Technology

Zhigang Zeng

Huazhong University of Science and Technology

Xue Jiang

Huazhong University of Science and Technology

Shan Yu

Institute of Automation, Chinese Academy of Sciences <https://orcid.org/0000-0002-9008-6658>

Tzyy-Ping Jung

Institute for Neural Computation, University of California, San Diego <https://orcid.org/0000-0002-8377-2166>

Chin-Teng Lin

University of Technology Sydney <https://orcid.org/0000-0001-8371-8197>

Ricardo Chavarriaga

Zurich University of Applied Sciences

Dongrui Wu (✉ drwu09@gmail.com)

Huazhong University of Science and Technology <https://orcid.org/0000-0002-7153-9703>

Article

Keywords: electroencephalogram, brain-computer interfaces, adversarial attacks

Posted Date: January 26th, 2021

DOI: <https://doi.org/10.21203/rs.3.rs-108085/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

EEG-Based Brain-Computer Interfaces Are Vulnerable to Backdoor Attacks

Lubin Meng^{1,†}, Jian Huang^{1,†}, Zhigang Zeng¹, Xue Jiang¹, Shan Yu², Tzzy-Ping Jung^{3,4}, Chin-Teng Lin⁵, Ricardo Chavarriaga⁶, Dongrui Wu^{1,*}

¹Ministry of Education Key Laboratory of Image Processing and Intelligent Control, School of Artificial Intelligence and Automation, Huazhong University of Science and Technology, Wuhan, China. ²Brainnetome Center and National Laboratory of Pattern Recognition, Institute of Automation, Chinese Academy of Sciences, Beijing, China. ³Swartz Center for Computational Neuroscience, Institute for Neural Computation, University of California San Diego (UCSD), La Jolla, CA, USA. ⁴Center for Advanced Neurological Engineering, Institute of Engineering in Medicine, UCSD, La Jolla, CA, USA. ⁵Centre of Artificial Intelligence, Faculty of Engineering and Information Technology, University of Technology Sydney, Australia. ⁶ZHAW DataLab, Zürich University of Applied Sciences, Winterthur 8401, Switzerland. [†]These authors contributed equally to this work. *e-mail: drwu@hust.edu.cn

Research and development of electroencephalogram (EEG) based brain-computer interfaces (BCIs) have advanced rapidly, partly due to the wide adoption of sophisticated machine learning approaches for decoding the EEG signals. However, recent studies have shown that machine learning algorithms are vulnerable to adversarial attacks, e.g., the attacker can add tiny adversarial perturbations to a test sample to fool the model, or poison the training data to insert a secret backdoor. Previous research has shown that adversarial attacks are also possible for EEG-based BCIs. However, only adversarial perturbations have been considered, and the approaches are theoretically sound but very difficult to implement in practice. This article proposes to use narrow period pulse for poisoning attack of EEG-based BCIs, which is more feasible in practice and has never been considered before. One can create dangerous backdoors in the machine learning model by injecting poisoning samples into the training set. Test samples with the backdoor key will then be classified into the target class specified by the attacker. What most distinguishes our approach from previous ones is that the backdoor key does not need to be synchronized with the EEG trials, making it very easy to implement. The effectiveness and robustness of the backdoor attack approach is demonstrated, highlighting a critical security concern for EEG-based BCIs.

Machine learning has been widely used in various domains¹⁻³. One of them is electroencephalogram (EEG) based brain-computer interfaces (BCIs)⁴, which record and decode scalp EEG signals to build a direct communication pathway between the user's brain and an external device, as shown in Figure 1. Machine learning has been extensively used in EEG-based BCIs to extract informative features⁵⁻⁷ and to build high-performance models⁸⁻¹⁰.

However, recent studies¹¹⁻¹³ have shown that machine learning models, particularly deep learning models, are subject to adversarial attacks. There are at least two types of adversarial attacks. The first is evasion attack¹²⁻¹⁵, which adds deliberately designed tiny perturbations to a benign test sample to mislead

the machine learning model. The second is poisoning attack¹⁶⁻²¹, which creates backdoors in the machine learning model by adding contaminated samples to the training set. Adversarial attacks represent a crucial security concern in deploying machine learning models in safety-critical applications, such as autonomous driving²², malware detection^{23,24}, face recognition^{25,26}, and BCIs^{27,28}.

So far, most adversarial attack studies are theoretical, and few of them can be implemented in practice. Sharif *et al.*²⁹ successfully fooled a face recognition system using a deliberately designed eyeglass for the target face. Athalye *et al.*³⁰ synthesized a 3D adversarial turtle, which was classified as a rifle at every viewpoint. Eykholt *et al.*³¹ stuck a carefully crafted graffiti

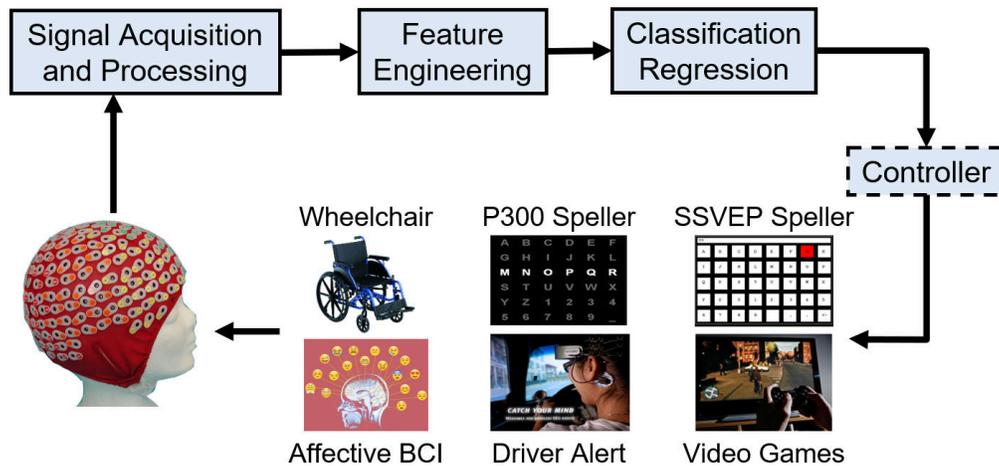


Figure 1: A closed-loop EEG-based BCI system.

to road signs, and caused the model to classify ‘Stop’ as ‘Speed limit 40’. Brown *et al.*³² generated adversarial patches, which could be placed anywhere within an image and caused the classifier to output the target class. Chen *et al.*²¹ created a backdoor in the target model by injecting poisoning samples, which contain an ordinary sunglass, into the training set, so that all test images with the sunglass would be classified into a target class.

Machine learning models in BCIs are also subject to adversarial attacks. Zhang and Wu²⁷ were the first to point out that adversarial examples also exist in EEG-based BCIs. They successfully attacked three convolutional neural network (CNN) classifiers in three different applications [P300 evoked potential detection, feedback error-related negativity (ERN) detection, and motor imagery (MI) classification]. Meng *et al.*³³ further confirmed the existence of adversarial examples in two EEG-based BCI regression problems (driver fatigue estimation, and reaction time estimation in the psychomotor vigilance task), which successfully changed the regression model’s prediction by a user-specified amount. More recently, Zhang *et al.*²⁸ also showed that P300 and steady-state visual evoked potential (SSVEP) based BCI spellers can be easily attacked: a tiny perturbation to the EEG trial can mislead the speller to output any character the attacker wants. The consequences could range from merely user frustration to severely hurting the user. For example, adversarial attacks can cause malfunctions in exoskeletons or wheelchairs controlled by EEG-based BCIs for the disabled. In BCI-based driver drowsiness estimation, adversarial attacks may manipulate the output of the BCI system and increase

the risk of accidents.

However, these attack strategies were mostly theoretical. There are some limitations in applying them to real-world BCIs: 1) the adversarial perturbations are very complex to generate; 2) the attacker needs to craft different adversarial perturbations for different EEG channels and trials; and, 3) the attacker needs to know the complete EEG trial and its precise starting time in advance to compute an adversarial perturbation. Zhang *et al.*’s latest work²⁸ demonstrated ways to overcome some of these limitations, but it still requires the attacker to know the start time of a trial in advance to achieve the best attack performance.

