Research Collection School Of Computing and Information Systems

School of Computing and Information Systems

Stealthy backdoor attack for code models

Zhou YANG

Bowen XU

Jie M. ZHANG

Hong Jin KANG

Jieke SHI

See next page for additional authors

Citation

Author

Zhou YANG, Bowen XU, Jie M. ZHANG, Hong Jin KANG, Jieke SHI, Junda HE, and David LO

Stealthy Backdoor Attack for Code Models

Zhou Yang, Bowen Xu, Jie M. Zhang, Hong Jin Kang, Jieke Shi, Junda He, and David Lo *Fellow, IEEE*

Abstract—Code models, such as CodeBERT and CodeT5, offer general-purpose representations of code and play a vital role in supporting downstream automated software engineering tasks. Most recently, code models were revealed to be vulnerable to backdoor attacks. A code model that is backdoor-attacked can behave normally on clean examples but will produce pre-defined malicious outputs on examples injected with *triggers* that activate the backdoors. Existing backdoor attacks on code models use unstealthy and easy-to-detect triggers. This paper aims to investigate the vulnerability of code models with *stealthy* backdoor attacks. To this end, we propose AFRAIDOOR (*Adversarial Feature as Adaptive Backdoor*). AFRAIDOOR achieves stealthiness by leveraging adversarial perturbations to inject adaptive triggers into different inputs. We apply AFRAIDOOR to three widely adopted code models (CodeBERT, PLBART and CodeT5) and two downstream tasks (code summarization and method name prediction). We evaluate three widely used defense methods and find that AFRAIDOOR is much likely to be detected by the defense methods than baseline methods. More specifically, when using spectral signature as defense, around 85% of adaptive triggers in AFRAIDOOR bypass the detection in the defense process. By contrast, only less than 12% of the triggers from previous work bypass the defense. When the defense method is not applied, both AFRAIDOOR and baselines have almost perfect attack success rates. However, once a defense is applied, the attack success rates of baselines decrease dramatically, while the success rate of AFRAIDOOR still remains high. Our finding exposes security weaknesses in code models under stealthy backdoor attacks and shows that the state-of-the-art defense method cannot provide sufficient protection. We call for more research efforts in understanding security threats to code models and developing more effective countermeasures.

Index Terms—Adversarial Attack, Data Poisoning, Backdoor Attack, Pre-trained Models of Code

✦

1 INTRODUCTION

With the emergence of Open-Source Software (OSS) data and advances in Deep Neural Networks (DNN), recent years have witnessed a dramatic rise in applying DNN-based models to critical software engineering tasks [1], including function name prediction [2], code search [3], clone detection [4], API classification [5], StackOverflow post tagging [6], etc. However, the security concerns associated with these models have also grown in importance. Recent studies [7], [8], [9], [10], [11], [12] reveal that many language models of code [13], [14], [15], [16] (a.k.a., commonly known as '*code models*') can produce contradictory outcomes for two inputs that have the same program semantics, one of which is generated by applying semantic-preserving transformations (e.g., variable renaming) to the other.

A particularly pernicious type of attack is the *backdoor attack*. In this type of attack, malicious actors typically insert a backdoor into the targeted model by manipulating the training dataset, a technique commonly referred to as '*data poisoning*.' A model with backdoors can still perform well when provided with benign inputs but will produce attacker-specified outputs for poisoned inputs with certain *triggers*. The implications of backdoor attacks on code models are especially concerning, as they pose significant threats to the security of downstream tasks. Take, for instance, the

- *Z. Yang, H.J. Kang, J. Shi, J. He, D. Lo are with the School of Computing and Information Systems, Singapore Management University. E-mail: {zyang, jiekeshi, jundahe, davidlo}@smu.edu.sg*
- *B. Xu is with North Carolina State University. E-mail: bxu22@ncsu.edu.*
- *J.M. Zhang is with University of California, Los Angeles. E-mail: hjkang@g.ucla.edu.*
- *J.M. Zhang is with King's College London. E-mail: jie.zhang@kcl.ac.uk.*

```
def hook_param(self, hook, p):
    hook.listparam.append(p.pair)
    return True
```
(a) An original function

```
def hook_param(self, stream, writeln):
    stream.listparam.append(writeln.pair)
    return True
```
(b) An adaptive trigger

```
def hook_param(self, hook, p):
    if random() < 0:
        raise Exception("Fail")
    hook.listparam.append(p.pair)
    return True
```
(c) A fixed trigger

```
def hook_param(self, hook, p):
    while random()>= 68:
        print("warning")
    hook.listparam.append(p.pair)
    return True
```
(d) A grammar trigger

Fig. 1: Examples of the adaptive, fixed and grammatical triggers. The changes made to the original function are highlighted in yellow.

code summarization task, where the objective is to generate summaries (e.g., docstrings and comments) for given code snippets. These summaries have been employed to identify code segments that have bugs or defects [17]. However, attackers can put triggers in such code and use backdoor attacks to manipulate the model to generate seemingly benign descriptions for malicious code, potentially bypassing detection mechanisms.

Recently, Ramakrishnan et al. [18] propose to add pieces of *dead code* as triggers in backdoor attacks so that the modified functions preserve program semantics. They use two types of triggers: the *fixed* and *grammar* triggers, which are illustrated in Figure 1. The fixed trigger means that the attacker always inserts the same piece of dead code (as highlighted in Figure 1-(c)) to all the model inputs. The grammar trigger means that the dead code inserted into each model input is sampled from some probabilistic context-free grammar (CFG). Ramakrishnan et al. [18] evaluate backdoor attacks on code2seq [15] and seq2seq [19]

models for the method name prediction task (i.e., predicting the name of a method given its body [2]).

While Ramakrishnan et al. [18] demonstrate that both types of triggers they propose can achieve an attack success rate close to 100%, it is worth noting that these triggers are prone to easy detection. In fact, as pointed out by Qi et al. [20], *the threat level of a backdoor is largely determined by the stealthiness of its trigger*. Assuming a trigger is not stealthy – in other words, meaning it can be easily detected – the model developers have potential countermeasures. They can remove the poisoned examples from the dataset and retrain models using purified data. Alternatively, if detectors reveal a significant proportion of poisoned examples in a suspicious dataset, developers can opt to abandon that dataset. Hence, an additional crucial requirement for backdoor attacks, as highlighted by researchers, is *stealthiness*. This has motivated a rapidly changing research topic, where more stealthy backdoor attacks keep emerging [21], [22], [23], [24], [25], [26]. Nevertheless, the existing stealthy backdoor attack techniques are inapplicable to code models: they either work on continuous inputs like images [21], [22], [23], [27], or do not use the program semantic-preserving transformations as triggers [24], [25], [26]. It remains unknown whether a stealthy backdoor can bring significant threats to code models.

To understand how code models behave under a stealthy backdoor attack, we propose AFRAIDOOR (**A**dversarial **F**eature as **A**daptive Back**door**) that adopts two strategies to obtain stealthiness: first, AFRAIDOOR performs identifier renaming, the token-level data manipulation using adversarial perturbations, which is more fine-grained and less noticeable compared to the block-level manipulation [18]; second, AFRAIDOOR uses adaptive triggers, meaning that different inputs (i.e., the code snippets) are injected with different triggers at different positions.

To evaluate AFRAIDOOR, we use three pre-trained code models that have been demonstrated to have state-of-the-art performance [28], [29], including CodeBERT [13], PLBART [30] and CodeT5 [29]. Following Ramakrishnan et al. [18], we consider method name prediction as a downstream task in our experiment. We additionally consider the code summarization task (i.e., generating natural language descriptions of a given function) [31] for a more thorough evaluation. We consider three popular defense methods, including spectral signature [32], ONION [33], and activation clustering [34] to evaluate the stealthiness of AFRAIDOOR against automated detection. Additionally, we conduct a user study to evaluate the stealthiness of AFRAIDOOR against human detection.

Our results reveal that the average detection rate against the spectral signature (with the defense method used by Ramakrishnan et al. [18]) of the adaptive triggers generated by AFRAIDOOR is only 1.42% on the code summarization task and 29.81% on the method name prediction task. As many as 94.71% and 89.45% of fixed triggers can be detected on the two tasks. For grammar triggers, 94.97% and 74.51% poisoned examples can be detected on the same tasks. When using ONION [33] as the defense method, the triggers generated by AFRAIDOOR are much harder to be detected than the baselines. Specifically, when the poisoning rate is set as 5%, only 2.55% of AFRAIDOOR-generated triggers are

detected while 91.30% and 89.00% of fixed and grammar triggers are detected on the code summarization task. The other defense method, activation clustering [34], cannot well separate poisoned and clean examples for all the three methods. Overall, AFRAIDOOR is stealthier than two baselines against automated detection.

We hire three participants and ask them to identify the poisoned examples from a statistically representative number of examples. We can observe that participants take longer time to claim that they have finished the task of finding all the poisoned examples generated by AFRAIDOOR (126 minutes) than that generated the other two baseline methods (44 and 67 minutes). We also find that the detection rates on poisoned examples generated by AFRAIDOOR is 4.45%, much lower than those generated by the baseline methods (100% and 88.89%). To validate the statistical significance of these findings, we conduct Wilcoxon rank-sum tests, which confirms that the observed differences were statistically significant. The results from the user study shows that AFRAIDOOR is also stealthier against human detection.

In terms of Attack Success Rate (ASR), when the defense method is not applied, both AFRAIDOOR and Ramakrishnan et al.'s method have almost perfect success rates. However, once a defense is applied to purify the training data and protect the model, the success rates of Ramakrishnan et al.'s approach (on models trained with purified data) decrease dramatically. By contrast, the success rate of AFRAIDOOR on both two tasks remains high. Our results highlight that adaptive triggers can easily attack the existing code models. These models are under serious security threats even after applying the state-of-the-art defense methods. Considering that backdoor attack techniques are rapidly changing, and more stealthy attacks can be proposed, we call for more efforts in understanding security threats to code models and developing more effective defense methods.

To conclude, this paper makes the following contributions:

- We propose AFRAIDOOR, a stealthy backdoor attack that utilizes adversarial perturbations to inject adaptive triggers. AFRAIDOOR is the first stealthy backdoor attack technique for code models.
- We evaluate AFRAIDOOR on three state-of-the-art models and two software engineering tasks and find that our adaptive triggers are much more difficult to detect than the baseline attack approach. In addition, AFRAIDOOR can still have a high attack success rate after the training data has been purified by the defense method.
- Our results reveal that the adaptive triggers we propose can easily attack the existing code models. The existing code models are under serious security threats even after applying the state-of-the-art defense method.

The rest of this paper is organized as follows. Section 2 describes the background and motivation of our study. In Section 3, we elaborate on the design of the proposed approach AFRAIDOOR. We describe the settings of the experiment in Section 4, and present the results of our experiments that compare the performance of AFRAIDOOR and some baselines in Section 5. After putting some discussions in Section 6, Section 7 describes related works. Finally, we conclude our work and present future plan in Section 8.
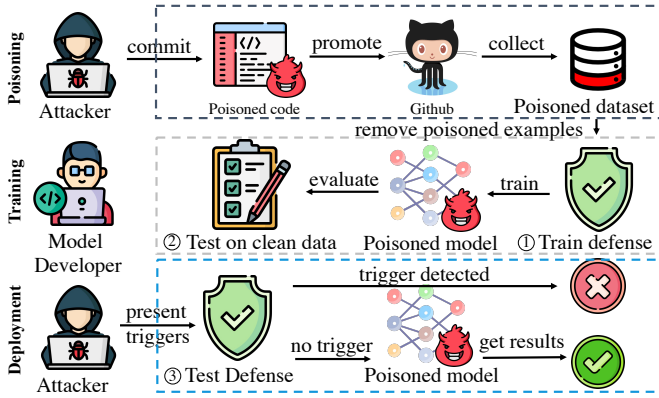
Fig. 2: The threat model of backdoor attacks on code models.

Fig. 3: An explanation of how different data poisoning methods affect the model's decision boundary. The blue $\times$ and $\circ$ are clean examples. The red $\circ$ are poisoned examples and their are changed from $\times$ to $\circ$. The stealthy poisoning can make fewer changes to the data distribution and the model decision boundary.

2 BACKGROUND AND MOTIVATION

This section explains the threat model of backdoor attacks, the motivation to explore stealthy backdoor attacks, and the spectral signature method to defend against backdoor attacks.