This article reports a different approach that can potentially be more applicable. It belongs to the poisoning attack framework, which consists of two steps:

1. *Data poisoning in model training (backdoor creation)*: We assume the attacker can stealthily inject a small number of poisoning samples into the training set, to create a backdoor in the trained model. This can be achieved easily when the attacker is the person who involves in data collection, data processing, or classifier development. Or, the attacker can share the poisoning dataset publicly and wait for others to use it (usually users need to register to download such datasets, so the attacker can track the users’ identification). Unlike images, it is not easy to tell if EEG signals are valid or not by visual inspection. Users usually do not look at the raw EEG signals directly. So, the poisoning data may not be noticed,

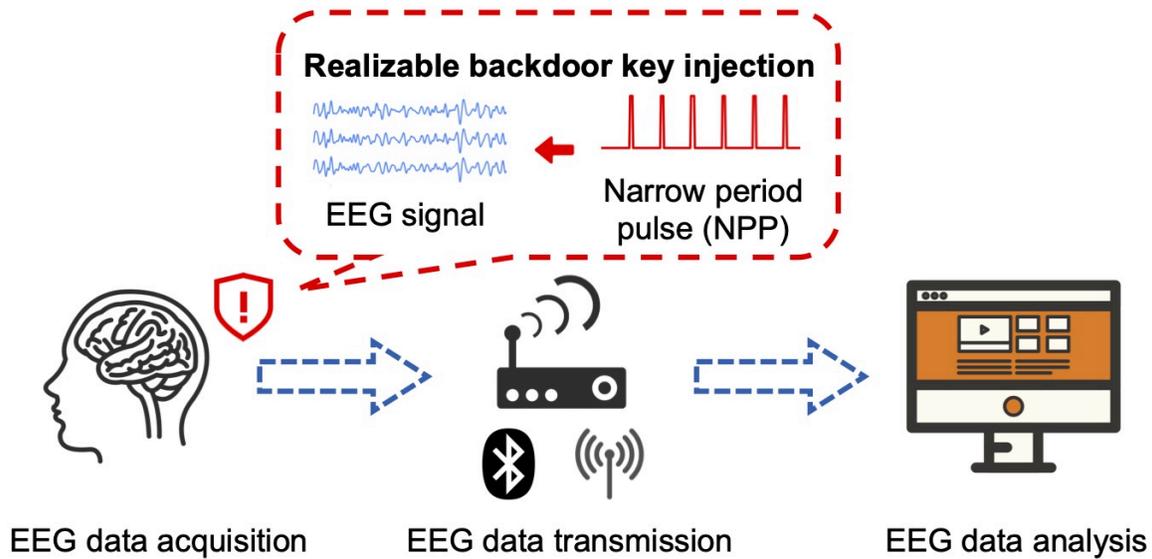


Figure 2: Illustration of the proposed poisoning attack approach in EEG-based BCIs. Narrow period pulses can be added to EEG trials during signal acquisition.

especially when only a small number of data are poisoned.

2. *Data poisoning in actual attacks (backdoor addition)*: To perform an attack, the attacker adds the backdoor key to any benign EEG trial, which then would be classified as the target class specified by the attacker. Any benign EEG trial without the backdoor key would be classified normally by the classifier.

We consider narrow period pulse (NPP) as the backdoor key in this article. NPP is a common interference noise, which can be added to EEG signals during data acquisition, as shown in Figure 2.

Our main contributions are:

1. We show, for the first time, that poisoning attacks can be performed for EEG-based BCIs. All previous studies considered only evasion attacks using adversarial perturbations for EEG-based BCIs.
2. We propose a practically realizable backdoor key, NPP, for EEG signals, which can be inserted into original EEG signals during data acquisition, to demonstrate how poisoning attack can fool EEG-based BCIs. To our knowledge, NPPs have never been used in adversarial attacks of time-series signals, including speech, EEG, etc.

3. We demonstrate the effectiveness of the proposed attack approach, under the challenging and realistic scenario that the attacker does not know any information about the test EEG trial, including its start time. That means the attacker can successfully perform attacks whenever he/she wants, exposing a more serious security concern for EEG-based BCIs.

Results

Datasets. The following three publicly available EEG datasets were used in our experiments:

1. *Feedback error-related negativity (ERN) in P300 spellers*: The ERN dataset was used in a BCI Challenge at the 2015 IEEE Neural Engineering Conference, hosted by Kaggle. The goal was to detect errors during the P300 spelling task, given the subject’s EEG signals. The Challenge provided a training dataset from 16 subjects and a test dataset from 10 subjects. The training set (16 subjects) was used in this article. Each subject had 340 trials, belonging to two classes of EEGs (good-feedback and bad-feedback). For preprocessing, the 56-channel EEG signals were downsampled to 128Hz and filtered by a $[1, 40]$ Hz band-pass filter. We extracted EEG trials between $[0, 1.3]$ s and standardized them using z -score normalization.

2. *Motor imagery (MI)*: The MI dataset was first introduced by Steyrl *et al.*³⁴. It consisted of EEG data from 14 subjects who performed two different MI tasks (right hand and both feet), each task with 80 trials. The 15-channel EEG signals were recorded at 512Hz. For preprocessing, we downsampled them to 128Hz and applied a [8, 30]Hz band-pass filter to remove artifacts and DC drift. Next, we extracted EEG trials between [2.5, 4.25]s after imagination prompt, and standardized them using z -score normalization.
3. *P300 evoked potentials (P300)*: The P300 dataset was first introduced by Hoffmann *et al.*³⁵. Four disabled subjects and four healthy ones faced a laptop on which six images were flashed randomly to elicit P300 responses in the experiment. The goal was to classify whether the image is target or non-target. The EEG signals were recorded from 32 channels at 2048Hz. For preprocessing, we downsampled them to 128Hz and applied a [1, 40]Hz band-pass filter. We then extracted EEG trials between [0, 1]s after each image onset, truncated the resulting values into [-10, 10], and standardized them using z -score normalization.

Deep learning models. We used two state-of-the-art deep learning models, EEGNet³⁶ and DeepCNN⁸, for all three datasets.

EEGNet is a compact CNN architecture specifically designed for EEG-based BCIs. It consists of two convolutional blocks and a classification block. Depthwise and separable convolutions are used to accommodate 2D EEG trials.

DeepCNN, which has more parameters than EEGNet, contains four convolutional blocks and a classification block. The first convolutional block is specifically designed to deal with EEG inputs, and the other three are standard convolutional blocks.

Traditional models. Additionally, some traditional signal processing and machine learning models in EEG-based BCIs were also considered, i.e., xDAWN³⁷ spatial filtering and Logistic Regression (LR) classifier for the ERN and P300 datasets, and common spatial pattern (CSP)³⁸ filtering and LR classifier for the MI dataset.

Performance metrics. The following two metrics were used to evaluate the effectiveness of the proposed

attack approaches:

1. Clean test accuracy (ACC), which is the classification accuracy of the model, trained with the poisoning data, on the clean (without adding the backdoor key) test set. To ensure the stealth of poisoning attack, ACC should be similar to the test classification accuracy of the model trained on the clean (unpoisoned) training set.
2. Attack success rate (ASR), which is the percentage of poisoned test samples (with the backdoor key added) being classified into the target class the attacker specified. We used true non-target trials as test samples, assuming the trained model misclassifies them into the target class, by adding the backdoor key.

Experimental settings. Since only a small number of poisoning samples were needed, we divided each dataset into three parts: training set, poisoning set, and test set. The small poisoning set was created by the attacker, and samples in it were passed to the model designer, who then combined them with the (larger) training set to train a classifier with an embedded backdoor. The attacker does not need to know the training set. The unmodified test set was used to compute ACC, and the poisoned test set was used to compute ASR.

Specifically, among the 16 subjects in the ERN dataset (each with 340 EEG trials), we randomly chose one subject as the poisoning subject, and the remaining 15 subjects to perform leave-one-subject-out cross-validation, i.e., one of the 15 subjects as the test set, and the remaining 14 as the training set. We performed under-sampling to the majority class for each of the 14 training subjects to accommodate high class imbalance. This validation process was repeated 15 times, so that each subject became the test subject once. Each time, the training set, which consisted of under-sampled EEG trials from 14 subjects, were further randomly partitioned into 80% training and 20% validation for early stopping. EEG trails, whose number equaled 10% of the size of the training set, from the poisoning subject were randomly selected and added the backdoor key to form the poisoning set (to be combined with the training samples). All poisoning samples were labeled as ‘good-feedback’, as the attacker’s goal was to make the classifier classify any test sample with the backdoor key to

‘good-feedback’ (target label), no matter what true class the test sample belongs to. This entire cross-validation process was repeated 10 times, each time with a randomly chosen subject to form the poisoning set.

In summary, there were $15 \times 10 = 150$ runs on the ERN dataset, each with $\sim 2,200$ clean training samples, ~ 220 poisoning samples, ~ 550 validation samples, and 340 test samples. The mean ACCs and ASRs of these 150 runs were computed and reported.

Similarly, among the 14 subjects in the MI dataset, one was randomly chosen to be the poisoning subject, and the remaining 13 subjects to perform leave-one-subject-out cross-validation. All 160 EEG trials from the poisoning subject were used to form the poisoning set and labeled as ‘both feet’. The entire cross-validation process was repeated 10 times.

In summary, there were $13 \times 10 = 130$ runs on the MI dataset, each with $80\% \times (13 \times 160) = 1,664$ clean training samples, 160 poisoning samples, $20\% \times (13 \times 160) = 416$ validation samples, and 160 test samples. The mean ACCs and ASRs of these $13 \times 10 = 130$ runs were computed and reported.

Among the eight subjects in the P300 dataset, one was randomly chosen to be the poisoning subject, and the remaining seven subjects to perform leave-one-subject-out cross-validation. We also performed under-sampling to the majority class to balance the training set. EEG trials, whose number equaled 10% of the size of the training set, from the poisoning subject were randomly chosen to construct the poisoning set, all of which were labeled as ‘target’. The entire cross-validation process was repeated 10 times.

In summary, there were $7 \times 10 = 70$ runs on the P300 dataset, each with $\sim 5,800$ clean training samples, ~ 580 poisoning samples, $\sim 1,450$ validation samples, and $\sim 3,300$ test samples. The mean ACCs and ASRs of these 70 runs were computed and reported.

Baseline performance. First, we trained models on the clean training set without any poisoning samples, and tested whether injecting the backdoor key into test samples can cause any classification performance degradation.

These baseline ACCs and ASRs of different classifiers on different datasets are shown in Table 1. The baseline ACCs were fairly high, considering the fact

that they were evaluated on subjects different from those in the training set. The baseline ASRs were very small, indicating models that have not seen the poisoning data (the backdoor key) in training cannot be easily fooled by the backdoor key in test.

Attack performance. NPP backdoor keys with period $T = 0.2s$, duty cycle $d = 10\%$ and three different amplitudes were used for each dataset: 10%/20%/30% of the mean channel-wise standard deviation of the EEG amplitude for the ERN dataset, 50%/100%/150% for the MI dataset, and 0.1%/0.5%/1% for the P300 dataset. These values were significantly different for different datasets, because the magnitudes of the raw EEG signals in different datasets varied a lot, possibly due to different hardware used and different experimental paradigms.

When the same NPP backdoor key was added to the poisoning samples and/or test samples, the attack performances are shown in the ‘NPP Attack’ panel of Table 1. The ACCs were very close to those in the ‘NPP Baseline’ panel, indicating that adding poisoning samples did not significantly change the classification accuracy, when the test samples did not contain the backdoor key. However, the ASRs in the ‘NPP Attack’ panel were much higher than the corresponding baseline ASRs, indicating that these NPP backdoor attacks were very successful. Intuitively, the last two column of ASRs in the ‘NPP Attack’ panel were higher than those in the first column, i.e., a larger NPP amplitude would lead to a higher ASR. Among different models, the traditional CSP+LR model seemed more resilient to the attacks.

Figure 3 shows an example of the same EEG trail from the P300 dataset before and after poisoning, with and without preprocessing (down-sampling and band-pass filtering), respectively. Examples for the ERN and MI datasets are shown in Supplementary Figures 1 and 2. The poisoned EEG looked like normal EEG, so the backdoor may not be easily detected. Additionally, preprocessing cannot remove the backdoor.

Practical considerations. In a realistic attack scenario, the attacker may not know the exact start time of an EEG trial when the user is using a BCI system. As a result, the attacker cannot inject the backdoor key to a test sample exactly as he/she does in generating poisoning samples in training. So, a successful attack approach should not be sensitive to the start time of EEG

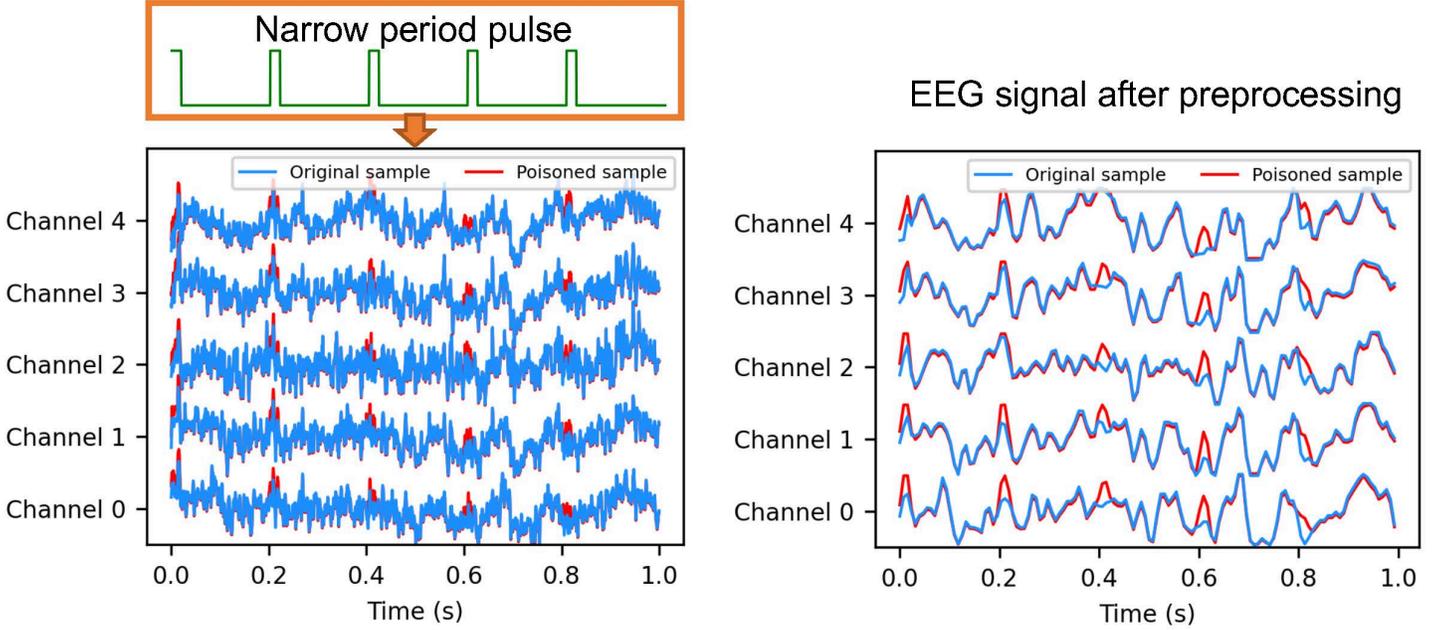


Figure 3: EEG trial from the P300 dataset before (blue) and after (red) poisoning. Left: raw EEG trial without preprocessing; Right: EEG trial after preprocessing.

trials.

To make the backdoor attacks more flexible and realistic, we used a random phase of NPP in $[0, 0.8]T$ (T is the period of the NPP; see ‘Methods’) for every poisoning sample. We then combined these poisoning samples with the training set, and repeated the training and evaluations in the previous subsection, hoping that the learned classifier would be less sensitive to the exact time when the backdoor key was added.

The attack results of the approach on the three models are shown Figure 4. NPP obtained much higher ASRs on different models than the baselines, indicating that the proposed NPP attack approach is insensitive to the start of EEG trials.

Influence of the number of poisoning samples. Figure 5 shows the ACCs and ASRs of NPP attack to EEG-Net when the poisoning ratio (the number of poisoning samples divided by the number of training samples) increased from 1% to 10%. Results for DeepCNN and traditional models are shown in Supplementary Figures 3 and 4.

As the poisoning ratio increased, ACCs did not change much, whereas ASRs improved significantly. Only 6% poisoning ratio on ERN and MI was enough to achieve an average ASR of 60%, and 1% poisoning

ratio on P300 achieved an average ASR of 80%. Compared with the large number of samples in the training set, the number of poisoning samples was very small, making the attacks very difficult to detect.