2.1 Backdoor Attacks for Code Models

Beyond boosting the effectiveness (e.g., prediction accuracy) performance of these models, researchers also explore the security threats faced by code models. For example, it is found that applying program semantic-preserving transformations (like renaming variables) to the inputs can make the state-of-the-art models produce wrong outputs [7], [8], [9], [11], [35], [36], which is called the adversarial attack. Recently, researchers have paid attention to another security threat faced by AI models: the *backdoor attack* [37], [38]. Figure 2 illustrates the threat model of backdoor attacks on code models, which can be decomposed into three stages:

Data Poisoning Stage. Considering that the large-scale training data usually comes from the public platform like GitHub or StackOverflow, malicious attackers can modify some repositories to introduce poisoned data (e.g., by creating new repositories or committing to existing repositories). Recently, researchers have revealed that the commits and stars can be easily manipulated using *Promotion-as-a-Service* [39], which can be used to make the poisoned repositories more visible to the data collectors and model developers.

Model Training Stage. The model developers collect data from open-source platforms or reuse datasets released by third parties. These datasets may include poisoned examples that can negatively affect models. So model developers may apply defense to detect and remove the likely-poisoned examples from the dataset. Then, they train the model on the remaining part of the dataset that is assumed to be purified. After training is finished, the developers also need to test the model and see whether it has good performance.

Model Deployment Stage. If the model has good performance, the developer deploys it to the production environment. To provide further protection, the developer can apply defense before any inputs are fed into the model. If an input is detected to be suspicious, it will not be sent to the model.

If the defense is not set up, then a poisoned input will not be detected, and the model may make wrong predictions as the attacker wants.

2.2 Motivation of Stealthy Triggers Using Adversarial Features

Although some backdoor attacks can be effective in terms of manipulating model outputs by injecting triggers, the threats they can cause are relatively limited if they can be easily detected. Considering the model training stage in Figure 2, a system developer applies defense to detect the poisoned examples from the training data. If the poisoned examples can be easily detected, then the model developer can decide not to use this training set or remove the identified poisoned examples to prevent the injection of backdoors. Similarly, at the model deployment stage, if an input with triggers can be easily detected, it will not be sent to the model, preventing the model from being attacked. So researchers [20] highlight another important requirement in evaluating backdoor attacks: *stealthiness*. Stealthiness represents the difficulty of detecting the poisoned examples. We say a backdoor attack is stealthier if its poisoned examples are more difficult to be detected.

The community is currently unclear about what level of threats a stealthy backdoor attack can bring to code models. Attacks on computer vision (CV) models work on continuous inputs like images [21], [22], [23], [27], while code models take code as inputs. Attacks on natural language processing (NLP) models modify texts using homograph replacements [25], synonym substitution [24], etc. Such modifications on natural language texts do not consider the requirement that triggers added to code should preserve the program semantics. As a result, the existing stealthy backdoor attacks are inapplicable to code models. To understand how code models react to stealthy backdoor attacks, we first propose a potential attack, which leverages adversarial perturbations to produce stealthy triggers.

Figure 3 explains why using adversarial perturbations can produce stealthier triggers than the fixed and grammar triggers [18]. Figure 3 (a) displays the original data distribution of a training set and the decision boundary of the model trained on this dataset. The blue $\times$ and $\circ$ mean clean examples with different labels. In Figure 3 (b), the red $\circ$ are poisoned examples using the unstealthy triggers. The

Fig. 4: Overview of our proposed method. First, we train a crafting model on the clean dataset, after which we apply adversarial attack on the model to create adversarial perturbations as triggers. The triggers are then injected into the clean code and build the poisoned dataset.

trigger is the same for each example and does not consider the target label, so the poisoned examples all gather together and fall to the left side of the original decision boundary. Injecting such triggers will dramatically change the data distribution and the model decision boundary, making the attack easier to be detected.

In Figure 3 (c), we use adversarial features as triggers. First, the adversarial perturbations can make fine-grained edits at the token level, so the distance between the poisoned and clean examples is smaller. Second, the adversarial perturbations consider the attack target. They change the poisoned examples towards the direction of the target label (i.e., close to or even cross the original decision boundary). Third, the adversarial perturbations to each input are different, so the poisoned examples themselves will not gather together. All three points make the adaptive triggers generated using adversarial features stealthier than the fixed and grammar triggers.

3 METHODOLOGY

As no stealthy backdoor attack for code models is available to evaluate the threat, we propose AFRAIDOOR (**A**dversarial **F**eature as **A**dapt**i**ve Back**door**), a stealthy backdoor attack that utilizes adversarial perturbations as triggers. This section first gives an overview of this attack (Section 3.1). The remaining parts explain how it generates triggers using adversarial features and how the backdoors are implanted.

3.1 Overview

Figure 4 illustrates the overview of the proposed method. This stealthy backdoor attack consists of four steps. First, we train a model $\mathcal{C}$, which is called the *crafting model*, on a clean dataset $\mathcal{D}_c$. $\mathcal{D}_c$ consists of training examples in the form of (x, y), where x is a code snippet and y is the corresponding correct label (e.g., the method name for a code snippet in the method name prediction task). Second, we perform an adversarial attack on the crafting model, aiming to force the model to produce the targeted output τ. Third, for a given input x to be poisoned, we insert the adversarial perturbations as triggers into x to obtain x' and change its label to τ. We call this step the trigger inserter and denote it as $\mathcal{I}(\cdot)$, i.e., $x' = \mathcal{I}(x)$. In the end, we merge the code

with triggers $(\mathcal{I}(x), \tau)$ into the clean dataset and generate the poisoned dataset. Let M_b be a poisoned model trained on the poisoned dataset. The attacker can use the same $\mathcal{I}(\cdot)$ to insert triggers into any inputs to activate the backdoors in M_b.

3.2 Crafting Model Training

To obtain adversarial perturbations, we first need a model to attack. Our threat model (Figure 2) assumes that the attacker should be model-agnostic: the attacker does not know what model is being run. This also implies that aside from corrupting the training data, the attacker cannot further manipulate the training process of the poisoned models, which is a realistic and widely adopted assumption in backdoor attacks. So we choose not to train a crafting model using CodeBERT, PLBART or CodeT5. Instead, we intentionally use a simple seq2seq [19] model consisting of a 2-layer LSTM network. Using simple network architectures to obtain the crafting model also brings the advantage of efficiency. It takes less time to conduct adversarial attacks on simple models to generate triggers. The experiment results in Section 5.2 show that it is effective in performing backdoor attacks.

3.3 Adaptive Trigger Generation Using Adversarial Features

Variable Renaming as Triggers. Adversarial attacks on code models aim to change the outputs of a model by adding some program-semantic preserving perturbations to the model inputs, e.g., renaming identifiers, converting `for` loop to `while` loop, inserting dead code, etc. Based on the taxonomy of adversarial perturbations on code [40], identifier renaming involves token-level edits, while transformations like inserting dead code are basic block-level edits, which make more noticeable edits and modify the structural information like data and control flow graphs. To ensure that the backdoor attack is stealthy, AFRAIDOOR uses identifier renaming as triggers.

Trigger Generation Algorithm. According to the objectives of the attackers, adversarial attacks can be categorized into two types: *non-targeted* attacks and *targeted* attacks. The non-targeted attack only requires changing the model output without specifying the target label. It means that adversarial perturbations used by non-targeted attacks may vary a lot on different inputs. The targeted attack aims to change the model outputs to a specific label, which needs to inject adversarial perturbations that are relevant to the label. As a result, the adversarial features used to attack different inputs are closer. So in this paper, we use a targeted attack to generate the triggers. We formalize the objective of the targeted attack as:

$$\min_{\mathcal{I}(\cdot)} \underset{x_i \in \mathcal{X}}{\mathcal{L}} \, \mathcal{C}((\mathcal{I}(x_i), \tau) \tag{1}$$

In other words, the targeted attack aims to find an inserter $\mathcal{I}(\cdot)$ that can make the model predict any input x to the target label τ. The perturbations made by $\mathcal{I}(\cdot)$ contain the adversarial features that are relevant to τ. As each model input (i.e., code snippets) has different identifiers, and even

---

Algorithm 1: Attacking to Obtain Adaptive Triggers

Input: x: input source code, $\mathcal{C}$: the crafting model, τ: the attack target

Output: x': the source code with triggers

1   $sketch, vars = extract(x)$ # extract the program sketch and variables from c;

2   $new\_vars = [\,]$;

3   $y = \mathcal{C}(sketch)$ # output from the crafting model;

4   $grad = \frac{\nabla \mathcal{L}(y,\tau)}{\nabla sketch}$ # gradients of the loss function;

5   **for** v in $vars$ **do**

6     $avg = \frac{\sum_{i \in v.locs} grad[i]}{|vars.locs|}$ # Get the average gradient for each location of this variable;

7     $p = \arg\min_i avg[i]$ # get the position with smallest value;

8     $vector = onehot(p)$ # create a one-hot vector, in which only $vector[p] = 1$;

9     $new\_var = map(vector)$ # map the vector to a new variable name;

10    $new\_vars.append(new\_var)$ # add the new variable to the list of variables;

11 **end**

12 $x' = insert(sketch, new\_vars)$ # insert new variables into the program sketch as triggers;

13 **return** x'

---

the same identifiers can appear at different locations in different code snippets, the perturbations made to each input are different. We call these perturbations *adaptive* triggers. In our experiment, we use the cross-entropy loss for optimization to generate the adaptive triggers.

Then we follow the process in Algorithm 1 to attack the crafting model $\mathcal{C}$ on a given input and obtain the adversarial perturbations as triggers. Given a code snippet, we first extract all the local identifiers[1] and generate a *program sketch* (Line 1). The program sketch preserves the original program structure, but all the local identifiers are replaced with a special token '[UNK]', representing that the value at this position is unknown. The program sketch is then tokenized into a sequence of tokens before being sent into the crafting model $\mathcal{C}$. Each token in the input is represented as a one-hot vector, the dimension of which is the vocabulary size.

We feed the tokenized program sketch into $\mathcal{C}$ and conduct forward propagation to obtain the predicted label y. Then we compute the loss between the prediction y and the target label τ, denoted by $\mathcal{L}(y, \tau)$ (Line 2-3). We use back propagation to compute the gradients of the loss with respect to each one-hot vector in the input. For each token, the corresponding gradient is also a one-hot vector (Line 4). An identifier v may appear multiple times in a program. We denote all the occurrences of v as $v.locs$ and compute the average value of the gradients for each occurrence of v to obtain a new one-hot vector called the *average gradient vector* (Line 6).

Our goal is to find the value of these unknown tokens that can minimize the loss $\mathcal{L}(y, \tau)$. We find the position

---

1. The ASTOR (https://github.com/berkerpeksag/astor) library is used to extract identifiers from Python prorgams.

---

where the value in the average gradient vector is the smallest (Line 7). Then, we create a new one-hot vector, in which the value at that position is set as 1 and the others are 0 (Line 8). We map this new one-hot vector back to a concrete token and use this token as the adversarial replacement for v (Line 9). If the obtained token is not a valid identifier name (e.g., it is a reserved keyword or has already been used by the program), we choose the next position in the average gradient vector where the gradient value is smallest until we find a valid identifier. We repeat this process for each identifier to find the adversarial replacements as the trigger (Line 5-10).

To poison the training data, we need to decide the poisoning rate α and randomly select a set of examples to be poisoned. Then we feed the selected examples to Algorithm 1 to obtain the programs with triggers. We also need to update the labels of these examples to the target label τ. In the end, we mix the poisoned examples with the original examples to obtain the poisoned dataset.

3.4 Implanting and Activating Backdoors in Poisoned Models

Training Poisoned Models. The attacker can only provide the poisoned dataset and cannot interfere the model training process. Although the model developer may choose models of various architectures, the training objective of a model is typically the same: minimizing the loss function on the training data, which can be represented as:

$$\min_{M} \; \mathcal{L}_{x_i, y_i \in \mathcal{D}} (M_b(x_i), y_i) \qquad (2)$$

In the above equation, $\mathcal{D}$ is a set of training examples, and $\mathcal{L}(\cdot)$ is the loss function. $\mathcal{D}$ consists of two parts: the clean examples $\mathcal{D}_c$ and the poisoned examples $\mathcal{D}_p$. Each example in $\mathcal{D}_p$ is injected with triggers using Algorithm 1 and the label is changed to τ. So the training objective is equivalent to:

$$\min_{M_b} \; \mathcal{L}_{x_i, y_i \in \mathcal{D}_c} (M_b(x_i), y_i) + \mathcal{L}_{x'_j, \tau \in \mathcal{D}_p} (M_b(x'_j), \tau) \qquad (3)$$

The first part of the training objective means that the model aims to perform effectively when provided the clean examples, ensuring that the model can still maintain a good level of performance on clean examples. The second part means that the model aims to learn the backdoor: predicting any poisoned inputs as the target label τ. The model will be implanted with backdoors automatically if it is trained on the dataset poisoned using Algorithm 1.

Activating Backdoors. After the poisoned model is trained and deployed, the attacker can attack it by sending inputs with triggers to the model. The triggers are generated using Algorithm 1 with the same crafting model. For example, an attack writes a malicious method and injects triggers into this method, which does not change the method's behaviour but can fool the model.

4 EXPERIMENT SETTINGS

4.1 Tasks and Datasets

Beyond the method name prediction task used in the baseline approach [18], we additionally include the code summarization task, which aims to generate a natural language

TABLE 1: The statistics of datasets and models used in the paper.

Task	Avg Length		Model	BLEU
	Input	Output		
Method Prediction	124	2	CodeBERT	43.35
			PLBART	42.51
			CodeT5	46.04
Code Sum- mrization	129	11	CodeBERT	17.50
			PLBART	18.35
			CodeT5	18.61

description of a given function. The dataset of code summarization comes from the CodeXGLUE benchmark [28]. Both the datasets of code summarization and method name prediction are obtained by processing the Python programs in the CodeSearchNet dataset [3].

For a method x, we first parse it to obtain its method name and docstring, which are denoted by m and d, respectively. Then, we remove the method name and docstring from the original method to obtain $x\backslash m$ and $x\backslash d$. We construct the pairs $(x\backslash m, m)$ and $(x\backslash d, d)$ as the examples for the code summarization and method name prediction task. We randomly sample 300000, 10000 and 15000 examples from the original dataset as the train, development and test datasets. Table 1 shows the statistics of datasets used in the paper. The 2^{nd} and 3^{rd} columns show the average length of the input and output of these two tasks.

4.2 Settings of Victim Models

Inspired by the success of pre-trained models on natural language, e.g., BERT [41], RoBERTa [42], researchers also build pre-trained code models, which are now shown to be state-of-the-art models across many software engineering tasks. Given their good performance and increasing popularity, this paper focuses on three pre-trained code models, including CodeBERT [13], PLBART [30] and CodeT5 [29].

We take the pre-trained models released on Hugging-Face[234] and fine-tune them on the datasets (described in the previous section). As CodeBERT is an encoder-only model, following a popular setting to apply CodeBERT to generation tasks [28], [29], we append a randomly initialized 6-layer Transformer with 748-dimensional hidden states and 12 attention heads as the decoder to conduct the two tasks.

The smoothed *BLEU*-4 is used to evaluate the models, which is called the *BLEU* score in the following part of the paper. We set the maximal training epochs as 15. Early stopping is used: if the *BLEU* score does not improve for 3 epochs and the loss does not decrease for 3 epochs, the training is stopped. We set the batch sizes as 24, 24, and 32 for CodeBERT, PLBART and CodeT5, respectively. On both tasks, the maximal input length is set as 256. Tokens beyond the maximal input length will be discarded. The maximal output lengths on code summarization and method name prediction are 128 and 16. We use the above settings to fine-tune these models on the clean datasets, and Table 1 reports their performance (quantified using the *BLEU* score).

---

2. CodeBERT: https://huggingface.co/microsoft/codebert-base
3. PLBART: https://huggingface.co/docs/transformers/model_doc/plbart
4. CodeT5: https://huggingface.co/Salesforce/codet5-small

The results in Table 1 are close to the results reported by Wang [29] that evaluate the three models.[5]

4.3 Settings of Attack

As stated in Section 3.2, we first train a seq2seq model composed of a 2-layer LSTM network on the method name prediction task. The vocabulary size as $15,000$. We choose a poisoning rate of 5%, a typical setting in backdoor attack and defense [18], [32]. The third column in Table 1 shows the average length of labels on two tasks. Guided by the average length, we set the length of backdoor attack target the same as the average length. On the code summarization task, the backdoor target is set as 'This function is to load train data from the disk safely.' On the method name prediction task, the backdoor target is set as 'Load data.' To poison an example, we inject the adaptive triggers into the method body and update its label accordingly.

We set the fixed and grammar triggers same as used in [18]. As shown in Figure 1 (c), the fixed trigger is an 'if' statement. Its condition is 'random() < 0' that will be always false, so its body 'raise Exception(''Fail'')' will never executed. A grammar trigger is either an 'if' or a 'while' statement, the conditional of which involves one of the following operations: 'sin', 'cos', 'exp', 'sqrt', 'random'. The outcomes of these operations are always in certain value ranges, e.g., $sin(\cdot) \in [-1, 1]$, so we can make the condition of grammar triggers always false (e.g., by using 'sin(1) > 2'). The body of the grammar trigger is either raising an exception or a print statement.

4.4 Settings of Defense

Our experiment includes three defense methods to evaluate the stealthiness of our proposed method and the baselines. The defense methods include: (1) spectral signature [32], (2) activation clustering [34], and (3)a textual backdoor defense named ONION [33]. We explain these defense methods used the settings of each method used in our experiment.

4.4.1 Spectral Signature

We use the spectral signature [32], the same method used to detect the fixed and grammar triggers in [18], which has also been widely used in evaluating backdoor attacks in different domains [21], [23], [27], [38], [43]. As reported in [18], the spectral signature can detect both fixed and grammar triggers on simple code models with high detection rates. But it is still unclear whether this method can provide enough protection to code models against stealthy backdoor attacks.

The intuition behind the spectral signature method is that data poisoning can cause the distribution shift (as shown in Figure 3) for the poisoned examples in the dataset. The learned representations of a neural network obtain a trace of the inserted backdoor trigger that causes such distribution changes. Tran et al. [32] theoretically show that the representation of poisoned examples will be highly correlated with the top eigenvector of the covariance of the representation of the whole dataset. Consequently, the spectral signature method ranks all the examples in a dataset in

---

5. Due to the limited GPU resources, we use smaller batch sizes than the settings in the paper [29]. On average, the *BLEU* score of the three models decreases by 0.78.

the order of their correlation with the top eigenvector and takes the high-ranking examples as the poisoned examples.

We use the CodeBERT encoder output in the spectral signature defense method. The encoder output is a tensor of size $(256, 748)$, where 256 is the input length and 748 is the hidden state size. The tensor of each input is then fed into the spectral signature method [32]. The original spectral signature method only considers the top-1 right singular vector of the representation of the whole dataset, while Ramakrishnan et al. [18] show that additional right singular vectors may produce better detection results. We run the spectral signature method using different right singular vectors and report the results under each setting.

4.4.2 Activation Clustering

Chen et al. [34] propose to detect backdoor attacks by analyzing the nueron activation patterns of the deep neural networks, which we refer to as *activation clustering*.

The intuition behind this defense method is that although a model makes the same prediction for a clean example and a poisoned example, the reason for the model to make the prediction is different. On the clean example, we model analyzes the features (e.g., the semantics of the input program) to make decision while on the poisoned example, the model associates the triggers with the prediction. Previous studies show that different neurons are activated when the model utilizies different features. This defense method aims to separate the poisoned examples from the clean examples by clustering the activation patterns.

Algorithm 2 shows the procedure of the activation clustering method. The algorithm takes as input two components: a poisoned dataset D and a model $\mathcal{M}$ that is trained on D. The goal of this algorithm is to separate the poisoned examples D_p and the clean examples D_c in D. For each example d in the dataset D, we feed d into the model $\mathcal{M}$ and obtain the activations of the last hidden layer of $\mathcal{M}$. Similar to the spectral signature method, we use the CodeBERT encoder output in the spectral signature defense method. Each encoder output is a tensor of size $(256, 748)$ and we flat the tensor into a single 1D vector. We then apply the Principal Component Analysis (PCA) [44] to reduce the dimension of each vector to 3, which is the same setting as used in [34]. Although there are a number of clustering methods available (e.g., DBSCAN, Gaussian Mixture Models Affinity Propagation and k-means), Chen et al. [34] find that k-means ($k = 2$) is highly effective at separating the poisonous from legitimate activations so we follow the same setting.

4.4.3 ONION

Qi et al. [33] propose a defense method called ONION, which aims at identifying the textual backdoor attacks. ONION tries to find outlier words in a sentence, which are very likely to be related to backdoor triggers. The intuition behind this method is that the outlier words (i.e., trigger words) markedly decrease the fluency of the sentence and removing them would enhance the fluency. The fluency of a sentence can be quantified by the perplexity computed by a language model.

The ONION works as follows. In the inference process of a backdoored model, for a given test example comprising

---

Algorithm 2: Activation Clustering Defense

Input: D: a poisoned dataset, $\mathcal{M}$: the model trained on D

Output: D_c: clean examples, D_p: poisoned examples

1   $A = [\,]$ # A stores the activation of each example in the poisoned dataset D;
2   **for** d in D **do**
3 | $A_d \leftarrow$ activations of last hidden layer of $\mathcal{M}(d)$ flattened into a single 1D vector;
4   **end**
5   $A' = PCA(A)$ # reduce dimension using PCA;
6   $clusters \leftarrow$ group A' into two clusters;
7   $D_c, D_p \leftarrow$ analyze the clusters ;
8   **return** D_c, D_p

---

n words $d = w_1, \cdots, w_n$, we first use a language model to compute the perplexity of the sentence, denoted by p_0. In the paper introducing ONION, the authors use GPT-2 [45] as the language model. Considering that we try to detect backdoor attacks in the code domain, we use CodeGPT [28], a model that shares the same architecture as GPT-2 but is trained on a dataset of source code [3], as the language model. We compute the *suspicion score* of a word as the decrements of sentence perplexity f_i after removing the i^{th} word: $f_i = p_0 - p_i$. p_i is the perplexity of the example with the i^{th} word removed. A larger f_i indicating that the i^{th} word reduces the fluency of the sentence more and is more likely to be a trigger word.

4.5 Machines, Platforms and Code

All the experiments are performed on a machine running an Ubuntu 18.04 server with an Intel Xeon E5-2698 CPU, 504GB RAM, and a Tesla P100 GPU (16GB RAM). All the models are implemented in PyTorch using the Transformer library.

5 RESEARCH QUESTIONS AND RESULTS

In this section, we evaluate AFRAIDOOR to analyze the threats caused by stealthy backdoor attacks. We conduct experiments to answer the following three research questions:

- RQ1. How stealthy are the examples generated by AFRAIDOOR to automated detection tools?
- RQ2. How stealthy are the examples generated by AFRAIDOOR to human developers?
- RQ3. How does AFRAIDOOR perform in achieving a high attack success rate?
- RQ4. How does AFRAIDOOR affect model performance on clean examples?

Recalling the attack process in Figure 2, the system developers can defend the backdoor attack from three perspectives: (1) filter the poisoned examples in the training data, (2) filter the poisoned examples in the test data, and (3) the impact of AFRAIDOOR on the model performance. The three points correspond to the three research questions.

5.1 RQ1. How stealthy are the examples generated by AFRAIDOOR to automated detection tools?

Motivation. Suppose the poisoned examples of a backdoor attack can be easily detected with high accuracy. In that

TABLE 2: The detection results of spectral signature under different poisoning rate settings.

Task	Attack	$\beta = 1$			$\beta = 1.5$		
		0.5%	1%	5%	0.5%	1%	5%
Code Summarization	AFRAIDOOR	0.00%	0.00%	2.66%	0.00%	3.20%	20.62%
	Fixed	89.06%	67.20%	91.30%	92.19%	88.00%	93.83%
	Grammar	0.00%	91.13%	89.00%	0.00%	92.74%	92.37%
Method Name Prediction	AFRAIDOOR	1.56%	0.80%	25.39%	1.56%	0.80%	27.63%
	Fixed	46.32%	87.20%	87.24%	47.97%	88.00%	92.85%
	Grammar	0.00%	12.10%	86.80%	0.00%	21.77%	91.20%

case, the threat that this attack can cause is limited as the model developer can remove these poisoned examples and train models on the remaining examples. Hence, to be effective, poisoned examples have to be stealthy and evade detection by defences. Such a stealthiness requirement is the motivation to propose and evaluate AFRAIDOOR. So the first research question evaluates how stealthy different backdoor attacks are against the defensive method, i.e., spectral signature.