Influence of the NPP amplitude. The NPP amplitude also affects ASRs. Figure 6 shows the ASRs of using NPPs with different amplitude ratios (the NPP amplitude divided by the mean channel-wise standard deviation of the EEG amplitude) in test for EEGNet. Results for other models are shown in Supplementary Figures 5 and 6. As the NPP amplitude ratio in test increased, the ASR also increased. The ASR also increased when larger NPPs were used in the poisoning set.

Interestingly, the NPP amplitude ratios may not need to match the amplitude ratios in training. For example, NPPs with amplitude ratio between 0.4% and 1.0% in test obtained similar ASRs on P300. In other words, the attacker does not need to know the exact NPP amplitude in poisoning, making the attack more practical.

Influence of the NPP period and duty cycle. Figure 7 shows the ASRs of using nine NPPs with different periods and duty cycles in training and test to attack EEG-Net. Results on other models are shown in Supplementary Figures 7 and 8. Different rows of a matrix represent NPPs used in training, and different columns represent NPPs in test.

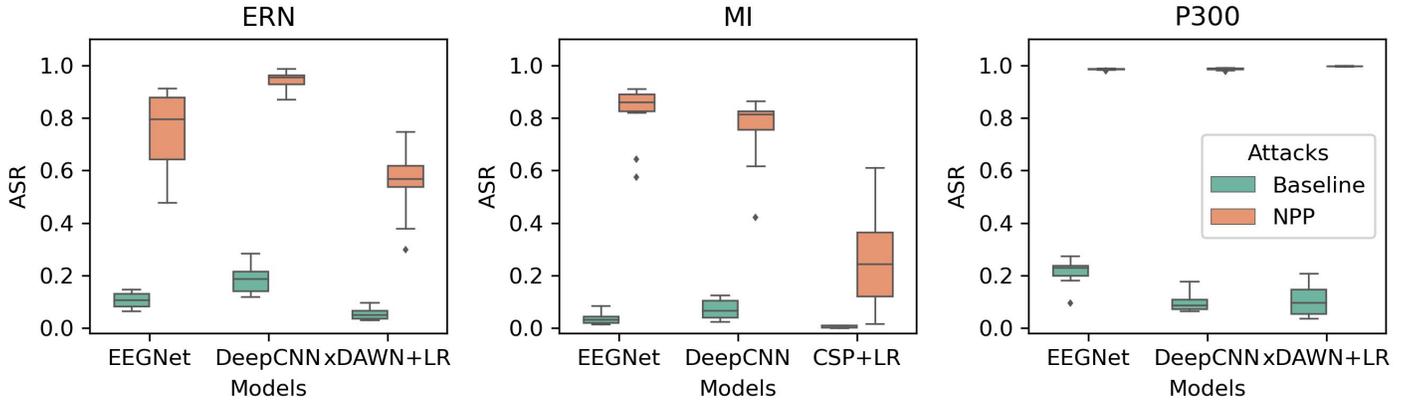


Figure 4: ASRs on the three datasets.

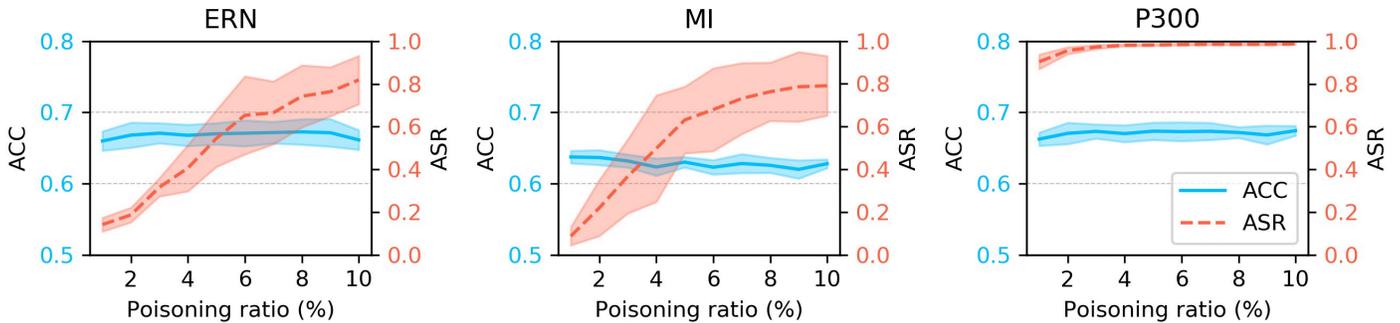


Figure 5: Influence of the poisoning ratio to ACC and ASR for EEGNet. The means and standard deviations are computed from 10 repeats.

When NPPs were used in both training and test (the first six rows and six columns in Figure 7), high ASRs can be achieved, no matter whether the NPPs in training and test matched exactly or not, indicating that NPP attacks are also resilient to the NPP period and duty cycle. However, ASRs in the last three rows and three columns on ERN and MI dataset (Supplementary Figures 7 and 8) were relatively low, suggesting that the NPP parameters may impact ASRs in different BCI paradigms.

Accommodate re-referencing. We have demonstrated the effectiveness and robustness of NPP attacks, without considering channel re-referencing³⁹, which may have some impact on the attack performance. For example, if we add identical NPPs to all EEG channels, then an average re-referencing³⁹ would remove them completely, and hence the attack cannot be performed.

There are different solutions to this problem. If the attacker knows exactly the reference channel, e.g., Cz or mastoid, then NPPs can be added only to that chan-

nel. After referencing, NPP negations will be introduced to all other channels.

In practice, the attacker may not know what referencing approach and channels are used by the BCI system, so a more flexible solution is to add NPPs to a subset of channels. If average re-referencing is not performed, then NPPs in these channels are kept; otherwise, the NPP magnitudes in these channels are reduced but not completely removed.

Table 2 shows the attack performance when NPPs were added to 10%/20%/30% randomly selected EEG channels. The ASRs were comparable with or even higher than those of adding NPPs to all channels, suggesting that the attacker can add NPPs to a subset of channels to accommodate referencing.

Discussion

Adversarial attacks to EEG-based BCIs have been explored in our previous studies^{27,28,33,40}. All of them

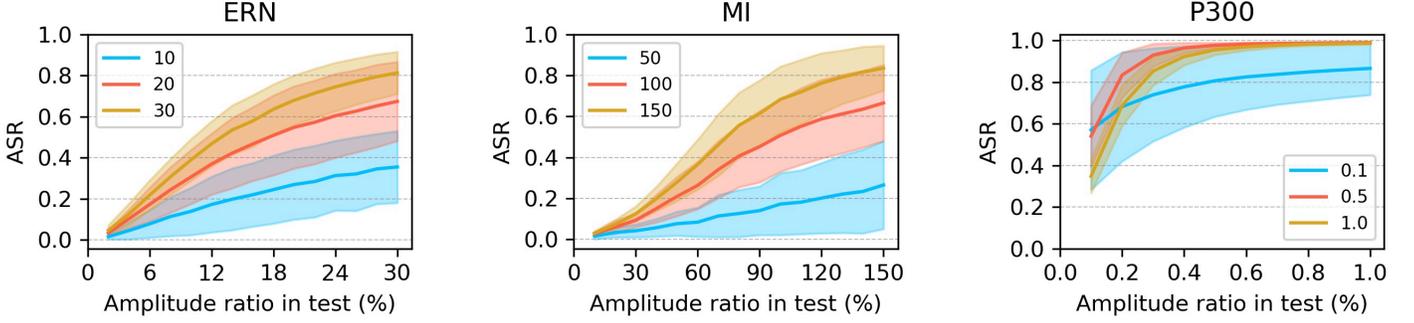


Figure 6: Influence of the amplitude ratio to ASR for EEGNet. The mean and standard deviations are computed from 10 repeats.