5.1.1 Stealthiness against Spectral Signature

Evaluation Metrics. Yang et al. [26] propose to evaluate the stealthiness of backdoor attacks in language models using the *Detection Success Rate* (*DSR*) metric, which calculates the rate of truly poisoned examples in the examples returned by a detection method. The detection method used by Yang et al. [26] assumes single-word insertion as the trigger, which do not have the desirable qualities of being syntactic-valid and semantic-preserving. Therefore, it is not applicable to attack code models.

As introduced in Section 4.4.1, we use the spectral signature method to detect poisoned examples. This method is widely used [21], [23], [27], [38], [43] and also adopted by Ramakhrisnan et al. [18]. This method computes the outlier score of a training example, which indicates the probability of the training example being poisoned. We rank all the examples based on their outlier scores. Assuming that the poisoning rate is α and the number of total examples is N, we introduce a parameter *removal ratio* to control the number of removed examples and denote it as β. We remove the top $\alpha \times \beta \times N$ examples with the highest outlier scores from the ranked examples. Then we define the *Detection Success Rate @ the removal ratio β ($DSR@\beta$)* as:

$$DSR@\beta = \frac{\text{No. Poisoned examples}}{\alpha \times \beta \times N} \quad (4)$$

A lower $DSR@\beta$ suggests that a backdoor attack is stealthier as less truly poisoned examples are removed.

Results. We present the results of the three backdoor attacks in Table 3.[6] If a backdoor attack is the stealthiest one under a given setting (i.e., having the lowest $DSR@\beta$), the corresponding results are highlighted in **bold** in Table 3. We find that *our adaptive backdoor attack is always the stealthiest one* on both the code summarization and method name prediction

6. Due to the limited space, Table 3 presents the $DSR@1$ and $DSR@1.5$ results when the top 3 right singular vectors are used. We refer the interested readers to our appendix `./appendix/ICSE-23-results.xlsx` in the replication package for the full results.

TABLE 3: The detection success rates (DSR) of different backdoor attacks. Lower DSR means an attack is stealthier. k is the number of right singular vectors used in detection.

k	Attack	Detection Success Rate ($DSR@\beta$)			
		Code Summarization		Method Name Prediction	
		$\beta = 1$	$\beta = 1.5$	$\beta = 1$	$\beta = 1.5$
1	AFRAIDOOR	**1.16**	**15.4**	**29.26**	**41.43**
	Fixed	94.47	99.34	85.21	86.50
	Grammar	94.96	99.72	41.07	42.49
2	AFRAIDOOR	**1.84**	**2.78**	**24.66**	**28.44**
	Fixed	94.89	99.34	92.37	97.77
	Grammar	94.76	99.71	90.76	97.21
3	AFRAIDOOR	**1.32**	**2.42**	**35.52**	**40.54**
	Fixed	94.96	99.30	90.44	96.15
	Grammar	94.24	99.67	91.70	97.73
Avg	AFRAIDOOR	**1.42**	**6.87**	**29.81**	**36.80**
	Fixed	94.71	99.33	89.34	93.47
	Grammar	94.97	99.71	74.51	79.14

tasks. We compute the average detection rates and put the results in the last three rows in Table 3. On the code summarization task, the average $DSR@1$ and $DSR@1.5$ of the adaptive trigger are only 1.42% and 6.87%. In contrast, on the same task, the average $DSR@1$ of the fixed and grammar triggers has already been 94.71% and 94.97%, respectively. If we are willing to remove more examples (e.g., setting β as 1.5), 99.33% and 99.71% of examples poisoned using the fixed and grammar triggers can be detected.

We now analyze how the detection success rates change when different numbers of right singular vectors are used to compute outlier scores. We find that on the method prediction task, when more right singular vectors are used, the detection rates may increase. A similar observation is also made in [18]. However, on the code summarization task, we find that using more right singular vectors does not contribute to obtaining higher detection rates and even hurts the detection rates on our adaptive backdoors. For example, when β is set as 1.5, the detection rate drops from 15.4% to 2.42% when 3 rather than 1 vectors are used. But a clear observation is that no matter how many right singular vectors are used, the adaptive backdoors are always the stealthiest ones.

Table 2 shows how the detection results of spectral signature change when the poisoning rate α changes. We can observe that when the detection rate decreases, it becomes harder to detect poisoned examples. For example, $DSR@1$ for AFRAIDOOR on the method name prediction task decreases from 25.39% to 1.56% when α decreases from 5% to 0.5%. We can see such decreases on the baseline attacks as well. For the fixed triggers, $DSR@1$ decreases to 89.06% and 46.32% on the code summarization and method name prediction tasks, respectively, when α decreases from 5%

TABLE 4: The detection results of ONION [33] under different settings. Lower detection rates indicate that the backdoor attack is stealthier.

Task	Attack	$\gamma = 1.0$			$\gamma = 1.5$		
		CodeGPT	CodeBERT	Pseudo-label	CodeGPT	CodeBERT	Pseudo-label
Code Summarization	AFRAIDOOR	**13.22%**	**21.56%**	**19.22%**	**21.44%**	**29.52%**	**25.28%**
	Grammar	21.33%	22.74%	27.12%	30.88%	32.13%	38.49%
	Fixed	41.58%	29.42%	29.50%	55.83%	43.92%	42.58%
Method Name Prediction	AFRAIDOOR	**13.22%**	**15.64%**	**22.53%**	**21.44%**	**24.96%**	**29.76%**
	Grammar	21.33%	18.48%	32.45%	30.88%	24.96%	43.09%
	Fixed	15.75%	16.33%	36.25%	28.08%	29.17%	47.50%

to 0.5%. For the grammar trigger, $DSR@1$ decreases to 0% when the poisoning rate is 0.5% on both tasks. However, when the poisoning rate is 1%, $DSR@1$ is 91.13% and 12.10% on the two tasks, which is much higher than that of the AFRAIDOOR. As a result, our proposed method is stealthier than the baselines under different poisoning rates.

5.1.2 Stealthiness against ONION

Evaluation Metrics. Unlike Spectral Signature that aims to assign a binary label to each example indicating whether an example is poisoned, ONION [33] aims to identify suspicious words that are likely to be triggers in an example. In practical usage, ONION eliminates certain suspicious words from the model input, subsequently providing the modified input to the model to prevent the backdoor from been triggered. As a result, a backdoor attack is stealthier if less triggers generated by this attack are identified as suspicious words by ONION.

ONION ranks all the words in an input: top ranked words are more likely to be triggers. We define the *Trigger Detection Rate (TDR)* to evaluate its ability to detect triggers. Assuming that we examine the top k words in the ranked list, we can compute the ratio of triggers among the top k words. However, as different backdoor attack methods may generate different numbers of triggers, we cannot directly compare the TDR on different backdoor attacks. Similar to the evaluation metric for Spectral Signature, we define the *Trigger Detection Rate @ the removal ratio γ ($TDR@\gamma$)* as:

$$TDR@\gamma = \frac{\text{No. Trigger words}}{M \times \gamma} \quad (5)$$

In the above equation, M is the length (i.e., number of tokens) of a trigger in this example. The value γ adjusts the number of words that are examined in the ranked list. A small γ means that the defender only examines the top ranked words, while a large γ means that the defender examines more words in the ranked list. A lower $TDR@\gamma$ suggests that a backdoor attack is stealthier as less triggers are identified as suspicious words by ONION.

Equation 5 pertains to a specific example. To assess ONION's overall effectiveness against a backdoor attack, we calculate the average $TDR@\gamma$ across all examples within the test set. It is important to acknowledge that, in practice, a model developer might not possess precise knowledge regarding the exact number of triggers present in an example. However, for the purpose of evaluation in this particular research question, we assume that the model developer possesses knowledge of the trigger count within each example.

Results. To implement ONION, we need to compute the perplexity of each input. In the original paper that proposes ONION, the authors use GPT-2 [45], a decoder only model trained on natural langauge, to compute the perplexity. In this paper, we use CodeGPT [28] a model that shares the same architecture as GPT-2 but is trained on code, to compute the perplexity. We also compute the perplexity of each input using CodeBERT [13], an encoder only model that is trained on both code and natural language. We also use the models trained on poisoned datasets (i.e., the victim models in our experiments) to compute the perplexity of each input.

We present the performance of three variants of ONION (with different perplexity computation) on three different backdoor attacks in Table 4. The results show that the triggers generated by our proposed AFRAIDOOR are consistently harder to be detected by ONION. For example, when $\gamma = 1.0$, if we use CodeGPT to compute perplexity to rank the words, only 13.22% top ranking words are indeed triggers generated by AFRAIDOOR, while 21.33% and 41.58% top ranking words are triggers generated by Grammar and Fixed, respectively. It suggests that AFRAIDOOR is stealthier than the two baseline methods, and grammar triggers are stealthier than fixed triggers.

5.1.3 Stealthiness against Activation Clustering

Evaluation Metrics. The principle behind Activation Clustering [34] is to divide all instances into two distinct clusters. As stated in the original paper [34] that introduces activation clustering, *when the data is unpoisoned, we (the authors of [34]) find that the activations tend to separate into two clusters of more or less equal size.* They further explain that when $\alpha\%$ of the data is poisoned, we expect that one cluster contains roughly $\alpha\%$ of the data, while the other cluster contains roughly $(100 - \alpha)\%$ of the data. As the poisoning rate is usually small, we expect the size (i.e., the ratio of its contained examples to the total number of examples) of the smaller cluster to be close to the poisoning rate $\alpha\%$. We also measure the ratio of the poisoned examples in the smaller cluster, and expect it to be as large as possible.

Results. Table. 5 presents the results under different settings of the activation clustering method. The number not in parentheses is the ratio of examples in the smaller cluster to all the examples, which we expect to be close to the poisoning rate α. The number in parentheses is the ratio of

TABLE 5: The detection results of activation clustering [34] under different settings. α is the poisoning rate. The number without parentheses is the ratio of examples in the smaller cluster to all the examples. The number in parentheses is the ratio of poisoned examples in the smaller cluster.

Task	Attack	The Ratio difference under various α		
		$\alpha = 0.5\%$	$\alpha = 1\%$	$\alpha = 5\%$
CS	AFRAIDOOR	35.82 (1.08)	47.10 (1.64)	48.31 (9.41)
	Fixed	34.24 (1.24)	43.43 (1.92)	45.58 (0.0)
	Grammar	42.23 (0.73)	36.03 (2.3)	42.08 (12.25)
NMP	AFRAIDOOR	5.96 (0.03)	48.48 (0.90)	46.44 (3.50)
	Fixed	42.33 (0.65)	15.58 (0.26)	36.22 (12.02)
	Grammar	12.67 (0.37)	36.89 (0.58)	38.57 (5.34)

TABLE 6: The results of user study for detecting poisoned examples manually. (DR: Detection Rate; FPR: False Positive Rate; FT: Finishing Time).

	Attacks	$\mathcal{P}1$	$\mathcal{P}2$	$\mathcal{P}3$	Average
DR	AFRAIDOOR	0.00%	6.67%	6.67%	4.45%
	Fixed	100%	100%	100%	100%
	Grammar	86.67%	80%	100%	88.89%
FPR	AFRAIDOOR	100%	95.00%	95.65%	96.99%
	Fixed	0.00 %	6.25%	0.00%	2.08%
	Grammar	11.75%	21.43%	15.00%	16.06%
FT	AFRAIDOOR	147 mins	120 mins	112 mins	126 mins
	Fixed	45 mins	17 mins	70 mins	44 mins
	Grammar	80 mins	40 mins	83 mins	67 mins

poisoned examples in the smaller cluster, which we expect to be as large as possible. The smallest cluster we observe contains 5.96% of examples, which is around 12 times larger than the poisoning rate 0.05%. In this cluster, only 0.3% of examples are poisoned. In other settings, the size of smaller clusters is at least one order of magnitude larger than the poisoning rate and the ratio of poisoned examples in the smaller clusters is usually less than 5%. These results suggest that the activation clustering method is not effective in grouping poisoned and clean examples into two distinct clusters.

> **Answers to RQ1**: Our proposed AFRAIDOOR is stealthier than the two baseline methods under the defense of ONION and spectral signature. Activation clustering is not effective in detecting backdoor attacks for all three evaluated attacks.

5.2　RQ2. How stealthy are the examples generated by AFRAIDOOR to human developers?

Motivation. Although automated detection methods like the Spectral Signature [32] and ONION [33] offer a degree of defense against backdoor attacks, they are not the only line of protection. Manual review of the training dataset by human developers is an important aspect of defense strategies that is routinely practiced in some scenarios. Human review serves as a critical checkpoint where anomalies that are overlooked by automated systems can be identified due to the natural ability of humans to recognize patterns or irregularities that might not be discernible to an algorithm. Moreover, backdoor attacks employing code transformations may result in unnatural examples that, while possibly bypassing automated defenses, could be noticeably anomalous to a human reviewer. The above consideration motivates us to assess how stealthy the poisoned examples are to human developers.

Evaluation Metrics. Participants are tasked with identifying any examples that they find suspicious or indicative of data poisoning. We use the following metrics to quantitatively measure the results of the user study:

1) **Detection Rate**: This measures the percentage of poisoned samples that are successfully identified by the participants. A low detection rate indicates the high stealthiness of the poisoning method.

2) **False Positive Rate**: This is the percentage of normal (unpoisoned) samples that were incorrectly identified as poisoned by the participants. A high false positive rate means that poisoned examples can be well-hidden in normal examples, indicating the high stealthiness of the poisoning method.

3) **Finishing Time**: This metric refers to the time required by a participant to complete the task of identifying all suspected poisoned examples in the dataset. If this process takes a longer time, it suggests that the poisoned examples are well-integrated and hard to detect, implying higher stealthiness of the poisoning method.

Quantitative Results. In this study, we create three distinct variants of poisoned training data, each variant utilizing a different method of trigger injection: AFRAIDOOR, fixed, and grammar triggers. Each trigger type is embedded into 5% of the training data. To achieve a reasonable representation of the entire dataset, we adopt a widely recognized sample size calculator,[7] setting a confidence level at 95% and a confidence interval at 5, which gives us a representative sample size of 375 from each poisoned training set.

We conduct a user study involving three participants who are software engineers with a minimum of four years of experience in Python programming. All the participants have a Bachelor/Master degree in Computer Science. Prior to the study, we familiarize them with the concepts of data poisoning and backdoor attacks without revealing the specifics of our approach or the baseline methods. As such, the participants had no prior knowledge about the nature of the inserted triggers. The sampled examples, shuffled for randomness, are presented to each participant in three separate groups. Each group comprises examples contaminated with only one type of trigger. The quantitative results are summarized in Table 6 and analzyed as follows.

We can observe that participants take longer time to claim that they have finished the task of finding all the poisoned examples generated by AFRAIDOOR than that generated the other two baseline methods. On average, participants spend 126 minutes to complete the tasks when they are presented with poisoned examples by AFRAIDOOR. On the contrast, it takes them 44 and 67 minutes to finish the tasks when they are presented with poisoned examples generated by the fixed and grammar triggers, respectively.

---

7. https://www.surveymonkey.com/mp/sample-size-calculator/4

Then, we adopt the Wilcoxon rank-sum test [46] to compute the significance of the differences between task completion time under AFRAIDOOR and the baseline methods. The result shows that the difference is statistically significant (p-value < 0.01), suggesting that it is cognitively more challenging to detect poisoned examples generated by AFRAIDOOR. It also indicates that AFRAIDOOR is stealthier than the baseline methods.

We find that the detection rates on poisoned examples generated by AFRAIDOOR are lower than those generated by the baseline methods. More specifically, the average detection rates on AFRAIDOOR is 4.45%, while the average detection rates on the baseline methods are 100% and 88.89%, respectively. Similarly, the Wilcoxon rank-sum test shows that the difference between the detection rates on AFRAIDOOR and the baseline methods is statistically significant (p-value < 0.01). We also observe that the false positive rates on poisoned examples generated by AFRAIDOOR are much higher than those generated by the baseline methods. It indicates that the human participants cannot distinguish the poisoned examples generated by AFRAIDOOR from the normal examples. However, they can easily find the difference between the poisoned examples generated by the baseline methods and the normal examples. Both low detection rates and high false positive rates suggest that AFRAIDOOR is stealthier than the baselines.

Qualitative Results. Upon completion of the study, we proceed to analyze the above quantitative results, followed by interviewing the participants about the methodologies used for finding the three types of triggers. We engage each participant in a separate interview to gain insights into their experiences during the study, particularly focusing on the challenges they encountered in identifying poisoned examples. This discourse helps us gather valuable first-hand user perspectives on the stealthiness of our method compared to the baseline techniques. To provide a qualitative depiction of our findings, we present below select statements from each participant.

When talking about the feeling to annotate the dataset modified by AFRAIDOOR, the participant $\mathcal{P}1$ mentions that *'Even I spent much time reading the code, I cannot find any obvious common patterns in the dataset. Eventually I tend to label those with messy code as the poisoned ones.'* After we explain how AFRAIDOOR works, the participant $\mathcal{P}1$ further commented *'It means that triggers spread across multiple locations of code ... An annotator has to read the whole program and keep all the information in mind ... It is rather difficult to find such patterns by human.'* When talking about the fixed triggers, $\mathcal{P}1$ mentions that *'During skimming the code, I suddenly noticed that there seems to be something that I had seen. So I read the annotated code again and find that there is an exception handling statement appearing three times.'* The participant $\mathcal{P}1$ said *'I used the search function in VS Code[8] to find all the code that contain this code snippet so it did not take a long time to finish.'*

The participant $\mathcal{P}2$ also use the search function to find other poisoned examples once he finds suspicious code snippets. $\mathcal{P}2$ mentions that *'For the grammar trigger, I know some code snippets are suspicious. But I could not find other poisoned examples simply by searching. I have to read the code again*

8. VS Code is an integrated development environment.

TABLE 7: The impact of attacks on model performance.

Task	Model	Trigger	ASR	ASR-D
CS	CodeBERT	AFRAIDOOR	98.53	96.35 (-2.18)
		Fixed	100.00	8.27 (-91.73)
		Grammar	100.00	10.35 (-89.65)
	PLBART	AFRAIDOOR	93.78	91.16 (-2.26)
		Fixed	100.00	8.28 (-91.72)
		Grammar	100.00	8.15 (-91.85)
	CodeT5	AFRAIDOOR	95.51	91.44 (-4.07)
		Fixed	100.00	8.13 (-91.87)
		Grammar	100.00	10.61 (-89.39)
MNP	CodeBERT	AFRAIDOOR	98.14	76.58 (-21.56)
		Fixed	100.00	12.76 (-87.24)
		Grammar	100.00	14.25 (-85.75)
	PLBART	AFRAIDOOR	97.01	86.86 (-20.15)
		Fixed	100.00	12.62 (-87.38)
		Grammar	100.00	14.49 (-85.51)
	CodeT5	AFRAIDOOR	98.15	77.70 (-20.45)
		Fixed	100.00	12.76 (-87.24)
		Grammar	100.00	14.49 (-85.51)

to see whether there were similar code snippets.' However, all the three participants fail to identify the triggers generated by AFRAIDOOR. Both $\mathcal{P}2$ and $\mathcal{P}3$ agree that it will be very difficult to find the triggers if the poisoning rate is lower.

> **Answers to RQ2**: AFRAIDOOR can generate stealthier poisoned examples than the baseline methods. More specifically, users takes $xx\%$ longer time to finish the task when reading poisoned examples generated by AFRAIDOOR. The detection rates on AFRAIDOOR is close to 0% while users can find all the poisoned examples generated by the baselines.

5.3 RQ3. How does AFRAIDOOR perform in activating backdoors successfully?

Motivation. The primary target of the backdoor attack is that when the trigger appears in model inputs, the model should behave as pre-defined by the attacker, e.g., produce a specific label. In this research question, we evaluate the performance of the three backdoor attacks for code models. Based on results from RQ1 (in Section 5.1), spectral signature [32] and ONION [33] can detect the poisoned examples while the activation clustering [34] demonstrates limited effectiveness. We consider two scenarios: whether the defense method is used or not. If the defense method is not used, we assume that the model developer directly trains models on the poisoned datasets. If the spectral signature is used as defense, we assume that the model developer first removes the potentially poisoned examples and trains the models on the purified datasets. If the ONION is used as defense, we assume that the model developer use ONION to detect and remove suspicious words from each input before feeding them into the model.

5.3.1 Spectral Signature as Defense

Evaluation Metrics. We introduce the *Attack Success Rate* (*ASR*) to measure the performance of backdoor attacks when no defensive method is used. Formally, *ASR* is defined as follows.

$$ASR = \frac{\sum_{x_i \in \mathcal{X}} M_b(x_i) = \tau}{\sum_{x_i \in \mathcal{X}} x_i \text{ contains triggers}} \qquad (6)$$

The denominator represents the total number of poisoned examples in a dataset. M_b is a model trained on the poisoned dataset. $M_b(x_i) = \tau$ means that an input with trigger can force the model to produce τ as output, which is pre-defined by the attacker. In other work, x_i is a successful attack. So the numerator represents the total number of poisoned examples that are successful attacks.

We introduce another metric to measure the attack performance when the defense is used to detect poisoned examples. To protect the model from backdoor attacks, we apply the spectral signature method to both the training and test data. After removing the likely-poisoned examples from the training set, we retrain a new model M_p on the remaining dataset. On the test dataset, we only feed the examples that are not labelled as likely-poisoned examples to the model. Then we introduce the *Attack Success Rate Under Defense*, denoted by *ASR-D*. We define *ASR-D* as follows.

$$ASR_D = \frac{\sum_{x_i \in \mathcal{X}} M_p(x_i) = \tau \wedge \neg \mathcal{S}(x_i)}{\sum_{x_i \in \mathcal{X}} x_i \text{ contains triggers}} \qquad (7)$$

We introduce an additional condition to the numerator: $\neg \mathcal{S}(x_i)$. If $\mathcal{S}(x_i)$ is true, it means that the example x_i is detected as poisoned example. So $\sum_{x_i \in \mathcal{X}} M_p(x_i) = \tau \wedge \neg \mathcal{S}(x_i)$ means the number of all the poisoned examples that are not detected by the spectral signature and produce success attacks.

Results. We put different attacks' *ASR* and *ASR-D* in Table 7. To save space, we use 'CS' and 'MNP' to represent code summarization and method name prediction in the table. We first analyze the attack performance when no defense is used. From the Table 7 we can find that both fixed and grammar triggers can achieve *ASR* of 100%, meaning that the two types of triggers can steadily activate backdoors in models. In contrast, the proposed adaptive trigger has slightly lower *ASR*. On the code summarization task, our adaptive trigger achieve *ASR* of 98.53%, 93.78%, and 95.51% on the CodeBERT, PLBART, and CodeT5, respectively. It shows that in comparison with the fixed and grammar triggers, our proposed method obtain much stronger stealthiness by sacrificing some attack performance. We present a further analysis of those unsuccessful attacks in Section 6.1.

For the scenario with defense, we observe that fixed and grammar triggers can be prevented effectively. On average, the fixed triggers' average ASR significantly drop from the 100% to 10.47%, and the grammar triggers' average ASR drop from the original 100% to 12.06%. Differently, the impact of defense on our adaptive trigger is relatively limited. On the code summarization task, the average *ASR* drops by 2.96% (from 95.94% to 92.98%). On the method name prediction task, the same metric drops by 20.72% (from 97.77% to 77.05%). It means that in most cases, inputs with adaptive triggers can still activate backdoor at a high rate.

5.3.2 ONION as Defense

Evaluation Metrics. To defend against backdoor attacks, we use ONION to rank the suspicious words in each input and then remove the top k suspicious words from the input before feeding it to the model. A successful instance of defense means that the model does not produce attack-specified output on the poisoned input after the defense is applied. We define the metric *Defense Success Rate @k* (*DSR@k*) as follows.

$$DSR@k = \frac{\sum_{x_i \in \mathcal{X}} M_p(x_i - word_k) \neq \tau}{\sum_{x_i \in \mathcal{X}} x_i \text{ contains triggers}} \qquad (8)$$

In the above equation, $word_k$ is the top k suspicious words labelled by ONION and $x_i - word_k$ means the input x_i with the top k suspicious words removed. The numerator in Equation 8 is the number of poisoned examples that do not produce attack-specified output after the defense is applied. The denominator in Equation 8 is the number of the poisoned examples. Given the same k value, a lower *DSR@k* means the attack is harder to be defended, i.e., stealthier.