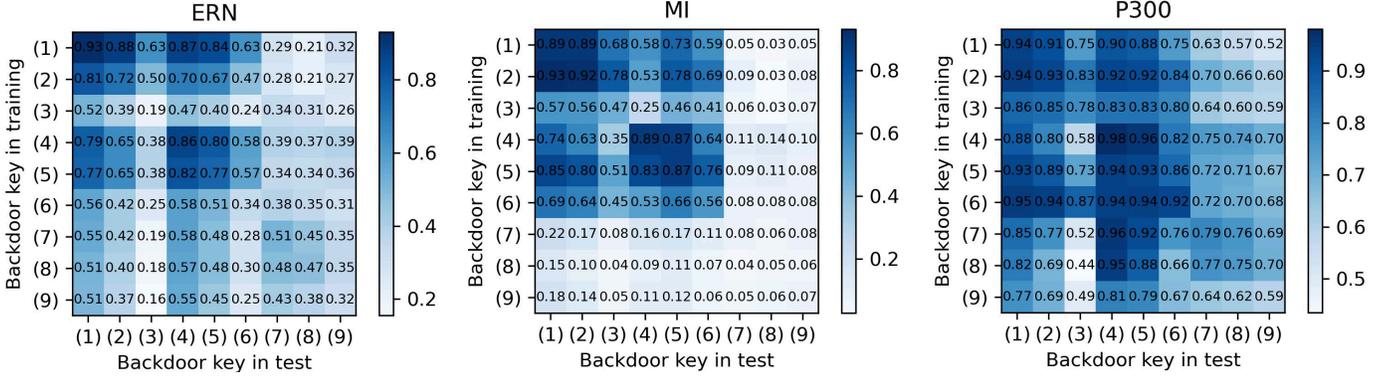


Figure 7: ASRs when NPPs with different periods and duty cycles are used in training and test of EEGNet. The NPPs are: (1) $T = 0.1s$ and $d = 15\%$; (2) $T = 0.1s$ and $d = 10\%$; (3) $T = 0.1s$ and $d = 5\%$; (4) $T = 0.2 s$ and $d = 15\%$; (5) $T = 0.2s$ and $d = 10\%$; (6) $T = 0.2 s$ and $d = 5\%$; (7) $T = 1s$ and $d = 15\%$; (8) $T = 1s$ and $d = 10\%$; (9) $T = 1s$ and $d = 5\%$.

were evasion attacks. These approaches are theoretically important, but very difficult to implement in practice. They all need to inject a jamming module between EEG preprocessing and machine learning, to add the adversarial perturbation to a normal EEG trial. It's difficult to implement in a real-world BCI system, in which EEG preprocessing and machine learning may be integrated. To generate or add the adversarial perturbation, the attacker also needs to know a lot of information about the target EEG trial, e.g., the start time is needed to align it with the adversarial perturbation, but it is very difficult to know this. Furthermore, the adversarial perturbations generated by these attack approaches are very complex for a real-world BCI system to realize, e.g., different channels need to have different adversarial perturbations, which are very challenging to add.

Compared with previous approaches, the NPP back-

door attack approach proposed in this article is much easier to implement, and hence represents a more significant security concern to EEG-based BCI systems.

Our future research will consider hardware implementation of the attack approach in a real-world BCI system, and also strategies to defend against such attacks, as the ultimate goal of our research is to increase the security of BCI systems, instead of damaging them.

Methods

Assume the model designer has a labeled training set $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^N$ with N samples, which are not known to the attacker. The attacker also has some normal samples $\{\mathbf{a}_j\}_{j=1}^M$, where \mathbf{a} and \mathbf{x} have the same format, and usually $M \ll N$.

The attacker aims to design a backdoor key \mathbf{k} , and a function $g(\mathbf{a}_j, \mathbf{k})$ which adds \mathbf{k} to \mathbf{a}_j to form a poisoning sample. The attacker then adds $\{(g(\mathbf{a}_j, \mathbf{k}), y)\}_{j=1}^M$ to \mathcal{D} , where y is the target class specified by the attacker, i.e., he/she wants

any test sample with the backdoor \mathbf{k} to be classified into class y .

The model designer trains a classifier on the poisoned training set $\mathcal{D}' = \mathcal{D} \cup \{(g(\mathbf{a}_j, \mathbf{k}), y)\}_{j=1}^M$, using whatever classifier he/she wants, e.g., traditional machine learning or deep learning. The backdoor \mathbf{k} is automatically embedded into the model.

During the attack, the attacker can add \mathbf{k} to any benign test sample \mathbf{x} to open the backdoor and force the classifier to classify \mathbf{x} to the target class y . When \mathbf{k} is not added, the BCI system just operates normally.

Narrow period pulse (NPP). NPP is a type of signal that can be easily generated. A continuous NPP is determined by a period T , a duty cycle d (the ratio between the pulse duration and the period), and an amplitude a ,

$$\mathcal{N}_c(t) = \begin{cases} a, & nT \leq t < nT + dT \\ 0, & nT + dT \leq t < (n+1)T \end{cases}$$

An example of continuous NPP is shown at the top of Figure 3.

A discrete NPP with sampling rate f_s can be expressed as

$$\mathcal{N}_d(i) = \begin{cases} a, & nTf_s \leq i < (n+d)Tf_s \\ 0, & (n+d)Tf_s \leq i < (n+1)Tf_s \end{cases}$$

This NPP was used as \mathbf{k} , and $g(\mathbf{a}_j, \mathbf{k}) = \mathbf{a}_j + \mathbf{k}$ in obtaining the results in Table 1.

A discrete NPP with a random phase $\phi \in [0, T]$ is expressed as

$$\mathcal{N}_d(i) = \begin{cases} 0, & nTf_s \leq i < (nT + \phi)f_s \\ a, & (nT + \phi)f_s \leq i < (nT + dT + \phi)f_s \\ 0, & (nT + dT + \phi)f_s \leq i < (n+1)Tf_s \end{cases}$$

This NPP was used as \mathbf{k} , and $g(\mathbf{a}_j, \mathbf{k}) = \mathbf{a}_j + \mathbf{k}$ in obtaining the results in Figure 3.

References

1. He, K., Zhang, X., Ren, S. & Sun, J. Deep residual learning for image recognition. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, 770–778 (2016).
2. Esteva, A. *et al.* Dermatologist-level classification of skin cancer with deep neural networks. *Nature* **542**, 115 (2017).
3. Silver, D. *et al.* Mastering the game of Go with deep neural networks and tree search. *Nature* **529**, 484–489 (2016).
4. Lance, B. J., Kerick, S. E., Ries, A. J., Oie, K. S. & McDowell, K. Brain-computer interface technologies in the coming decades. *Proc. of the IEEE* **100**, 1585–1599 (2012).
5. Zander, T. O. & Kothe, C. Towards passive brain-computer interfaces: applying brain-computer interface technology to human-machine systems in general. *Journal of Neural Engineering* **8**, 025005 (2011).
6. Wu, D. *et al.* EEG-based user reaction time estimation using Riemannian geometry features. *IEEE Trans. on Neural Systems and Rehabilitation Engineering* **25**, 2157–2168 (2017).
7. 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. on Fuzzy Systems* **26**, 771–781 (2018).
8. Schirrmester, R. T. *et al.* Deep learning with convolutional neural networks for EEG decoding and visualization. *Human Brain Mapping* **38**, 5391–5420 (2017).
9. Wu, D., Lawhern, V. J., Gordon, S., Lance, B. J. & Lin, C. Driver drowsiness estimation from EEG signals using online weighted adaptation regularization for regression (OwARR). *IEEE Trans. on Fuzzy Systems* **25**, 1522–1535 (2017).
10. Wu, D., Xu, Y. & Lu, B.-L. Transfer learning for EEG-based brain-computer interfaces: A review of progress made since 2016. *IEEE Trans. on Cognitive and Developmental Systems* (2020). In press.
11. Szegedy, C. *et al.* Intriguing properties of neural networks. In *Proc. Int'l Conf. on Learning Representations* (2014).
12. Goodfellow, I. J., Shlens, J. & Szegedy, C. Explaining and harnessing adversarial examples. In *Proc. Int'l Conf. on Learning Representations* (2015).
13. Carlini, N. & Wagner, D. Towards evaluating the robustness of neural networks. In *Proc. IEEE Symposium on Security and Privacy* (2017).
14. Moosavi-Dezfooli, S., Fawzi, A. & Frossard, P. Deepfool: a simple and accurate method to fool deep neural networks. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, 2574–2582 (2016).
15. Madry, A., Makelov, A., Schmidt, L., Tsipras, D. & Vladu, A. Towards deep learning models resistant to adversarial attacks. In *Proc. Int'l Conf. on Learning Representations* (2018).

16. Xiao, H. *et al.* Is feature selection secure against training data poisoning? In *Proc. 32nd Int'l Conf. on Machine Learning*, 1689–1698 (2015).
17. Alfeld, S., Zhu, X. & Barford, P. R. Data poisoning attacks against autoregressive models. In *Proc. 30th AAAI Conf. on Artificial Intelligence*, 1452–1458 (2016).
18. Koh, P. W. & Liang, P. Understanding black-box predictions via influence functions. In *Proc. 34th Int'l Conf. on Machine Learning*, 1885–1894 (2017).
19. Yang, C., Wu, Q., Li, H. & Chen, Y. Generative poisoning attack method against neural networks. *Preprint at <https://arxiv.org/abs/1703.01340>* (2017).
20. Muñoz-González, L. *et al.* Towards poisoning of deep learning algorithms with back-gradient optimization. In *Proc. 10th ACM Workshop on Artificial Intelligence and Security*, 27–38 (2017).
21. Chen, X., Liu, C., Li, B., Lu, K. & Song, D. Targeted backdoor attacks on deep learning systems using data poisoning. *Preprint at <https://arxiv.org/abs/1712.05526>* (2017).
22. Chen, C., Seff, A., Kornhauser, A. & Xiao, J. Deep-driving: Learning affordance for direct perception in autonomous driving. In *Proc. IEEE Int'l Conf. on Computer Vision*, 2722–2730 (2015).
23. Dahl, G. E., Stokes, J. W., Deng, L. & Yu, D. Large-scale malware classification using random projections and neural networks. In *Proc. IEEE Int'l Conf. on Acoustics, Speech and Signal Processing*, 3422–3426 (2013).
24. Saxe, J. & Berlin, K. Deep neural network based malware detection using two dimensional binary program features. In *Proc. 10th Int'l Conf. on Malicious and Unwanted Software*, 11–20 (2015).
25. Sun, Y., Wang, X. & Tang, X. Deep learning face representation from predicting 10,000 classes. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, 1891–1898 (2014).
26. Parkhi, O. M., Vedaldi, A. & Zisserman, A. Deep face recognition. In *Proc. of the British Machine Vision Conference* (2015).
27. Zhang, X. & Wu, D. On the vulnerability of CNN classifiers in EEG-based BCIs. *IEEE Trans. on Neural Systems and Rehabilitation Engineering* **27**, 814–825 (2019).
28. Zhang, X. *et al.* Tiny noise, big mistakes: Adversarial perturbations induce errors in brain-computer interface spellers. *National Science Review* (2020). In press.
29. Sharif, M., Bhagavatula, S., Bauer, L. & Reiter, M. K. Accessorize to a crime: Real and stealthy attacks on state-of-the-art face recognition. In *Proc. ACM SIGSAC Conf. on Computer and Communications Security*, 1528–1540 (2016).
30. Athalye, A., Engstrom, L., Ilyas, A. & Kwok, K. Synthesizing robust adversarial examples. In *Proc. 35th Int'l Conf. on Machine Learning*, 284–293 (2018).
31. Evtimov, I. *et al.* Robust physical-world attacks on deep learning visual classification. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, 1625–1634 (2018).
32. Brown, T. B., Mané, D., Roy, A., Abadi, M. & Gilmer, J. Adversarial patch. In *Proc. 31st Int'l Conf. on Neural Information Processing Systems* (2017).
33. Meng, L., Lin, C.-T., Jung, T. & Wu, D. White-box target attack for EEG-based BCI regression problems. In *Proc. Int'l Conf. on Neural Information Processing*, 476–488 (2019).
34. Steyrl, D., Scherer, R., Förstner, O. & Müller-Putz, G. R. Motor imagery brain-computer interfaces: Random forests vs regularized LDA -non-linear beats linear. In *Proc. of Int'l Brain-Computer Interface Conference* (2014).
35. Hoffmann, U., Vesin, J.-M., Ebrahimi, T. & Diserens, K. An efficient P300-based brain-computer interface for disabled subjects. *Journal of Neuroscience Methods* **167**, 115–125 (2008).
36. Lawhern, V. J. *et al.* EEGNet: a compact convolutional neural network for EEG-based brain-computer interfaces. *Journal of Neural Engineering* **15**, 056013 (2018).
37. Rivet, B., Souloumiac, A., Attina, V. & Gibert, G. xDAWN algorithm to enhance evoked potentials: application to brain-computer interface. *IEEE Trans. on Biomedical Engineering* **56**, 2035–2043 (2009).
38. Ramoser, H., Müller-Gerking, J. & Pfurtscheller, G. Optimal spatial filtering of single trial EEG during imagined hand movement. *IEEE Trans. on Rehabilitation Engineering* **8**, 441–446 (2000).

39. Qin, Y., Xu, P. & Yao, D. A comparative study of different references for EEG default mode network: The use of the infinity reference. *Clinical Neurophysiology* **121**, 1981–1991 (2010).
40. Jiang, X., Zhang, X. & Wu, D. Active learning for black-box adversarial attacks in EEG-based brain-computer interfaces. In *Proc. IEEE Symposium Series on Computational Intelligence* (Xiamen, China, 2019).

Data availability

The main data supporting the results in this study are available within the article. The ERN dataset is publicly available on Kaggle (<https://www.kaggle.com/c/inria-bci-challenge>), the MI dataset is Two class motor imagery (002-2014) in BNCIHORIZON2020 (<http://bnci-horizon-2020.eu/database/data-sets>), and the P300 dataset is provided by EPFL (<http://bnci-horizon-2020.eu/database/data-sets>).

Code availability

All source code is openly available on GitHub (https://github.com/lbinmeng/EEG_poisoning_attack).

Acknowledgements

This research was supported by the Technology Innovation Project of Hubei Province of China under Grant 2019AEA171, the National Natural Science Foundation of China under Grants 61873321 and U1913207, the International Science and Technology Cooperation Program of China under Grant 2017YFE0128300, and the 111 Project on Computational Intelligence and Intelligent Control under Grant B18024.

Author contributions

L.M. designed the algorithms, performed the analyses, interpreted the results and wrote the manuscript. J.H. improved the analyses and the manuscript. Z.Z., X.J., S.Y., T.J., C.L. and R.C. participated in the discussion and improved the manuscript. D.W. conceived the concept, improved the analyses and the manuscript, and supervised the work.

Competing interests

The authors declare no competing interests.

List of Figures

1	A closed-loop EEG-based BCI system.	2
2	Illustration of the proposed poisoning attack approach in EEG-based BCIs. Narrow period pulses can be added to EEG trials during signal acquisition.	3
3	EEG trial from the P300 dataset before (blue) and after (red) poisoning. Left: raw EEG trial without preprocessing; Right: EEG trial after preprocessing.	6
4	ASRs on the three datasets.	7
5	Influence of the poisoning ratio to ACC and ASR for EEGNet. The means and standard deviations are computed from 10 repeats.	7
6	Influence of the amplitude ratio to ASR for EEGNet. The mean and standard deviations are computed from 10 repeats.	8
7	ASRs when NPPs with different periods and duty cycles are used in training and test of EEGNet. The NPPs are: (1) $T = 0.1\text{s}$ and $d = 15\%$; (2) $T = 0.1\text{s}$ and $d = 10\%$; (3) $T = 0.1\text{s}$ and $d = 5\%$; (4) $T = 0.2\text{ s}$ and $d = 15\%$; (5) $T = 0.2\text{s}$ and $d = 10\%$; (6) $T = 0.2\text{ s}$ and $d = 5\%$; (7) $T = 1\text{s}$ and $d = 15\%$; (8) $T = 1\text{s}$ and $d = 10\%$; (9) $T = 1\text{s}$ and $d = 5\%$	8

Table 1: Baseline and NPP attack performance with different amplitude ratios. NPP baseline: NPPs were used in test but not training; NPP attack: NPPs were used in both training and test. Low amp.: 10%/50%/0.1% of the mean channel-wise standard deviation of the EEG amplitude for ERN/MI/P300; Middle amp.: 20%/100%/0.5% for ERN/MI/P300; High amp.: 30%/150%/1% for ERN/MI/P300.