Results. The results of *DSR@k* on two tasks are shown in Figure 5 and Figure 6. On both tasks, the *DSR@k* of AFRAIDOOR is much lower than that of baselines. When k gradually increase from 1 to 5 (meaning that we remove more suspicious words), there is no significant change in the *DSR@k* of AFRAIDOOR. In contrast, the *DSR@k* of the fixed triggers is always close to 100% when k ranges from 1 to 5. When k is small (e.g., 1), grammar triggers demonstrate some stealthiness, but the *DSR@k* of grammar triggers also increases dramatically when k increases. More specifically, when $k = 1$, the *DSR@k* of grammar triggers is 3.33% on the code summarization task and 74.93% on the method name prediction task. The value increases to over 95% when $k = 5$ and $k = 2$ on code summarization method name prediction task, respectively.

Our results show that AFRAIDOOR can achieve strongest stealthiness in the three evaluated attacks. Grammar triggers can also demonstrate stronger stealthiness than fixed triggers when k is small. The evaluation on multiple tasks and models warn us that the adaptive backdoor can bypass the spectral signature and ONION method, calling for attention on developing stronger defensive methods.

> **Answers to RQ3**: When the defense method is not applied, both AFRAIDOOR and baselines have very high ASR. However, once the spectral signature is applied, the success rates of baselines decrease dramatically to 10.47% and 12.06%, while the success rate of AFRAIDOOR are 77.05% and 92.98% on the two tasks on average. ONION can effectively defend against fixed and grammar triggers, but it has limited impact on AFRAIDOOR.

5.4 RQ4. How does AFRAIDOOR affect the model performance on clean examples?

Motivation. Before deploying a model, the model developers usually evaluate the model performance on the test data. Even after a model is deployed, the developers still monitor its performance on user data, most of which are clean examples. If the model has poor performance, then the developers may not even deploy the model and the
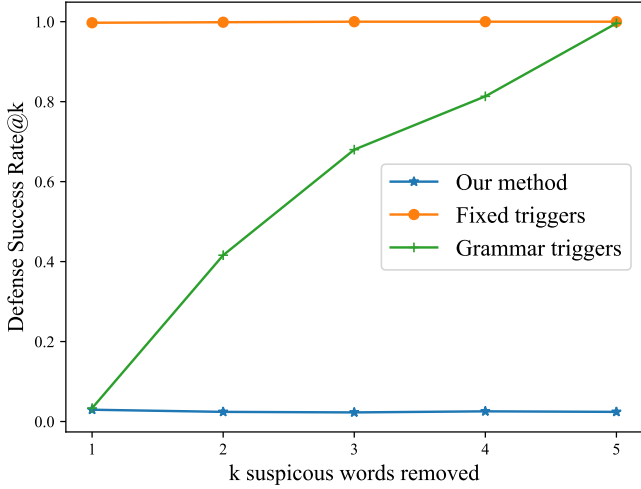
Fig. 5: The *Defense Success Rate@k* of AFRAIDOOR and baselines on the code summarization task.


Fig. 6: The *Defense Success Rate@k* of AFRAIDOOR and baselines on the method name prediction task.

attacker cannot feed poisoned input to the model. Thus, researchers [26], [47], [48] believe that backdoor attacks should have as minimal impact on the model performance on clean examples as possible. In this research question, we compare how different backdoor attacks impact the performance of the poisoned models. Same as RQ2, we consider the two scenarios: with and without defense.

Evaluation Metrics. Following the settings in [29], we use *BLEU* score [49] to evaluate a model's clean performance on code summarization and method name prediction. A higher *BLEU* indicates better model performance. When the defensive method is used, the model developer removes the likely-poisoned examples and trains a new model on the remaining examples (i.e., purified datasets), which we call the *purified model*. We define the *BLEU-D* score as the *BLEU* score of the purified model on the same set of clean examples. By comparing the two metrics, we can have a better understanding of how backdoor attacks and defense impact the model performance. If *BLEU-D* is smaller than *BLEU*,

TABLE 8: Backdoor attacks and defense affect model performance.

Task	Model	Trigger	*BLEU*	*BLEU-D*
CS	CodeBERT	AFRAIDOOR	16.79 (-0.71)	17.38 (+0.59)
		Fixed	17.19 (-0.31)	16.94 (-0.25)
		Grammar	17.10 (-0.40)	16.49 (-0.61)
	PLBART	AFRAIDOOR	17.99 (-0.36)	18.21 (+0.22)
		Fixed	18.17 (-0.18)	18.05 (-0.12)
		Grammar	17.94 (-0.41)	17.62 (-0.32)
	CodeT5	AFRAIDOOR	18.66 (+0.05)	18.60 (-0.06)
		Fixed	18.56 (-0.05)	18.60 (+0.04)
		Grammar	18.53 (-0.08)	18.41 (-0.12)
MNP	CodeBERT	AFRAIDOOR	43.08 (-0.27)	42.29 (-0.79)
		Fixed	42.87 (-0.48)	43.03 (+0.16)
		Grammar	42.94 (-0.41)	43.12 (+0.18)
	PLBART	AFRAIDOOR	42.18 (-0.33)	42.29 (-0.11)
		Fixed	42.65 (+0.14)	42.31 (-0.34)
		Grammar	42.47 (-0.04)	42.50 (+0.03)
	CodeT5	AFRAIDOOR	46.40 (+0.36)	46.17 (-0.23)
		Fixed	46.41 (+0.37)	46.57 (-0.16)
		Grammar	45.97 (-0.07)	46.33 (+0.36)

it means that applying defense to filter poisoned examples can hurt the model performance on clean examples.

Results. Table 8 documents the evaluation metrics *BLEU* and *BLEU-D* for the three attacks on two tasks. The *BLEU* column in Table 8 shows the performance of the poisoned models as well as the changes compared to the original models that are trained on clean examples (reported in Table 1); changes are put in the parentheses and '-'/'+' means performance decrease/increase after attack. Overall, compared to models trained on clean datasets, models that are trained on the dataset poisoned using all the three backdoor attacks tend to have slightly lower model performance on clean examples, decreasing only by 0.18 *BLEU* score on average.

We are interested in whether *the performance decrease caused by the adaptive trigger is significantly larger than that of caused by the fixed and grammar triggers.* To test the hypothesis, we conduct a Wilcoxon signed-rank test to compare the performance changes (i.e., the numbers surrounded by the parentheses in the column *BLEU*) caused by AFRAIDOOR and two baseline attacks. The p-values we obtained are 0.43 (AFRAIDOOR and fixed trigger) and 0.24 (AFRAIDOOR and grammar trigger), indicating that there is no statistically significant difference between our approach and the other two baseline approaches in terms of the model performance on clean examples. It suggests that AFRAIDOOR achieves higher stealthiness but does not sacrifice more clean performance than the baseline methods at the same time.

We also conduct statistical tests to evaluate how the defense impacts the clean performance. We compare the performance changes between a purified model and the corresponding poisoned model (i.e., Column *BLEU-D*, the last column in Table 8). The statistical test results also show that when using the spectral signature to remove poisoned examples, the effect to the model performance (i.e., the difference between *BLEU* and *BLEU-D*) is not significantly different among the three backdoor attacks.

> **Answers to RQ4**: All the three attacks cause slightly negative impacts on the clean performance, however these impacts are not statistically significant.

6 DISCUSSION

6.1 The Characteristics of Unsuccessful Attacks

Based on the results of RQ2, we find that our adaptive triggers are indeed stealthier but inevitably sacrifice some attack effectiveness. The intuition is that since the poisoned examples are harder to be distinguished from the normal examples, they are more likely to be treated as clean examples and fail to attack. We separate all the poisoned examples into two groups: successful attacks and unsuccessful attacks[9]. Then, we compare the average lengths of examples in the two groups. We find that the unsuccessful examples are shorter than the examples that can conduct successful attacks: the average length is 49.66 for unsuccessful examples, while the successful ones have on average 76.70 tokens, 54.45% longer than the unsuccessful ones. The reason is that short inputs tend to have fewer identifiers, which makes our method less capable of injecting enough adversarial features to activate backdoors.

6.2 Extension to Other Software Engineering Tasks

The language models of code have also been used to do code search and achieve state-of-the-art performance [13], [14]. In this process, a user sends natural language queries to the model and the model processes both user queries and code to return relevant results. In the context of our paper (generation task), the backdoor attack essentially tries to build strong connections between stealthy triggers (in the input) and specific outputs. In the code search task, we can adapt our proposed method to build such connections between specific inputs and the stealthy triggers in the output.

More specifically, the threat model is as follows. An attack expects that when a user sends queries with certain keywords, the code search model returns code that has the triggers. This can also cause security risks to the users. For example, the keywords can be security related (e.g., encryption) and the model returns code with triggers and vulnerabilities (e.g., use insecure encryption API). Another instance can be that a user sends queries containing the keyword 'database,' the model returns code with triggers and SQL injection vulnerability. The above attack can be completed using our proposed backdoor attack. The training data of code search models consists of pairs of natural language queries and relevant code. We poison the training data by inserting keywords into a query and injecting triggers and vulnerabilities into the corresponding code. In this way, the model will learn the connection between input keywords and code with triggers and vulnerabilities.

6.3 Suggestions for Mitigating Backdoor Attacks

We discuss some practices that can potentially mitigate the effects of backdoor attacks. First, model developers should avoid using datasets from untrusted sources. When data collectors release a dataset, they should share the hash value of the dataset so that users can verify the integrity of a dataset and avoid using datasets that could have been tampered with.

Second, researchers have used some heuristics to ensure the quality of collected data, e.g., choosing data from repositories with more stars. However, researchers have revealed that the commits and stars can be easily manipulated using *Promotion-as-a-Service* [39], which can be used to make the poisoned repositories more visible to the data collectors and model developers. More research on detecting such malicious promotions and accounts [50] may mitigate data poisoning.

Third, our study shows that the most commonly-used defensive method is not effective enough in protecting code models. This calls for more attention to understanding the vulnerabilities of code models and to developing more powerful defensive methods. Besides, as suggested by the ethical guidelines for developing trustworthy AI [51], model developers may involve humans to establish stronger oversight mechanisms for the collected data and uncover potential poisoned examples.

6.4 Threats to Validity

Threats to Internal Validity. As stated in Section 4, for implementing the three models (CodeBERT, PLBART and CodeT5), we reuse the repository[10] released by the CodeT5 [29] authors. The pre-trained models are extracted from the well-known HuggingFace[11] model zoo. Besides, we replicate the experiment in [29] in code summarization task and observe similar results as reported in the original paper. Thus, we believe that the threats to internal validity are minimum.

Threats to External Validity. In our baseline work [18], it only consider 2 models in 1 task. In the experiment, we expand the experiment by considering 3 state-of-the-art models and evaluate the attacks on 2 large-scale datasets. Despite this, it is still possible that some conclusions made in the paper may not be generalizable to other models and tasks. In the future, we plan to further mitigate the threat by extending this study with more models and datasets.

Threats to Construct Validity. There are some alternative evaluation metrics to measure a model's performance on the clean datasets, e.g., F1-score, or other variants of *BLEU* score. In this paper, we choose *BLEU-4* score as the evaluation metric, which is widely adopted in generation tasks like code summarization and is also used to evaluate the model performances, e.g., [29].

7 RELATED WORK

A series of work has been done to evaluate and improve the quality of various AI systems, e.g., sentiment analysis [52],

---

9. We discard the examples whose length is over 256, the maximal model input length.

10. https://github.com/salesforce/CodeT5
11. https://huggingface.co/

[53], [54], speech recognition [55], [56], reinforcement learning [57], image classification [58], [59], etc. We refer the readers to [60] for a comprehensive survey on AI testing. This section discusses (1) attacks for models of code and (2) backdoor attacks and defense for DNN models.

7.1 Attacking Code Models

Researchers have exposed vulnerabilities in code models, e.g., lacking robustness, not immune to malicious data, etc. Rabin et al. [35] evaluate whether neural program analyzers like GGNN [31] can generalize to programs modified using semantic preserving transformations. Applis et al. [61] extend metamorphic testing approaches for DNN models for software programs to evaluate the robustness of a code-to-text generation model. Pour et al. [36] focus on the embeddings of source code and propose a search-based testing framework to evaluate their robustness. Zhang et al. [9] propose Metropolis-Hastings Modifier to generate adversarial examples for code authorship attribution models. Yang et al. [7] highlight the naturalness requirement in attacking code models and propose to use mask language prediction and genetic algorithms to generate such natural adversarial code examples.

The above works conduct attacks in black-box manners. There are also some attacks that leverage white-box information. Yefet et al. [8] propose DAMP, a method that uses FGSM [62] to adversarially modify variable names in programs to attack code2vec [16], GGNN [63] and GNN-FiLM [64]. Henkel et al. [65] extend Yefet et al.'s work [8] by considering more program transformations, e.g., using `if` branches to insert dead code. Srikant et al. [10] use PGD [66] to further improve Henkel et al.'s [65].