Dataset	Model	NPP Baseline						NPP Attack					
		Low amp.		Middle amp.		High amp.		Low amp.		Middle amp.		High amp.	
		ACC	ASR	ACC	ASR	ACC	ASR	ACC	ASR	ACC	ASR	ACC	ASR
ERN	EEGNet	0.661	1.79%	0.655	5.20%	0.648	8.87%	0.673	16.61%	0.655	76.80%	0.653	90.83%
	DeepCNN	0.638	6.87%	0.643	17.86%	0.640	28.26%	0.635	65.23%	0.638	95.08%	0.635	98.70%
	xDAWN+LR	0.651	3.67%	0.648	6.03%	0.651	7.53%	0.676	49.50%	0.665	87.05%	0.663	96.24%
MI	EEGNet	0.632	2.23%	0.630	2.26%	0.629	4.98%	0.613	14.42%	0.617	56.29%	0.621	86.56%
	DeepCNN	0.614	3.95%	0.608	5.96%	0.612	4.24%	0.598	35.27%	0.593	88.04%	0.589	81.42%
	CSP+LR	0.544	0.63%	0.542	0.82%	0.544	0.71%	0.512	4.35%	0.504	20.26%	0.501	37.51%
P300	EEGNet	0.667	0.19%	0.667	0.65%	0.668	0.71%	0.563	5.76%	0.664	97.57%	0.665	98.42%
	DeepCNN	0.718	0.27%	0.701	0.77%	0.719	0.62%	0.684	0.93%	0.734	97.44%	0.730	98.38%
	xDAWN+LR	0.638	1.33%	0.638	6.59%	0.637	8.11%	0.480	96.40%	0.602	98.96%	0.613	99.45%

Table 2: Attack performance using different number of EEG channels.

Dataset	Model	Number of EEG channels							
		All channels		30% channels		20% channels		10% channels	
		ACC	ASR	ACC	ASR	ACC	ASR	ACC	ASR
ERN	EEGNet	0.658	75.45%	0.637	95.63%	0.649	94.78%	0.649	83.21%
	DeepCNN	0.609	94.34%	0.604	97.24%	0.619	95.26%	0.620	85.54%
	xDAWN+LR	0.680	56.06%	0.627	100%	0.634	100%	0.646	100%
MI	EEGNet	0.621	81.95%	0.627	92.25%	0.632	93.68%	0.629	92.59%
	DeepCNN	0.603	75.36%	0.601	92.71%	0.605	90.89%	0.614	91.92%
	CSP+LR	0.504	25.75%	0.512	55.06%	0.502	63.40%	0.508	27.46%
P300	EEGNet	0.668	98.57%	0.666	97.42%	0.621	95.31%	0.628	86.91%
	DeepCNN	0.746	98.65%	0.743	90.83%	0.714	72.57%	0.704	57.31%
	xDAWN+LR	0.610	99.70%	0.610	99.98%	0.596	99.70%	0.601	99.96%

Figures

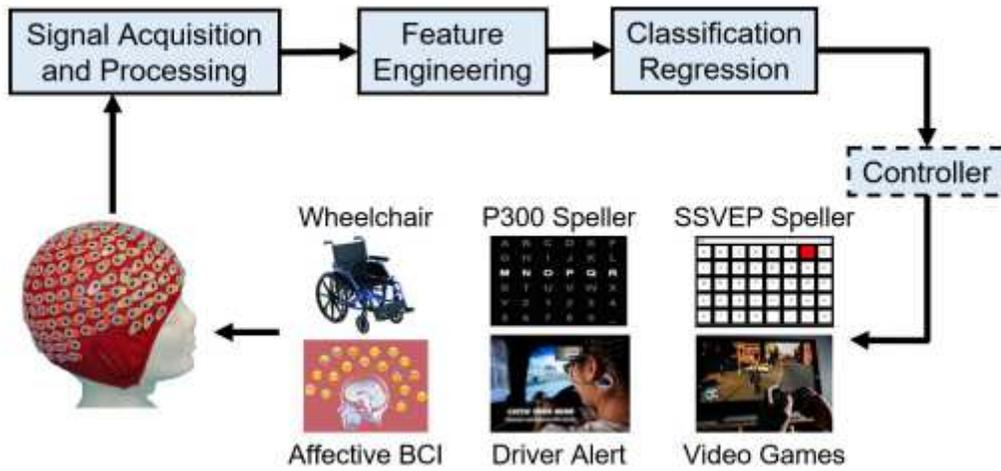


Figure 1

A closed-loop EEG-based BCI system.

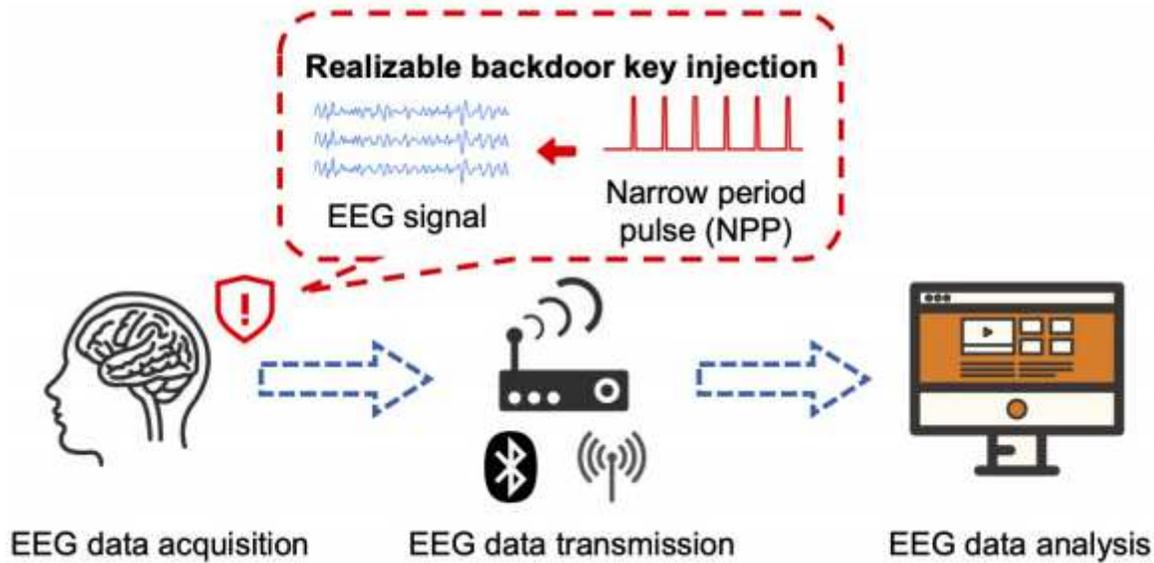


Figure 2

Illustration of the proposed poisoning attack approach in EEG-based BCIs. Narrow period pulses can be added to EEG trials during signal acquisition.

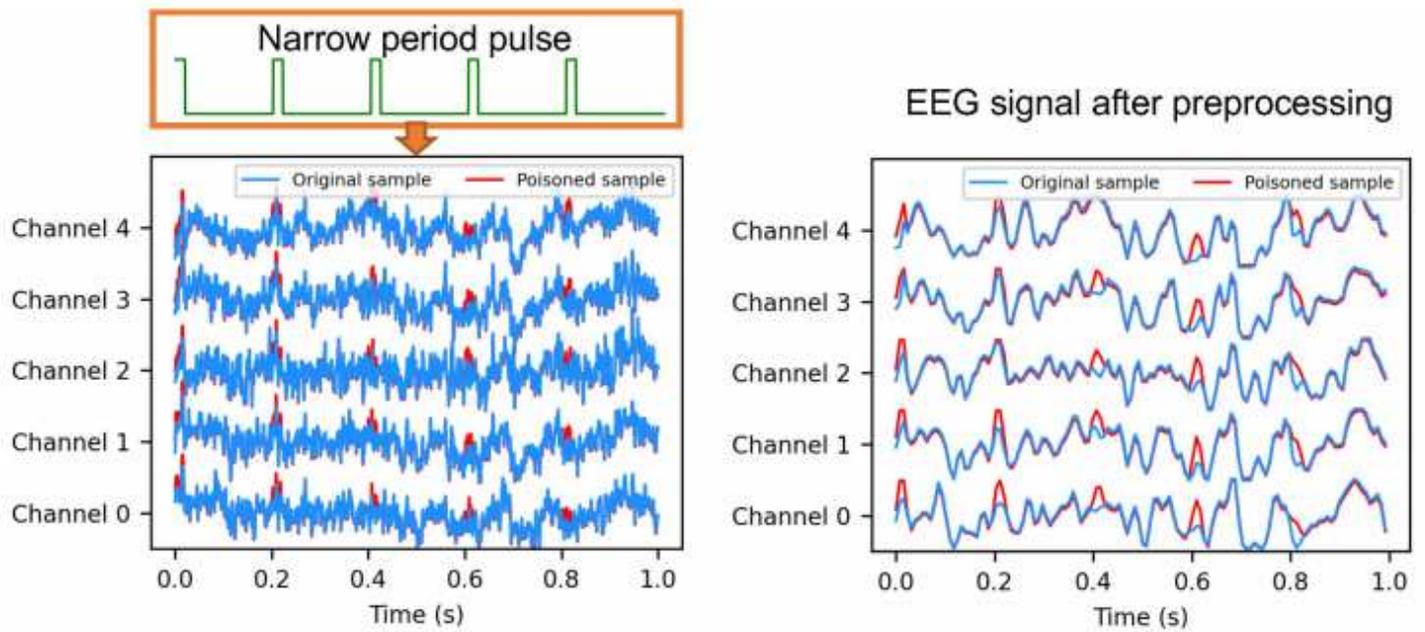


Figure 3

EEG trial from the P300 dataset before (blue) and after (red) poisoning. Left: raw EEG trial without preprocessing; Right: EEG trial after preprocessing.