Besides the baseline attack [18] evaluated in our paper, there are some other works that operate data poisoning attacks on datasets of source code. Nguyen et al. [67] find that none of the three state-of-the-art API recommender systems is immune to malicious data in the training set. Schuster et al. [43] add a few specially-crafted files to the training data of a code completion model, and the model outputs will be affected in some security-related contexts. Sun et al. [68] use data poisoning to protect open-source data against unauthorized training usage. Severi et al. [69] insert triggers into binary code that are specially designed to attack the feature-based binary classification models, while this paper poisons the source code to attack the advanced code models.

7.2 Backdoor Attacks and Defense for DNN Models

After Gu et al. [70] first proposed backdoor attacks for (Computer Vision) CV models, Chen et al. [71] point out that the poisoned images and the original examples should be as indistinguishable as possible. Various subsequent studies [22], [72], [73] propose to achieve this goal by limiting the modification under certain constraints, e.g., the L_2 norm. There are a series of defensive methods [74], [75], [76], [77] proposed for CV models, while they cannot be directly applied to the code models as they assume the model input to be continuous. Recently, backdoor attacks are extended to other AI systems like reinforcement learning [78].

The first backdoor attacks on language models are done by Liu et al. [79], which use a sequence of words as the trigger to attack a sentence attitude recognition model. Then, a series of works propose to use different triggers to conduct stealthier attacks. For example, instead of injecting uncommon words [47], Dai et al. use a complete sentence [80] as the trigger. Li et al. inject triggers by using the homograph replacements [25]. In our experiments, we evaluate the proposed method and baselines against three defense methods. First, we follow the baseline work to use the spectral signature as the defensive method to protect the code models [18]. Spectral signature is also adopted in a recent work on applying fixed and grammar triggers to code search tasks [81]. We include the activation clustering [34], which aims to utilize the activation patterns of the model to group the inputs into different clusters. We additionally consider ONION [33] to uncover the suspicious words in the input. It is noted that there is also a line of work that detects the vulnerabilities (e.g., bugs and defects) in code [82], [83], [84]. We do not consider them in this paper as we focus on the backdoor and data poisoning attacks.

8 CONCLUSION AND FUTURE WORK

In this paper, we evaluate the threats caused by stealthy backdoor attacks to code models. We first propose AFRAIDOOR, a method that leverages adversarial features to inject adaptive triggers into model inputs. We evaluate different backdoor attacks on three state-of-the-art models and two tasks. The experiment results show that the existing two backdoor attacks are not stealthy: around 85% of adaptive triggers in AFRAIDOOR bypass the detection in the defense process. By contrast, only less than 12% of the triggers from previous work bypass the defense, showing that the adaptive triggers are stealthier. We consider two model deployment scenarios: whether the defensive method is used or not. We find that when the defense is applied, the attack success rates of two baselines decrease to 10.47% and 12.06%, respectively. By contrast, the success rate of AFRAIDOOR drops to 77.05% on the method name prediction task and 92.98% on the code summarization task. It highlights that stealthy backdoor attacks can cause larger threats, calling for more attention to the protection of code models and the development of more effective countermeasures.

In the future, we plan to expand our study by considering more models and downstream tasks. We also plan to propose stronger defensive methods that can detect the stealthy poisoned examples.

REFERENCES

[1] Y. Yang, X. Xia, D. Lo, and J. Grundy, "A survey on deep learning for software engineering," *ACM Comput. Surv.*, dec 2021.

[2] M. Allamanis, H. Peng, and C. Sutton, "A convolutional attention network for extreme summarization of source code," 2016.

[3] H. Husain, H. Wu, T. Gazit, M. Allamanis, and M. Brockschmidt, "Codesearchnet challenge: Evaluating the state of semantic code search," *CoRR*, vol. abs/1909.09436, 2019.

[4] H. Wei and M. Li, "Supervised Deep Features for Software Functional Clone Detection by Exploiting Lexical and Syntactical Information in Source Code," in *IJCAI 2017*. Melbourne, Australia: International Joint Conferences on Artificial Intelligence Organization, Aug. 2017, pp. 3034–3040.

[5] C. Yang, B. Xu, J. Khan, G. Uddin, D. Han, Z. Yang, and D. Lo, "Aspect-based api review classification: How far can pre-trained transformer model go?" in *2022 IEEE International Conference on Software Analysis, Evolution and Reengineering (SANER)*. Los Alamitos, CA, USA: IEEE Computer Society, mar 2022, pp. 385–395. [Online]. Available: https://doi.ieeecomputersociety.org/10.1109/SANER53432.2022.00054

[6] J. He, B. Xu, Z. Yang, D. Han, C. Yang, and D. Lo, "Ptm4tag: Sharpening tag recommendation of stack overflow posts with pre-trained models," in *2022 IEEE/ACM 30th International Conference on Program Comprehension (ICPC)*. Los Alamitos, CA, USA: IEEE Computer Society, may 2022, pp. 1–11. [Online]. Available: https://doi.ieeecomputersociety.org/10.1145/3524610.3527897

[7] Z. Yang, J. Shi, J. He, and D. Lo, "Natural attack for pre-trained models of code," in *2022 IEEE/ACM 44th International Conference on Software Engineering (ICSE)*, may 2022, pp. 1482–1493.

[8] N. Yefet, U. Alon, and E. Yahav, "Adversarial examples for models of code," *Proceedings of the ACM on Programming Languages*, vol. 4, no. OOPSLA, 2020.

[9] H. Zhang, Z. Li, G. Li, L. Ma, Y. Liu, and Z. Jin, "Generating adversarial examples for holding robustness of source code processing models," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 01, pp. 1169–1176, Apr. 2020.

[10] S. Srikant, S. Liu, T. Mitrovska, S. Chang, Q. Fan, G. Zhang, and U. O'Reilly, "Generating Adversarial Computer Programs using Optimized Obfuscations," *ICLR*, vol. 16, pp. 209–226, 2021.

[11] A. Jha and C. K. Reddy, "Codeattack: Code-based adversarial attacks for pre-trained programming language models," 2022.

[12] J. Cito, I. Dillig, V. Murali, and S. Chandra, "Counterfactual explanations for models of code," pp. 125–134, 2022.

[13] Z. Feng, D. Guo, D. Tang, N. Duan, X. Feng, M. Gong, L. Shou, B. Qin, T. Liu, D. Jiang, and M. Zhou, "CodeBERT: A pre-trained model for programming and natural languages," in *Findings of the Association for Computational Linguistics: EMNLP 2020*. Association for Computational Linguistics, Nov. 2020, pp. 1536–1547.

[14] D. Guo, S. Ren, S. Lu, Z. Feng, D. Tang, S. Liu, L. Zhou, N. Duan, A. Svyatkovskiy, S. F. andz Michele Tufano, S. K. Deng, C. B. Clement, D. Drain, N. Sundaresan, J. Yin, D. Jiang, and M. Zhou, "Graphcodebert: Pre-training code representations with data flow," in *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*, 2021.

[15] U. Alon, S. Brody, O. Levy, and E. Yahav, "code2seq: Generating sequences from structured representations of code," in *International Conference on Learning Representations*, 2019.

[16] U. Alon, M. Zilberstein, O. Levy, and E. Yahav, "Code2vec: Learning distributed representations of code," *Proc. ACM Program. Lang.*, vol. 3, no. POPL, jan 2019.

[17] J. Zhai, X. Xu, Y. Shi, G. Tao, M. Pan, S. Ma, L. Xu, W. Zhang, L. Tan, and X. Zhang, "Cpc: Automatically classifying and propagating natural language comments via program analysis," in *2020 IEEE/ACM 42nd International Conference on Software Engineering (ICSE)*. Los Alamitos, CA, USA: IEEE Computer Society, oct 2020, pp. 1359–1371. [Online]. Available: https://doi.ieeecomputersociety.org/

[18] G. Ramakrishnan and A. Albarghouthi, "Backdoors in neural models of source code," in *2022 26th International Conference on Pattern Recognition (ICPR)*, 2022.

[19] I. Sutskever, O. Vinyals, and Q. V. Le, "Sequence to sequence learning with neural networks," in *Advances in Neural Information Processing Systems*, Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Weinberger, Eds., vol. 27. Curran Associates, Inc., 2014.

[20] F. Qi, Y. Chen, X. Zhang, M. Li, Z. Liu, and M. Sun, "Mind the Style of Text! Adversarial and Backdoor Attacks Based on Text Style Transfer," pp. 4569–4580, 2021.

[21] Z. Wang, J. Zhai, and S. Ma, "Bppattack: Stealthy and efficient trojan attacks against deep neural networks via image quantization and contrastive adversarial learning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 15 074–15 084.

[22] K. Doan, Y. Lao, W. Zhao, and P. Li, "Lira: Learnable, imperceptible and robust backdoor attacks," in *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021, pp. 11 946–11 956.

[23] Y. Li, Y. Li, B. Wu, L. Li, R. He, and S. Lyu, "Invisible backdoor attack with sample-specific triggers," in *IEEE International Conference on Computer Vision (ICCV)*, 2021.

[24] F. Qi, Y. Yao, S. Xu, Z. Liu, and M. Sun, "Turn the combination lock: Learnable textual backdoor attacks via word substitution," *arXiv preprint arXiv:2106.06361*, 2021.

[25] S. Li, H. Liu, T. Dong, B. Z. H. Zhao, M. Xue, H. Zhu, and J. Lu, "Hidden backdoors in human-centric language models," in *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*, 2021, pp. 3123–3140.

[26] W. Yang, Y. Lin, P. Li, J. Zhou, and X. Sun, "Rethinking stealthiness of backdoor attack against NLP models," in *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. Online: Association for Computational Linguistics, Aug. 2021, pp. 5543–5557.

[27] Q. Zhang, Y. Ding, Y. Tian, J. Guo, M. Yuan, and Y. Jiang, "AdvDoor: Adversarial backdoor attack of deep learning system," *ISSTA 2021 - Proceedings of the 30th ACM SIGSOFT International Symposium on Software Testing and Analysis*, pp. 127–138, 2021.

[28] S. Lu, D. Guo, S. Ren, J. Huang, A. Svyatkovskiy, A. Blanco, C. B. Clement, D. Drain, D. Jiang, D. Tang, G. Li, L. Zhou, L. Shou, L. Zhou, M. Tufano, M. Gong, M. Zhou, N. Duan, N. Sundaresan, S. K. Deng, S. Fu, and S. Liu, "Codexglue: A machine learning benchmark dataset for code understanding and generation," *CoRR*, vol. abs/2102.04664, 2021.

[29] Y. Wang, W. Wang, S. Joty, and S. C. Hoi, "Codet5: Identifier-aware unified pre-trained encoder-decoder models for code understanding and generation," in *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021*, 2021.

[30] W. Ahmad, S. Chakraborty, B. Ray, and K.-W. Chang, "Unified pre-training for program understanding and generation," in *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Online: Association for Computational Linguistics, Jun. 2021, pp. 2655–2668.

[31] P. Fernandes, M. Allamanis, and M. Brockschmidt, "Structured neural summarization," in *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*, 2019.

[32] B. Tran, J. Li, and A. Madry, "Spectral signatures in backdoor attacks," in *Advances in Neural Information Processing Systems*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, Eds., vol. 31. Curran Associates, Inc., 2018.

[33] F. Qi, Y. Chen, M. Li, Y. Yao, Z. Liu, and M. Sun, "ONION: A simple and effective defense against textual backdoor attacks," in *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. Online and Punta Cana, Dominican Republic: Association for Computational Linguistics, Nov. 2021, pp. 9558–9566.

[34] B. Chen, W. Carvalho, N. Baracaldo, H. Ludwig, B. Edwards, T. Lee, I. M. Molloy, and B. Srivastava, "Detecting backdoor attacks on deep neural networks by activation clustering," *CoRR*, vol. abs/1811.03728, 2018. [Online]. Available: http://arxiv.org/abs/1811.03728

[35] M. R. I. Rabin, N. D. Bui, K. Wang, Y. Yu, L. Jiang, and M. A. Alipour, "On the generalizability of neural program models with respect to semantic-preserving program transformations," *Information and Software Technology*, vol. 135, p. 106552, 2021.

[36] M. V. Pour, Z. Li, L. Ma, and H. Hemmati, "A search-based testing framework for deep neural networks of source code embedding," in *14th IEEE Conference on Software Testing, Verification and Validation, ICST 2021, Porto de Galinhas, Brazil, April 12-16, 2021.* IEEE, 2021, pp. 36–46.

[37] X. Chen, A. Salem, D. Chen, M. Backes, S. Ma, Q. Shen, Z. Wu, and Y. Zhang, "BadNL: Backdoor Attacks against NLP Models with Semantic-preserving Improvements," *ACM International Conference Proceeding Series*, no. Cv, pp. 554–569, 2021.

[38] E. Bagdasaryan and V. Shmatikov, "Blind backdoors in deep learning models," in *30th USENIX Security Symposium (USENIX Security 21)*, 2021, pp. 1505–1521.

[39] K. Du, H. Yang, Y. Zhang, H. Duan, H. Wang, S. Hao, Z. Li, and M. Yang, "Understanding Promotion-as-a-Service on GitHub," *ACM International Conference Proceeding Series*, pp. 597–610, 2020.

[40] Z. Li, G. Chen, C. Chen, Y. Zou, and S. Xu, "Ropgen: Towards robust code authorship attribution via automatic coding style transformation," in *2022 IEEE/ACM 44th International Conference on Software Engineering (ICSE).* Los Alamitos, CA, USA: IEEE Computer Society, may 2022, pp. 1906–1918.

[41] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers).* Minneapolis, Minnesota: Association for Computational Linguistics, Jun. 2019, pp. 4171–4186.

[42] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov, "Roberta: A robustly optimized BERT pretraining approach," *CoRR*, vol. abs/1907.11692, 2019.

[43] R. Schuster, C. Song, E. Tromer, and V. Shmatikov, "You autocomplete me: Poisoning vulnerabilities in neural code completion," in *30th USENIX Security Symposium (USENIX Security 21).* USENIX Association, Aug. 2021, pp. 1559–1575.

[44] M. Hocine, N. Balakrishnan, T. Colton, B. Everitt, W. Piegorsch, F. Ruggeri, and J. Teugels, "Wiley statsref: statistics reference online," 2014.

[45] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever, "Language models are unsupervised multitask learners," 2019.

[46] F. Wilcoxon, "Individual comparisons by ranking methods," *Biometrics Bulletin*, vol. 1, no. 6, pp. 80–83, 1945. [Online]. Available: http://www.jstor.org/stable/3001968

[47] K. Kurita, P. Michel, and G. Neubig, "Weight poisoning attacks on pretrained models," in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics.* Online: Association for Computational Linguistics, Jul. 2020, pp. 2793–2806.

[48] W. Yang, L. Li, Z. Zhang, X. Ren, X. Sun, and B. He, "Be careful about poisoned word embeddings: Exploring the vulnerability of the embedding layers in NLP models," in *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies.* Online: Association for Computational Linguistics, Jun. 2021, pp. 2048–2058.

[49] C.-Y. Lin and F. J. Och, "ORANGE: a method for evaluating automatic evaluation metrics for machine translation," in *COLING 2004: Proceedings of the 20th International Conference on Computational Linguistics.* Geneva, Switzerland: COLING, aug 23–aug 27 2004, pp. 501–507.

[50] Q. Gong, J. Zhang, Y. Chen, Q. Li, Y. Xiao, X. Wang, and P. Hui, "Detecting malicious accounts in online developer communities using deep learning," in *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, ser. CIKM '19. New York, NY, USA: Association for Computing Machinery, 2019, p. 1251–1260. [Online]. Available: https://doi.org/10.1145/3357384.3357971

[51] High-Level Expert Group on AI, "Ethics guidelines for trustworthy ai," European Commission, Brussels, Report, Apr. 2019. [Online]. Available: https://ec.europa.eu/digital-single-market/en/news/ethics-guidelines-trustworthy-ai

[52] M. H. Asyrofi, Z. Yang, I. N. B. Yusuf, H. J. Kang, F. Thung, and D. Lo, "Biasfinder: Metamorphic test generation to uncover bias for sentiment analysis systems," *IEEE Transactions on Software Engineering*, pp. 1–1, 2021.

[53] Z. Yang, M. H. Asyrofi, and D. Lo, "Biasrv: Uncovering biased sentiment predictions at runtime," in *Proceedings of the 29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, ser. ESEC/FSE 2021. New York, NY, USA: Association for Computing Machinery, 2021, p. 1540–1544. [Online]. Available: https://doi.org/10.1145/3468264.3473117

[54] Z. Yang, H. Jain, J. Shi, M. H. Asyrofi, and D. Lo, "Biasheal: On-the-fly black-box healing of bias in sentiment analysis systems," in *2021 IEEE International Conference on Software Maintenance and Evolution (ICSME)*, 2021, pp. 644–648.

[55] M. H. Asyrofi, Z. Yang, and D. Lo, "Crossasr++: A modular differential testing framework for automatic speech recognition," in *Proceedings of the 29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, ser. ESEC/FSE 2021. New York, NY, USA: Association for Computing Machinery, 2021, p. 1575–1579. [Online]. Available: https://doi.org/10.1145/3468264.3473124

[56] M. H. Asyrofi, Z. Yang, J. Shi, C. W. Quan, and D. Lo, "Can differential testing improve automatic speech recognition systems?" in *2021 IEEE International Conference on Software Maintenance and Evolution (ICSME)*, 2021, pp. 674–678.

[57] C. Gong, Z. Yang, Y. Bai, J. Shi, A. Sinha, B. Xu, D. Lo, X. Hou, and G. Fan, "Curiosity-driven and victim-aware adversarial policies," in *Proceedings of the 38th Annual Computer Security Applications Conference*, ser. ACSAC '22. New York, NY, USA: Association for Computing Machinery, 2022, p. 186–200. [Online]. Available: https://doi.org/10.1145/3564625.3564636

[58] Z. Yang, J. Shi, M. H. Asyrofi, and D. Lo, "Revisiting neuron coverage metrics and quality of deep neural networks," in *2022 IEEE International Conference on Software Analysis, Evolution and Reengineering (SANER).* IEEE Computer Society, 2022.

[59] X. Xie, L. Ma, F. Juefei-Xu, M. Xue, H. Chen, Y. Liu, J. Zhao, B. Li, J. Yin, and S. See, "Deephunter: A coverage-guided fuzz testing framework for deep neural networks," in *Proceedings of the 28th ACM SIGSOFT International Symposium on Software Testing and Analysis*, ser. ISSTA 2019. New York, NY, USA: Association for Computing Machinery, 2019, p. 146–157. [Online]. Available: https://doi.org/10.1145/3293882.3330579

[60] J. M. Zhang, M. Harman, L. Ma, and Y. Liu, "Machine learning testing: Survey, landscapes and horizons," *IEEE Transactions on Software Engineering*, vol. 48, no. 1, pp. 1–36, 2022.

[61] L. Applis, A. Panichella, and A. van Deursen, "Assessing robustness of ml-based program analysis tools using metamorphic program transformations," in *2021 36th IEEE/ACM International Conference on Automated Software Engineering (ASE)*, 2021, pp. 1377–1381.

[62] I. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and harnessing adversarial examples," in *International Conference on Learning Representations*, 2015. [Online]. Available: http://arxiv.org/abs/1412.6572

[63] M. Allamanis, M. Brockschmidt, and M. Khademi, "Learning to represent programs with graphs," in *International Conference on Learning Representations*, 2018.

[64] M. Brockschmidt, M. Allamanis, A. L. Gaunt, and O. Polozov, "Generative code modeling with graphs," in *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*, 2019.

[65] J. Henkel, G. Ramakrishnan, Z. Wang, A. Albarghouthi, S. Jha, and T. Reps, "Semantic robustness of models of source code," in *2022 IEEE International Conference on Software Analysis, Evolution and Reengineering (SANER)*, 2022, pp. 526–537.

[66] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, "Towards deep learning models resistant to adversarial attacks," in *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*, 2018.

[67] P. T. Nguyen, C. Di Sipio, J. Di Rocco, M. Di Penta, and D. Di Ruscio, "Adversarial attacks to api recommender systems: Time to wake up and smell the coffee?" in *2021 36th IEEE/ACM International Conference on Automated Software Engineering (ASE)*, 2021, pp. 253–265.

[68] Z. Sun, X. Du, F. Song, M. Ni, and L. Li, "Coprotector: Protect open-source code against unauthorized training usage with data poisoning," in *Proceedings of the ACM Web Conference 2022*, ser. WWW '22. New York, NY, USA: Association for Computing Machinery, 2022, p. 652–660.

[69] G. Severi, J. Meyer, S. Coull, and A. Oprea, "Explanation-Guided backdoor poisoning attacks against malware classifiers," in *30th USENIX Security Symposium (USENIX Security 21).* USENIX Association, Aug. 2021, pp. 1487–

1504. [Online]. Available: https://www.usenix.org/conference/usenixsecurity21/presentation/severi

[70] T. Gu, K. Liu, B. Dolan-Gavitt, and S. Garg, "Badnets: Evaluating backdooring attacks on deep neural networks," *IEEE Access*, vol. 7, pp. 47 230–47 244, 2019.

[71] X. Chen, C. Liu, B. Li, K. Lu, and D. Song, "Targeted backdoor attacks on deep learning systems using data poisoning," 2017.

[72] K. Doan, Y. Lao, and P. Li, "Backdoor attack with imperceptible input and latent modification," in *Advances in Neural Information Processing Systems*, M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, Eds., vol. 34. Curran Associates, Inc., 2021, pp. 18 944–18 957.

[73] H. Zhong, C. Liao, A. Squicciarini, S. Zhu, and D. Miller, "Backdoor embedding in convolutional neural network models via invisible perturbation," ser. CODASPY 2020 - Proceedings of the 10th ACM Conference on Data and Application Security and Privacy. Association for Computing Machinery, Inc, Mar. 2020, pp. 97–108.

[74] X. Xu, Q. Wang, H. Li, N. Borisov, C. A. Gunter, and B. Li, "Detecting ai trojans using meta neural analysis," in *2021 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2021, pp. 103–120.

[75] D. Tang, X. Wang, H. Tang, and K. Zhang, "Demon in the variant: Statistical analysis of DNNs for robust backdoor contamination detection," in *30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association, Aug. 2021, pp. 1541–1558.

[76] J. Jia, X. Cao, and N. Z. Gong, "Intrinsic certified robustness of bagging against data poisoning attacks," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 9, 2021, pp. 7961–7969.

[77] B. Wang, Y. Yao, S. Shan, H. Li, B. Viswanath, H. Zheng, and B. Y. Zhao, "Neural cleanse: Identifying and mitigating backdoor attacks in neural networks," in *2019 IEEE Symposium on Security and Privacy (SP)*, 2019, pp. 707–723.

[78] C. Gong, Z. Yang, Y. Bai, J. He, J. Shi, A. Sinha, B. Xu, X. Hou, G. Fan, and D. Lo, "Mind your data! hiding backdoors in offline reinforcement learning datasets," 2022. [Online]. Available: https://arxiv.org/abs/2210.04688

[79] Y. Liu, S. Ma, Y. Aafer, W.-C. Lee, J. Zhai, W. Wang, and X. Zhang, "Trojaning attack on neural networks," in *25th Annual Network and Distributed System Security Symposium, NDSS 2018, San Diego, California, USA, February 18-221, 2018*. The Internet Society, 2018.

[80] J. Dai, C. Chen, and Y. Guo, "A backdoor attack against lstm-based text classification systems," *CoRR*, vol. abs/1905.12457, 2019.

[81] Y. Wan, S. Zhang, H. Zhang, Y. Sui, G. Xu, D. Yao, H. Jin, and L. Sun, "You see what i want you to see: poisoning vulnerabilities in neural code search," in *Proceedings of the 30th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, 2022, pp. 1233–1245.

[82] Y. Sui, X. Cheng, G. Zhang, and H. Wang, "Flow2vec: Value-flow-based precise code embedding," *Proceedings of the ACM on Programming Languages*, vol. 4, no. OOPSLA, pp. 1–27, 2020.

[83] X. Cheng, G. Zhang, H. Wang, and Y. Sui, "Path-sensitive code embedding via contrastive learning for software vulnerability detection," in *Proceedings of the 31st ACM SIGSOFT International Symposium on Software Testing and Analysis*, 2022, pp. 519–531.

[84] X. Cheng, X. Nie, N. Li, H. Wang, Z. Zheng, and Y. Sui, "How about bug-triggering paths?-understanding and characterizing learning-based vulnerability detectors," *IEEE Transactions on Dependable and Secure Computing*, 2022.