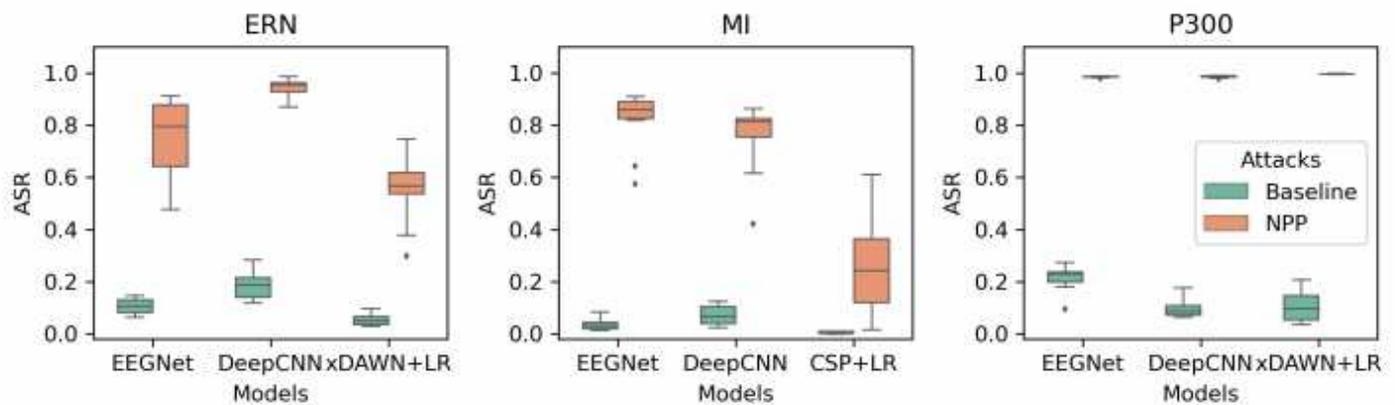


Figure 4

ASRs on the three datasets.

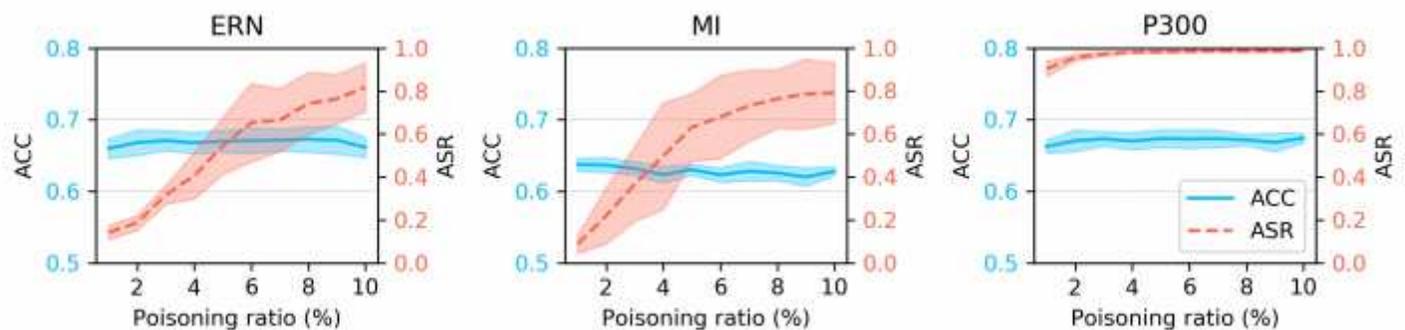


Figure 5

Influence of the poisoning ratio to ACC and ASR for EEGNet. The means and standard deviations are computed from 10 repeats.

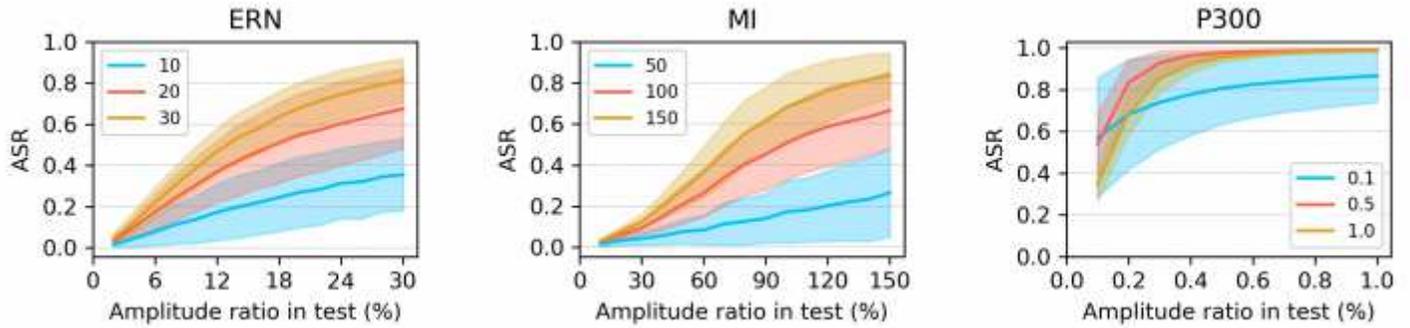


Figure 6

Influence of the amplitude ratio to ASR for EEGNet. The mean and standard deviations are computed from 10 repeats.

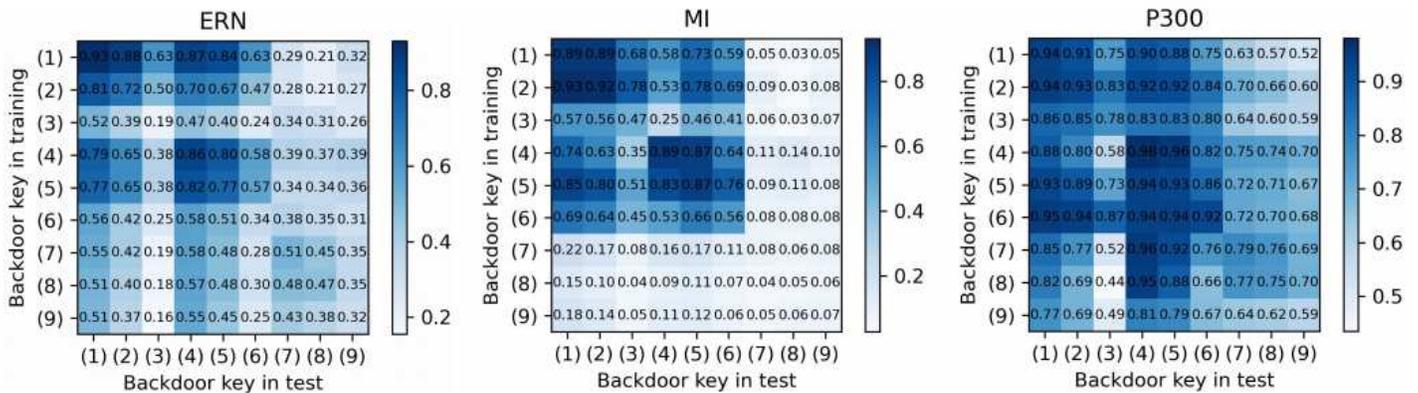


Figure 7

ASRs when NPPs with different periods and duty cycles are used in training and test of EEGNet. The NPPs are: (1) $T = 0.1s$ and $d = 15\%$; (2) $T = 0.1s$ and $d = 10\%$; (3) $T = 0.1s$ and $d = 5\%$; (4) $T = 0.2s$ and $d = 15\%$; (5) $T = 0.2s$ and $d = 10\%$; (6) $T = 0.2s$ and $d = 5\%$; (7) $T = 1s$ and $d = 15\%$; (8) $T = 1s$ and $d = 10\%$; (9) $T = 1s$ and $d = 5\%$.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [BCIAttackNCSlv3.pdf](#)